Some recent developments in stochastic volatility modelling

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

Elisa Nicolato
Department of Mathematical Sciences,
University of Aarhus, Ny Munkegade, DK-8000 Aarhus C, Denmark
`elisa@imf.au.dk`

Neil Shephard
Nuffield College, Oxford OX1 1NF, UK
`neil.shephard@nuf.ox.ac.uk`
`www.LevyProcess.org`

December 7, 2001

Abstract

This paper reviews and puts in context some of our recent work on stochastic volatility modelling for financial economics. Here our main focus is on: (i) the relationship between subordination and stochastic volatility, (ii) OU based volatility models, (iii) exact option pricing, (iv) realised power variation and realised variance, (v) building multivariate models.

Contents

1. Introduction

This paper reviews and puts in context our recent work on stochastic volatility (SV). Some of the ideas are based on general SV models, while other pieces focus on the Lévy process based Ornstein-Uhlenbeck (OU) driven volatility models we have been developing in the last few years.

The points we will make about general SV models will be based around two issues. First we will give a discussion of the relationship between subordination and stochastic volatility. We feel that having these connections spelt out clarifies a number of issues which arise in the modern modelling of univariate and multivariate processes. Second, and more originally, we review some of our work on "realised power variation", defined as the sum of powers of absolute high frequency returns. This statistic generalises the well known realised volatility or variance concept used in financial econometrics. Our analysis provides the asymptotic distribution of realised power variation as the number of high frequency observations goes off to infinity.

Our recent work on OU based SV models has focused on two issues. The first is the development of exact option pricing formulae based on these models. The work has provided the first example of such exact results outside the CIR based models which have dominated this literature. Second we have made considerable progress on various multivariate modelling issues using OU based volatility models. The advances will be discussed here.

The structure of this paper is as follows. In Section 2, we talk about the relationship between SV models and time change models. Section 3 will study derivatives written on OU based SV models. Section 4 will review the work on realised volatility, while Section 5 looks at some recent results on the new idea of realised power variation. Section 6 discusses various multivariate SV models and their uses and Section 7 collects short discussions of a number of interesting other issues. Finally, Section 8 concludes.

2. SV models and chronometers

2.1. SDE and time change models

There are several ways of introducing stochastic volatility into the Samuelson-Black-Scholes framework for the log-price of an asset

$$y^*(t) = \beta t + \sigma w(t), \tag{2.1}$$

where w is Brownian motion and β (the drift) and σ (the volatility) are parameters.

One approach is by changing clock time t to a random time τ^*, which is any non-decreasing process with $\tau^*(0) = 0$, so that (2.1) becomes

$$y^*(t) = \beta\tau^*(t) + w(\tau^*(t)),$$

where we have dropped σ as this may be incorporated in $\tau^*(t)$. Here τ^* is called a *chronometer* (in our work we reserve the name subordinator for a Lévy process with non-negative increments). It is then natural to add an extra drift term so as to obtain a process of the form

$$y^*(t) = \mu t + b^\beta(\tau^*(t)), \tag{2.2}$$

where b^β denotes a Brownian motion with drift β, i.e.

$$b^\beta(t) = \beta t + b(t).$$

2

If τ^* has continuous sample paths then so will y^*, otherwise both will have jumps.

Another approach changes σ in (2.1) into a positive process $\sigma(t)$ such that

$$\int_0^t \sigma^2(s)\mathrm{d}s < +\infty \quad \text{a.s for every } t > 0$$

and reformulates (2.1) into

$$y^*(t) = \mu t + \beta \tau^*(t) + \int \sqrt{\tau(s)}\mathrm{d}w(s), \tag{2.3}$$

where we have written $\tau(t)$ for the *variance process* $\sigma^2(t)$ and have introduced a drift term, in analogy with (2.2), and where

$$\tau^*(t) = \int_0^t \tau(s)\mathrm{d}s, \tag{2.4}$$

i.e. τ^* is the *integrated variance*. It plays a key role. It is often assumed that the processes w and τ are independent, and we shall do so here. One implication of this model is that log-prices always have continuous sample paths.

Noting that (2.3) is the solution to the stochastic differential equation

$$\mathrm{d}y^*(t) = \mu \mathrm{d}t + \beta \tau(t)\mathrm{d}t + \sqrt{\tau(t)}\mathrm{d}w(t),$$

we shall refer to the above-mentioned two approaches as the SDE approach and the time-change approach, respectively

We may, in fact, rephrase the model (2.3) in the form (2.2). More specifically, given (2.3) there exists a Brownian motion b such that y^*, τ and b satisfy (2.2) and such that b and τ are independent. Conversely, if τ^* in (2.2) is of the form (2.4) then there exists a Brownian motion w such that y^*, τ and w satisfy (2.3) and such that w and τ are independent (cf. Barndorff-Nielsen and Shephard (2001a)).

Not all chronometers τ^* are of the form (2.4) and in that sense (2.2) is more general than (2.3). In particular, if τ^* is a *subordinator*, i.e. a non-decreasing Lévy process, then (2.2) cannot be represented as in (2.3). An example in case is the *NIG* (normal inverse Gaussian) Lévy process which comes about by (2.2) through choosing τ^* as the *IG* (inverse Gaussian) subordinator, i.e. the Lévy process $z(t)$ for which $z(1)$ follows the inverse Gaussian law $IG(\delta, \gamma)$ with probability density

$$\frac{\delta}{\sqrt{2\pi}} e^{\delta\gamma} x^{-3/2} \exp\left\{-\frac{1}{2}(\delta^2 x^{-1} + \gamma^2 x)\right\}. \tag{2.5}$$

A more general case of this is the generalised hyperbolic Lévy process, which takes τ^* as the *GIG* (generalised inverse Gaussian) subordinator.

From the viewpoint of detailed modelling, there may be some appeal in combining the two approaches, by considering log price processes $y^*(t) = y_0^*(t) + y_1^*(t)$ with y_0^* and y_1^* independent and respectively of SDE and subordination type (e.g. SV plus *NIG* Lévy process). The resulting processes have some similarities to the jump-diffusion processes discussed by, for example, Duffie, Pan, and Singleton (2000). The subordination approach may also be broadened to random time changes of Lévy processes in general, i.e. $y^*(t) = \mu t + z(\tau^*(t))$ where z is a Lévy process. This has been discussed recently by Carr, Geman, Madan, and Yor (2001). We will return to these issues in Subsection 7.3.

Two modifications of the SDE approach will be considered in Sections 3, 4 and 5. In the first, a term is added to (a particular subclass of) (2.3) in order to model the so-called leverage effect and in the second the term $\mu t + \beta \tau^*(t)$ is changed to a more general type of process $\alpha(t)$. Before dealing with these issues, we will now go on to discuss the question of choice of volatility process and chronometer.

2.2. Volatility and integrated variance processes

Main considerations in the choice of volatility process τ are: (i) The distributions of increments of $y^*(t)$ should have the kind of shapes typically observed on the financial markets. This means they should be fat tailed, as Figure 2.1 illustrates which plots the log of non-parametric density estimators of daily exchange rates, together with the corresponding fit of generalised hyperbolic distributions to the same data. These plots bring out the well known feature that returns have log-densities whose tails decay not faster than linearly. (ii) The timewise correlation structure of the increments of y^* should show the kind of quasi long range behaviour that is also typically observed. (Note that irrespective of the choice of τ, the form (2.3) accommodates possible asymmetry of the distributions, by the parameter β.) This can be illustrated by using realised variances, sums of intra-day squared returns, which is discussed in Section 4 of this paper. SV models can reproduce both of these features.
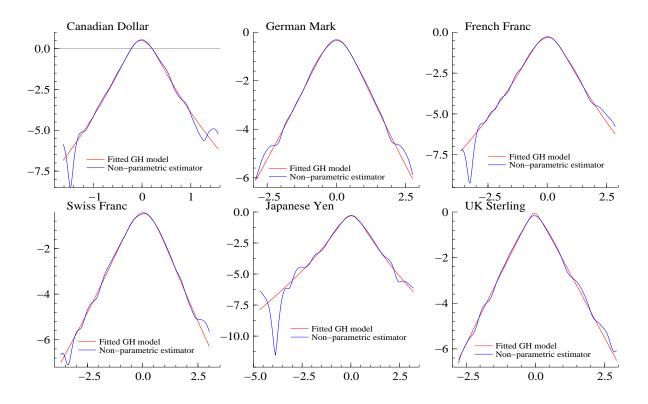

Figure 2.1: *Log of the estimates of the unconditional density of the daily returns for six exchange rates against the US Dollar. 26 July 1985 to 28 July 2000. Also plotted is the log-density for the ML fit of the five parameter generalised hyperbolic distribution.*

In the financial economics literature almost all volatility models are the solution to non-linear SDEs driven by Brownian motion. The non-linearity is essential in order that the volatility does not go negative. The most common is the constant elasticity of variance (CEV) process

$$\mathrm{d}\sigma^2(t) = -\lambda \left\{ \sigma^2(t) - \xi \right\} \mathrm{d}t + \omega\sigma(t)^\eta \mathrm{d}b(\lambda t), \qquad \eta \in [1, 2],$$

where $b(t)$ is standard Brownian motion. Of course the special case of $\eta = 1$ delivers the square root process, while when $\eta = 2$ we have Nelson's GARCH diffusion. These models have been heavily favoured by Meddahi and Renault (2002) in this context.

In some recent work we have taken a radically different route. We construct non-negative OU (Ornstein-Uhlenbeck type) processes by replacing Browian motion with non-negative Lévy

processes (subordinators). This model class generalises by the use of superposition (ie. the addition of independent non-Gaussian OU processes). The linearity of these models allows us to carry out many analytic calculations (e.g. analytic option pricing) which are not possible with most non-linear diffusions.

Non-Gaussian OU processes are constructed in the following way. Let D be a distribution on the positive halfline. If D is selfdecomposable[1], and only then, there exists a subordinator z such that whatever the value of $\lambda > 0$ the SDE

$$d\tau(t) = -\lambda\tau(t)dt + dz(\lambda t), \qquad (2.6)$$

has a stationary solution. In this case, the law of $\tau(t)$ does not depend on λ, and we refer to the process τ as a D-OU process. Notice $\tau(t)$ is kept from being negative by being driven by $z(t)$, a process with non-negative increments. A simulated example of the paths that the $\tau(t)$ and $z(\lambda t)$ processes follow is given in Figure 2.2. This shows the case where D is a $\Gamma(\nu, \alpha)$. There are many other interesting choices for D. In particular generalised inverse Gaussian and its important special case the inverse Gaussian.
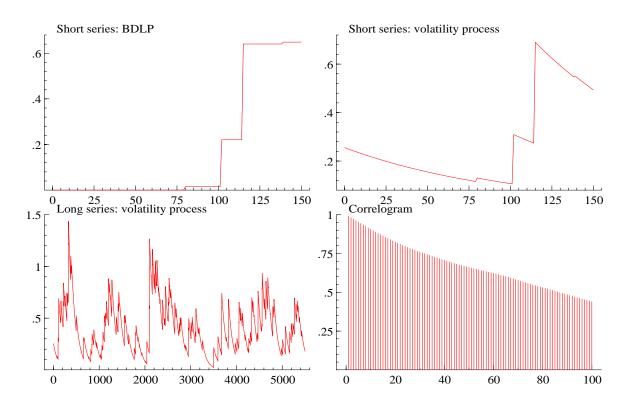

Figure 2.2: OU process with $\Gamma(\nu, \alpha)$ marginals. *Throughout, $\nu = 3$, $\alpha = 8.5$, $\lambda = 0.01$ and $\Delta = 1$. Top left: plot of $z(\lambda n\Delta)$ against n. Top right: plot of $\tau(n\Delta)$ against n. Same graph but for longer series in bottom left. Bottom right: as a numerical check we also present the empirical autocorrelation function for $\tau(n\Delta)$.*

The solution of (2.6) may be expressed as

$$\tau(t) = e^{-\lambda t}\tau(0) + e^{-\lambda t}\int_0^t e^{\lambda s}dz(\lambda s). \qquad (2.7)$$

---

[1] A probability distribution D on $\mathbf{R}$ is *selfdecomposable* if and only if the characteristic function ϕ of D satisfies $\phi(\zeta) = \phi(c\zeta)\phi_c(\zeta)$ for all $\zeta \in \mathbf{R}$ and all $c \in (0, 1)$ and for some family of characteristic functions $\{\phi_c : c \in (0, 1)\}$.

An immediate consequence of (2.7) is that setting

$$\epsilon(s,t) = \lambda^{-1}\left\{1 - e^{-\lambda(t-s)}\right\} \tag{2.8}$$

the integrated variance process $\tau^*(t)$ takes the simple form

$$\tau^*(t) = \epsilon(0,t)\tau(0) + \int_0^t \epsilon(s,t)\mathrm{d}z(\lambda s). \tag{2.9}$$

Addition of independent OU processes, with different λ but such that the law of the one-dimensional marginals of the resulting process is D, is referred to as a D-supOU process. An alternative parametric modelling approach is to assume that $z(1)$ follows a distribution D, then we say that the corresponding OU process is of type OU-D while the corresponding superposition based model is written as supOU-D.

One of the main advantages of the OU processes is that we are able to derive the conditional cumulant function of $\tau^*(t)|\tau(0)$ as well as the marginal cumulant function of $\tau^*(t)$, in terms of the cumulant function of either D or $z(1)$. Moreover, there is a simple relation between the two latter functions; specifically, if we denote these by $\acute{\kappa}$ and κ, respectively, then (see Barndorff-Nielsen and Shephard (2001a), Barndorff-Nielsen and Shephard (2001b))

$$\acute{\kappa}(\zeta) = \int_0^\infty \kappa(e^{-s}\zeta)\mathrm{d}s \quad \text{and} \quad \kappa(\zeta) = \zeta\acute{\kappa}'(\zeta). \tag{2.10}$$

The NIG law gives in many cases a consistently good fit to returns. If the chronometer $\tau^*(t)$ in (2.2) or (2.3) followed the $IG(\delta,\gamma)$ law then $y^*(t)$ would have the $NIG(\alpha,\beta,t\mu,t\delta)$ distribution (where $\alpha = \sqrt{\beta^2+\gamma^2})^2$. The IG-OU and, more generally, the IG-supOU variance processes τ, that have had a central place in our work, have the property that $\tau(t)$ has law $IG(\delta,\gamma)$ (for all t, since τ is stationary). Then the law of $\tau^*(t)$ is not exactly IG for any t but it is shown in Barndorff-Nielsen and Shephard (2001c) that it is in fact close to being IG.

The analogous conclusion holds more generally for the TS (tempered stable) and normal tempered stable laws and associated stochastic volatility processes, but it is not the case if instead of a TS law one uses, for instance, the log normal. This is also shown in Barndorff-Nielsen and Shephard (2001c).

The TS distributions are obtained from the positive ω-stable laws ($0 < \omega < 1$) by exponential tilting. A further power tilting leads to the modified stable laws, of which the GIG laws constitute the special case for which $\omega = \frac{1}{2}$. If u is modified stable and independent of a standard normal variable ε then, by definition, $\mu + \beta u + \sqrt{u}\varepsilon$ (compare to (2.3)) follows a normal modified stable law; the generalised hyperbolic laws correspond to $\omega = \frac{1}{2}$. The modified stable and normal modified stable laws are introduced and studied in Barndorff-Nielsen and Shephard (2001d).

3. Derivative assets analysis

In this section we investigate the class of OU based stochastic volatility models from the viewpoint of derivative asset analysis. In particular we assume that on a given filtered probability space $(\Omega, \mathcal{F}_t, P)$ with P denoting the *physical* probability measure, the log-price process has dynamics given by

$$\mathrm{d}y^*(t) = (\mu + \beta\,\tau(t))\mathrm{d}t + \sqrt{\tau}(t)\mathrm{d}w(t) + \rho\mathrm{d}z(\lambda t) \tag{3.1}$$

---

[2]Four parameters is the minimum number needed to accomodate location, spread, asymmetry and tail heavyness.

where the variance process[3] $\tau(t)$ is described by an Ornstein-Uhlenbeck type process as in (2.6), z is the subordinator (independent of the Brownian motion w) driving $\tau(t)$ and ρ is a non positive real parameter which accounts for possible leverage effects (e.g. Black (1976) and Nelson (1991)). For simplicity, we also assume that z has no deterministic drift and that its Lévy measure is absolutely continuous with respect to the Lebesgue measure. Hence the cumulant generating function of $z(1)$ can be written in the form

$$k(\theta) = \log \mathrm{E}[\mathrm{e}^{\theta z(1))}] = \int_0^{+\infty} \left(\mathrm{e}^{\theta x} - 1 \right) u(x)\mathrm{d}x$$

where $u(x)$ denotes the density of the Lévy measure. We also assume that there is a constant continuously compounded interest rate r.

3.1. Sets of equivalent martingale measures

Consider a derivative product written on an asset with log-price given by (3.1), maturity T and payoff $h(y^*(T))$ and let C_t denote its arbitrage-free price at time $t \leq T$. According to the fundamental theorem of asset pricing, C_t is given by

$$C_t = \mathrm{E}^Q[\mathrm{e}^{-r(T-t)}h(y^*(T))|\mathcal{F}_t] \tag{3.2}$$

where the expectation is taken with respect to an equivalent martingale measure (EMM) or risk neutral measure Q, i.e. a probability measure equivalent to P under which the discounted price process $\exp\left(-rt + y^*(t)\right)$ evolves as a martingale. As is typical for financial market models allowing for stochastic volatility, the model in (3.1) is arbitrage free but incomplete, meaning that there exists an infinite number of EMMs. The structure of a general EMM and some relevant subsets of EMMs under the model (3.1) are studied in Nicolato and Venardos (2001). A similar analysis has been independently developed by Hubalek and Tompkins (2000b).

Of special interest is the set $\mathcal{M}'$ consisting of those equivalent martingale measures Q under which $z(\lambda t)$ remains a Lévy process. Let $Q \in \mathcal{M}'$ and let $u^Q(\mathrm{d}x)$ denote the Lévy measure of z under Q. Then

$$u^Q(\mathrm{d}x) = \varphi(x)u(x)\mathrm{d}x, \tag{3.3}$$

where the density $\varphi(x)$ is a strictly positive function such that

$$\int_{\mathbb{R}_+} \left(\sqrt{\varphi(x)} - 1 \right)^2 u(x)\mathrm{d}x < +\infty. \tag{3.4}$$

The dynamics of the log-price process $y^*(t)$ under Q are given by

$$\mathrm{d}y^*(t) = \left\{ r - \lambda k^Q(\rho) - \frac{1}{2}\tau(t) \right\} \mathrm{d}t + \sqrt{\tau(t)}\mathrm{d}w^Q(t) + \rho\mathrm{d}z(\lambda t),$$
$$\mathrm{d}\tau(t) = -\lambda\tau(t)\mathrm{d}t + \mathrm{d}z(\lambda t), \tag{3.5}$$

where $w^Q(t)$ is a Q-Brownian motion, $w^Q(t)$ and $z(\lambda t)$ are independent and $k^Q(\theta)$ is the cumulant function of $z(1)$ under Q, i.e.

$$k^Q(\theta) = \int_0^{+\infty} \left(\mathrm{e}^{\theta x} - 1 \right) \varphi(x)u(x)\mathrm{d}x.$$

Conversely, for any positive function $\varphi(x)$ satisfying (3.4) there exists $Q \in \mathcal{M}'$ under which z has Lévy measure $u^Q(\mathrm{d}x)$ given by (3.3).

---

[3]Notice that when $z(1)$ has finite second moment, the instantaneous variance of log-returns is given by $(\tau(t) + \rho^2\lambda\mathrm{Var}(z(1)))\mathrm{d}t$.

Process and $k^Q(\theta)$	Integral $\int_t^T k^Q[f(s,\theta)]\mathrm{d}s$ setting $f_1 = \theta\rho + \frac{1}{2}(\theta^2 - \theta)\left(1 - e^{-\lambda(T-t)}\right)$ $f_2 = \theta\rho + \frac{1}{2}(\theta^2 - \theta)$
$\mathrm{IG}(\gamma,\delta)$-OU $\dfrac{\delta\theta}{\sqrt{\gamma^2 - 2\theta}}$	$\dfrac{\delta}{\lambda}\left(\sqrt{\gamma^2 - 2f_1} - \sqrt{\gamma^2 - 2\theta\,\rho}\right) + \dfrac{2\delta f_2}{\lambda\sqrt{2f_2 - \gamma^2}}\times$ $\left(\arctan\left(\sqrt{\dfrac{\gamma^2 - 2\theta\rho}{2f_2 - \gamma^2}}\right) - \arctan\left(\sqrt{\dfrac{\gamma^2 - 2f_1}{2f_2 - \gamma^2}}\right)\right)$
$\Gamma(\alpha,\nu)$-OU $\dfrac{\nu\theta}{\alpha - \theta}$	$\left(\alpha\log\left[\dfrac{\alpha - f_1}{\alpha - \theta\rho}\right] + f_2\lambda(T-t)\right) \times \dfrac{\nu}{\lambda(\alpha - f_2)}$

Table 3.1: The closed form solution for the integral appearing in the Laplace transform of $y^*(T)|y^*(t), \tau(t)$ for the IG-OU and Γ-OU models.

In other words, in a risk neutral world governed by a $Q \in \mathcal{M}'$, the log price process and its volatility are again described by a model of the type (3.1), albeit with different parameters and possibly different invariant law for the volatility process.

In subsection 3.3 it will be argued that it is sufficient to consider only equivalent martingale measures in $\mathcal{M}'$. From the viewpoint of analytical tractability, this choice has clear advantages. For example, the Markov property of the pair $(y^*(t), \tau(t))$ implies the price at time t of the claim $h(y^*(T))$ depends only on the value of the underlying and of the volatility process at time t and not on the whole history up to time t. Moreover, for processes as in (3.5) (or, more generally as in (3.1)) the conditional cumulant function of the log-price $y^*(T)$ given $y^*(t), \tau(t)$ has a simple expression. More precisely, when the log price has risk neutral dynamics given by (3.5), the conditional cumulant function

$$k_t(\theta) = \log \mathrm{E}^Q\left[\exp(\theta y^*(T))|\mathcal{F}_t\right] = \log \mathrm{E}^Q\left[\exp(\theta y^*(T))|y^*(t), \tau(t)\right]$$

is given by

$$k_t(\theta) = \left[\theta\left(y^*(t) + \left(r - k^Q(\rho)\right)(T-t)\right) + (f(t,\theta) - \rho\,\theta)\,\tau(t) + \lambda\int_t^T k^Q(f(s,\theta))\mathrm{d}s\right], \quad (3.6)$$

where

$$f(s,\theta) = \rho\theta + \frac{1}{2}\left(\theta^2 - \theta\right)\epsilon(s,T) \tag{3.7}$$

and $\epsilon(s,t)$ is given in (2.8). For many concrete specifications of the volatility process, the expression in (3.6) can be computed in terms of elementary functions. The results for the IG-OU and the Γ-OU cases are given in Table 3.1. As we shall see below, the availability of the conditional cumulant function in closed form can be exploited for computing the pricing formula (3.2) rapidly.

3.2. Option pricing

In this section we assume that the equivalent martingale measure Q used for pricing the claim $h(y^*(T))$ has been selected in the subset $\mathcal{M}'$ and we investigate how to evaluate the expected value in formula (3.2).

The first approach we discuss is along the lines of Hull and White (1987). The key observation is that for the model (3.5) the conditional law of the log-price $y^*(T)$ given $(y^*(t), \tau(t))$ can be represented as a normal variance-mean mixture. In fact, defining the *effective log price* and the *effective variance* over the time interval $[t, T]$ as

$$y^*_{eff} = y^*(t) + \rho \left(z(\lambda T) - z(\lambda t) \right) - \lambda k^Q(\rho) \left(T - t \right), \tag{3.8}$$

$$\tau^*_{eff} = \frac{\tau^*(T) - \tau^*(t)}{T - t} = (T - t)^{-1} \left\{ \epsilon(t, T)\tau(t) + \int_t^T \epsilon(s, T) \mathrm{d}z(\lambda s) \right\} \tag{3.9}$$

with $\epsilon(t, T)$ given by (2.8), and integrating (3.5) over $[t, T]$ it is easy to see that

$$y^*(T)|(y^*(t), \tau(t)) \stackrel{\mathcal{L}}{=} r + y^*_{eff} - \frac{1}{2}\tau^*_{eff}(T - t) + \sqrt{(T - t)\tau^*_{eff}}\, \mathcal{N} \tag{3.10}$$

where $\mathcal{N}$ is a standard normal random variable independent of $(y^*_{eff}, \tau^*_{eff})$.

Let $C_t^{\mathrm{BS}}(y, \tau)$ denote the Black-Scholes price for the contract $h(y^*(T))$ when the current price of the underlying is y and the deterministic variance is τ. Namely, $C_t^{\mathrm{BS}}(y, \tau)$ is the derivative price when the conditional law of the log price at expiry is given by (3.10) where the effective quantities $(y^*_{eff}, \tau^*_{eff})$ are treated as constant and are given by y and τ respectively.

Then computing (3.2) by iterated conditional expectations yields

$$C_t = \mathrm{E}^Q[C_t^{\mathrm{BS}}(y^*_{eff}, \tau^*_{eff})|y^*(t), \tau(t)]. \tag{3.11}$$

If the Black-Scholes price $C_t^{\mathrm{BS}}(y, \tau)$ is known in closed form, like in the case of European call options, expression (3.11) can be evaluated as a sample average across simulations of the pair $(y^*_{eff}, \tau^*_{eff})$. The random terms in the effective quantities involve only integrals of deterministic functions with respect to a Lévy process. Hence $(y^*_{eff}, \tau^*_{eff})$ can be efficiently simulated using the work of Marcus (1987) and Rosinski (1991) on infinite series representations of these type of integrals. A self-contained exposition of the required results is given in Barndorff-Nielsen and Shephard (2001b), whereas recent developments are surveyed in Rosinski (2000).

An alternative and extremely powerful approach to the evaluation of (3.2) is provided by so-called transform based methods. The main idea has been essentially introduced in this context by Heston (1993) and it has been further developed by Carr and Madan (1998) and Raible (1998). It applies whenever the Laplace transform of the conditional log-price at expiry has an analytical expression. As discussed in the previous section, this is the case for the IG-OU and Γ-OU specifications of the variance process $\tau(t)$.

The crux of the matter is that the derivative price C_t can be expressed, under some regularity conditions, as follows

$$C_t = \frac{\mathrm{e}^{-r(T-t)}}{2\pi\mathrm{i}} \int_{c-\mathrm{i}\infty}^{c+\mathrm{i}\infty} \mathrm{e}^{k_t(\theta)} \widehat{h}(\theta) \mathrm{d}\theta, \tag{3.12}$$

where $k_t(\theta)$ is given in (3.6), $\widehat{h}(\theta)$ is the Laplace transform of the pay-off function

$$\widehat{h}(\theta) = \int_{-\infty}^{+\infty} \mathrm{e}^{-\theta x} h(x) \mathrm{d}x$$

and c is a real number belonging to the domains of existence of both k_t and $\widehat{h}$.

When also $\widehat{h}$ can be expressed in terms of elementary functions, as in the case of European call options, the pricing formula (3.12) can be computed numerically using automatic integration software packages or more sophisticated procedures as described in Carr and Madan (1998) and Raible (1998).

Nicolato and Venardos (2001) have reanalysed the 87 European call options on the S&P500 index observed in November 1993 discussed in Duffie, Pan, and Singleton (2000). The dataset is based on 6 maturities. The fit to the volatility smiles of the OU-based SV option pricing model is illustrated in Figure 3.1 using a Γ-OU case. In the OU framework, the empirical investigation of more flexible pricing models based on superpositions of OU processes is a topic of ongoing research.
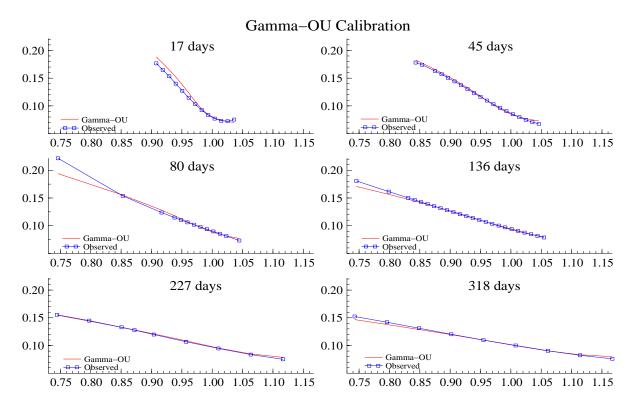

Figure 3.1: *Empirical volatility smiles at various maturities plus fitted values using the Γ-OU option pricing model.*

3.3. Range of prices

When dealing with incomplete market models, it is important to determine the set of prices obtained by formula (3.2) as the pricing measure Q varies over the whole set of equivalent martingale measures. In particular, it has been shown by El Karoui and Quenez (1995) and by Kramkov (1996) that the supremum of possible prices corresponds to the minimum initial value of a self-financing strategy which super-replicates the contingent claim.

Eberlein and Jacod (1997) and Jakubénas (1998) determine the range of prices for European call options when the log-price of the underlying is given by a Lévy process. Incomplete markets described by jump diffusion processes are examined by Bellamy and Jeanblanc (2000) while the case of stochastic volatility in the context of diffusion processes is analysed by Frey and Sin (1999) and Cvitanic, Pham, and Touzi (1999).

The set of viable prices of a call option when the underlying is of the form (3.1) is studied in Nicolato and Venardos (2001). As in most of the above mentioned incomplete markets, it turns out that the supremum of possible derivative prices coincides with the obvious upper bound given by the current price of the underlying. Therefore, the minimal strategy to super-hedge the European call is to hold a long position in the underlying asset. On the other hand, the

infimum does not coincide with the trivial lower bound. It is given by the Black-Scholes function evaluated at the current price of the underlying and at the infimum of the effective volatility, i.e.

$$\inf_Q C_t = C_t^{BS}\left\{y_t^*, \frac{\epsilon(t,T)}{T-t}\tau(t)\right\}.$$

As discussed in Sections 3.1 and 3.2, it is particularly convenient to consider only structure preserving equivalent martingale measures. One may ask whether this choice is too restrictive in the sense that the range of prices is narrowed when the pricing measures are let to vary only in $\mathcal{M}'$. Nicolato and Venardos (2001) show that this is not the case. Indeed even smaller subclasses of equivalent martingale measures give rise to the same range of prices.

4. Realised variance

In this and the following sections we study log-price processes of the form

$$y^*(t) = \alpha(t) + \int_0^t \tau^{1/2}(s)\mathrm{d}w(s), \tag{4.1}$$

where w is Brownian motion and where τ - the variance - and α - the mean or '*risk*' process - are assumed to be jointly independent of w. Further, we assume that α is a process of zero quadratic variation. For example we can think of α as a process of bounded variation with continuous paths.

Over an interval of time of length $\Delta > 0$, say a day, returns are defined as

$$y_n = y^*(\Delta n) - y^*((n-1)\Delta), \qquad n = 1, 2, ... \tag{4.2}$$

which implies that whatever the model for τ and α, it follows that

$$y_n|\tau_n, \alpha_n \sim N(\alpha_n, \tau_n),$$

are conditionally independent through time, where we have defined

$$\tau_n = \tau^*(n\Delta) - \tau^*\{(n-1)\Delta\}$$

and $\alpha_n = \alpha(n\Delta) - \alpha((n-1)\Delta)$. Barndorff-Nielsen and Shephard (2002) call τ_n the *actual variance* of the process.

4.1. Returns and realised variance

Suppose our desire is to estimate τ_n. We know that quadratic variation reveals integrated variance

$$[y^*](t) = \tau^*(t),$$

so theoretically we can compute

$$\tau_n = [y^*](n\Delta) - [y^*]((n-1)\Delta),$$

see the concurrent independent work of Barndorff-Nielsen and Shephard (2001a) and Andersen and Bollerslev (1998). However, in practice market microstructure effects (e.g. discreteness of prices, bid/ask bounce) imply the SV models are typically misspecified over short time horizons. Thus in econometrics the above asymptotic concept is replaced by the estimator, called *realised variance* (or realised volatility),

$$[y_M^*]_n^{[2]} = \sum_{j=1}^M y_{j,n}^2,$$

11

where
$$y_{j,n} = y^* \left((n-1)\Delta + \frac{\Delta j}{M} \right) - y^* \left((n-1)\Delta + \frac{\Delta(j-1)}{M} \right),$$
is the j-th intra-Δ return.

Realised variance estimators have a long history in financial economics. Examples include Poterba and Summers (1986), Schwert (1989), Taylor and Xu (1997) and Christensen and Prabhala (1998), Andersen, Bollerslev, Diebold, and Labys (2001) and Andersen, Bollerslev, Diebold, and Ebens (2001). A stimulating discussion of this literature is given by Andersen, Bollerslev, and Diebold (2002). However, until recently no econometric theory had been developed to understand the distribution of realised variance error $[y_M^*]_n^{[2]} - \tau_n$ other than the quadratic variation implication that as $M \to \infty$ so

$$[y_M^*]_n^{[2]} - \tau_n \xrightarrow{p} 0.$$

Barndorff-Nielsen and Shephard (2001e) established a theory of power variation (see further in sections below) which in particular implies that realised variance has a mixed normal asymptotic distribution and, moreover, that

$$\frac{[y_M^*]_n^{[2]} - \tau_n}{\sqrt{\frac{2}{3}\sum_{j=1}^{M} y_{j,n}^4}} \xrightarrow{\mathcal{L}} N(0,1). \tag{4.3}$$

This is interesting for it means that the distribution of $[y_M^*]_n^{[2]} - \tau_n$ can be calculated (approximately) without knowing the form of α or τ. This is an extension of an earlier result given in Barndorff-Nielsen and Shephard (2002) which assumed $\alpha(t) = \mu t + \beta \tau^*(t)$. An asymptotically equivalent form for this result is

$$\frac{\log[y_M^*]_n^{[2]} - \log \tau_n}{\sqrt{\frac{2}{3} \frac{\sum_{j=1}^{M} y_{j,n}^4}{\left(\sum_{j=1}^{M} y_{j,n}^2\right)^2}}} \xrightarrow{\mathcal{L}} N(0,1). \tag{4.4}$$

Simulation evidence in Barndorff-Nielsen and Shephard (2001f) suggests that for small M the asymptotic approximation in (4.4) is more accurate than the one in (4.3).

To illustrate this result we have used the same return data employed by Andersen, Bollerslev, Diebold, and Labys (2001) in their empirical study of the properties of realised variance, although we have made slightly different adjustments to deal with some missing data (in the context of this paper the effect of these differences are tiny, but were made here to be consistent with our other work on this dataset). Full details of this are given in Barndorff-Nielsen and Shephard (2002). The data was kindly supplied to us by the Olsen group in Zurich. This United States Dollar/ German Deutsche Mark series covers the ten year period from 1st December 1986 until 30th November 1996. It records every five minutes the most recent quote to appear on the Reuters screen. Figure 4.1 shows the daily time series of the realised variance for $M = 144$. Here we report the first 50 days of the series, while the realised variance estimator is based on 10 minute returns.

The 95% daily confidence intervals for realised variance are based on the accurate log-based asymptotic result given in equation (4.4). We can see the important widening and closing of the 95% confidence intervals, with the intervals seemingly being very large when the volatility is high.

4.2. Second order properties of realised variance

Although the asymptotic distribution of $[y_M^*]_n^{[2]}$ is potentially helpful, it is also interesting to think about its exact second order properties. Barndorff-Nielsen and Shephard (2002) have
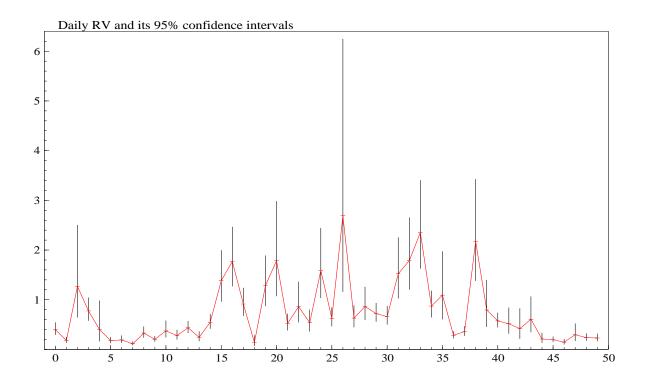
Figure 4.1: *Daily RV,* $[y_M^*]_n^{[2]}$, *drawn against* n *for the first 50 days of the sample. Also drawn as vertical bars are the 95% intervals based on the log transformation. Throughout* $M = 144$.

studied this when $\alpha(t) = 0$. Although this is a limitation, the above asymptotic theory suggests that if we incorrectly set $\alpha(t)$ to zero, the size of the approximation error is of small order for reasonably large values of M.

The second order properties are built up in three stages. The first stage is that we write $u_n = [y_M^*]_n^{[2]} - \tau_n$ so that $[y_M^*]_n^{[2]} = \tau_n + u_n$. It has the feature that

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

where $\varepsilon_{j,n} \stackrel{i.i.d.}{\sim} N(0,1)$ and is independent of

$$\tau_{j,n} = \tau^* \left((n-1)\Delta + \frac{\Delta j}{M} \right) - \tau^* \left((n-1)\Delta + \frac{\Delta(j-1)}{M} \right),$$

the intra-Δ variance. An implication is that u_n is a zero mean, white noise process which is uncorrelated with τ_n.

Second Barndorff-Nielsen and Shephard (2001a) showed that if we write (when they exist) ξ, ω^2 and r, respectively, as the mean, variance and the autocorrelation function of the process $\tau(t)$ then

$$\mathrm{E}(\tau_n) = \xi\Delta, \quad \mathrm{Var}(\tau_n) = 2\omega^2 r^{**}(\Delta) \quad \text{and} \quad \mathrm{Cov}\{\tau_n, \tau_{n+s}\} = \omega^2 \diamondsuit r^{**}(\Delta s), \quad (4.5)$$

where

$$\diamondsuit r^{**}(s) = r^{**}(s+\Delta) - 2r^{**}(s) + r^{**}(s-\Delta), \quad (4.6)$$

and

$$r^*(t) = \int_0^t r(u)\mathrm{d}u \quad \text{and} \quad r^{**}(t) = \int_0^t r^*(u)\mathrm{d}u. \quad (4.7)$$

13

Finally we note

$$
\begin{aligned}
\mathrm{Var}(u_n) &= 2M\mathrm{E}\left\{\tau_{1,n}^2\right\} & (4.8)\\
&= 2M\left\{\mathrm{Var}\left(\tau_{1,n}\right) + \mathrm{E}\left(\tau_{1,n}\right)^2\right\}\\
&= 2M\left\{2\omega^2 r^{**}\left(\Delta M^{-1}\right) + \left(\Delta M^{-1}\xi\right)^2\right\}. & (4.9)
\end{aligned}
$$

This has the feature that

$$
\mathrm{Var}\left(\sqrt{M}u_n\right) \to 2\Delta^2\left(\omega^2 + \xi^2\right) \qquad \text{as} \qquad M \to \infty.
$$

The results imply that $[y_M^*]_n^{[2]}$ is a serially dependent sequence with known autocovariance function, plus an uncorrelated white noise error term with known variance. This is enough to determine the entire second order properties of $[y_M^*]_n^{[2]}$, as well as allowing the linear filtering and smoothing theory reviewed by Whittle (1983) to be used to provide model based estimators of τ_n using the time series of realised volatilities. Further, a Gaussian quasi-likelihood can be constructed for the realised volatilities which can be used to estimate the parameters which index the continuous time SV model. This estimation is fast, free of discretisation errors and seems quite accurate for reasonable values of M such as 12 or more. Details of this, together with empirical and Monte Carlo illustrations, are given in Barndorff-Nielsen and Shephard (2001a). Figure 4.2 gives the main result, which shows the correlogram of $[y_M^*]_n^{[2]}$ using $M = 144$ (i.e. ten minute returns) of the Olsen group's Dollar/DM data based on ten years of data. This shows quite a high correlation which damps down gradually. The Figure shows the implied autocorrelation for the realised variances from a continuous time SV model based on a volatility process constructed by a superposition of J Ornstein-Uhlenbeck processes. The parameters of the models are determined using the above quasi-likelihood. We see that the one factor model is hopeless, two factors produce a good fit, while models with three factors are more compelling.

5. Power variation

In some very recent work Barndorff-Nielsen and Shephard (2001e) have extended the above analysis to allow detailed study of sums of powers of absolute high frequency returns. Let the log-price y^* again follow (4.1). Instead of working with multiple intervals of length Δ, in this section we will follow Barndorff-Nielsen and Shephard (2001e) and think of computing the sums of powers of absolute returns once, over the fixed interval 0 to t. Of course, one could use the result repeatedly over non-overlapping intervals to analyse the previous situation but that will not be our focus here. Instead our attention is on the development of the new asymptotic results reported in Barndorff-Nielsen and Shephard (2001e).

Throughout the following holds, r denotes a positive number and we assume that the processes $\tau = \sigma^2 > 0$ and α are of local bounded variation on $[0, \infty)$. This implies that τ and α are locally bounded Riemann integrable functions and that y^* is a semimartingale with a continuous local martingale component.

Moreover we shall refer to the following conditions on the variance and mean processes:

(V) The variance process $\tau = \sigma^2$ is (pathwise) locally bounded away from 0 and has, moreover, the property

$$
p\text{-}\lim_{\delta \downarrow 0} \delta^{1/2} \sum_{j=1}^{M} |\tau^r(\eta_j) - \tau^r(\xi_j)| = 0 \qquad (5.1)
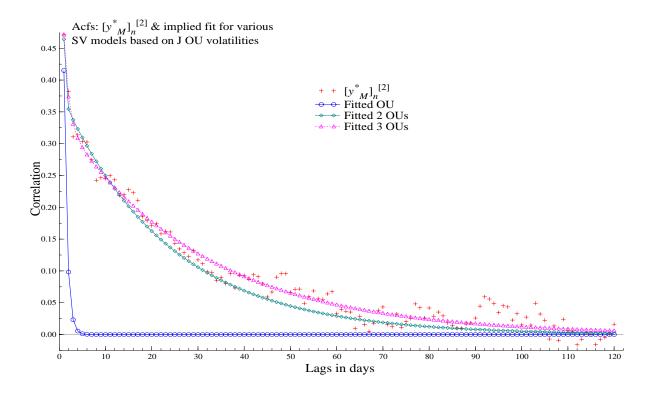$$

14

Figure 4.2: *Results from the fit of the SV model using $M = 144$ (10 minute returns). Autocorrelation function of $[y_M^*]_n^{[2]}$ and the fit of SV model with various superpositions of J Ornstein-Uhlenbeck processes.*

for some r (equivalently for every r) and any ξ_j and η_j such that

$$0 \le \xi_1 \le \eta_1 \le \delta \le \xi_2 \le \eta_2 \le 2\delta \le \cdots \le \xi_j \le \eta_j \le M\delta = t.$$

(M) The mean process α satisfies (pathwise)

$$\overline{\lim_{\delta\downarrow 0}} \max_{1 \le j \le M} \delta^{-1}|\alpha(j\delta) - \alpha((j-1)\delta)| < \infty. \qquad (5.2)$$

These regularity conditions are quite mild and are satisfied in particular if τ is of OU type and if α is an integral of an OU process. Of some special interest are cases where α is of the form

$$\alpha(t) = \int_0^t g(\sigma(s))\mathrm{d}s$$

for g a smooth function. Then regularity of τ will imply regularity of α.

Note that the assumptions allow τ to have, for example, deterministic diurnal effects, jumps, long memory, no unconditional mean or to be non-stationary.

The *realised power variation* is defined, for $r > 0$, as

$$[y_\delta^*]^{[r]}(t) = \sum_{j=1}^M |y_j(t)|^r .$$

where $y_j(\delta) = y^*(j\delta) - y^*((j-1)\delta)$.

The main result of the above mentioned paper is the following theorem.

Theorem 1. *For $\delta \downarrow 0$ and $r \geq \frac{1}{2}$, under conditions (V) and (M),*

$$\mu_r^{-1} \delta^{1-r/2} [y_\delta^*]^{[r]}(t) \xrightarrow{p} \tau^{r/2*}(t) \tag{5.3}$$

and

$$\frac{\mu_r^{-1} \delta^{1-r/2} [y_\delta^*]^{[r]}(t) - \tau^{r/2*}(t)}{\mu_r^{-1} \delta^{1-r/2} \sqrt{\mu_{2r}^{-1} v_r [y_\delta^*]^{[2r]}(t)}} \xrightarrow{\mathcal{L}} N(0,1), \tag{5.4}$$

where $\mu_r = \mathrm{E}\{|u|^r\}$ and $v_r = \mathrm{Var}\{|u|^r\}$, with $u \sim N(0,1)$. $\square$

This result tells us that, for $\delta \downarrow 0$, scaled realised power variation converges in probability to integrated power volatility and follows asymptotically a normal variance mixture distribution with variance distributed as

$$\delta \mu_r^{-2} v_r \tau^{r*}(t)$$

which is consistently estimated by the square of the denominator in (5.4). Hence the limit theory is statistically feasible and does not depend upon knowledge of α or σ^2.

Leading cases are *realised quadratic variation*, usually called realised volatility in the finance and econometrics literature, which results when $r = 2$

$$[y_\delta^*]^{[2]}(t) = \sum_{j=1}^M y_j^2(t),$$

in which case

$$\frac{\sum_{j=1}^M y_j^2(t) - \int_0^t \sigma^2(s)\mathrm{d}s}{\sqrt{\frac{2}{3} \sum_{j=1}^M y_j^4(t)}} \xrightarrow{\mathcal{L}} N(0,1) \tag{5.5}$$

and *realised absolute variation*

$$[y_\delta^*]^{[1]}(t) = \sum_{j=1}^M |y_j(t)|,$$

when

$$\frac{\mu_1^{-1} \sqrt{\delta} \sum_{j=1}^M |y_j(t)| - \int_0^t \sigma(s)\mathrm{d}s}{\sqrt{\delta\left(\mu_1^{-2} - 1\right) \sum_{j=1}^M y_j^2(t)}} \xrightarrow{\mathcal{L}} N(0,1), \tag{5.6}$$

where $\mu_1 = \sqrt{2/\pi}$.

6. Multivariate extensions

When building a multivariate model for financial data, there are some basic considerations that should be born in mind. Obviously, the univariate marginal processes should capture well-known stylized features of univariate financial time series. The multivariate model should also provide a flexible and realistic dependence structure between different assets. For example, a constant correlation matrix of returns is considered economically unsatisfactory. Finally, an analytically tractable continuous time model is clearly advisable, especially when issues like pricing and hedging derivatives written on several assets come into play.

Here we discuss a possible extension of SV models of the type (2.3) to a multivariate setting which has been originally suggested by Barndorff-Nielsen and Shephard (2001a) and further developed by Hubalek and Nicolato (2001). This approach is in the spirit of earlier work in discrete time by Diebold and Nerlove (1989) and Pitt and Shephard (1999). The vector of K log-price processes $y^*(t) = (y_1^*(t), \ldots, y_K^*(t))'$ satisfies

$$\mathrm{d}y^*(t) = (\mu + \Sigma(t)\beta)\mathrm{d}t + \Sigma^{\frac{1}{2}}(t)\mathrm{d}w(t) \tag{6.1}$$

16

where $w(t)$ is standard Brownian motion and $\Sigma(t)$ is a time varying spot covariance matrix. The spot covariance matrix is modeled indirectly via a factor structure

$$\Sigma(t) = I(t) + BJ(t)B' = \begin{bmatrix} I & B \end{bmatrix} \operatorname{diag}(\varphi(t)) \begin{bmatrix} I & B \end{bmatrix}',$$

where $\varphi(t)$ is a vector containing $K + k$ independent Ornstein-Uhlenbeck type processes, thus satisfying

$$\mathrm{d}\varphi_i(t) = -\lambda_i \varphi_i(t)\mathrm{d}t + \mathrm{d}z_i(\lambda_i t),$$

with $z(t) = (z_1(t), \ldots, z_{K+k}(t))'$ being independent Lévy processes which are also assumed to be independent of $w(t)$. The structure of the model means that there are K *idiosyncratic* and k *common factors* which are arranged in the diagonal matrices

$$I(t) = \operatorname{diag}(\varphi_1(t), \ldots, \varphi_K(t)), \quad J(t) = \operatorname{diag}(\varphi_{K+1}(t), \ldots, \varphi_{K+k}(t))$$

and $B \in \mathbb{R}^{K \times k}$ is the *factor loadings matrix*. This multivariate model corresponds to the straightforward economic interpretation that assets in a particular market are driven by market specific factors (common factors), but have a component that is specific to the particular asset under consideration (the corresponding idiosycratic factor).

An attractive feature of this model is that it has a simple subordination interpretation. In particular

$$y^*(t) \overset{\mathcal{L}}{=} \mu t + \left\{ \begin{bmatrix} I & B \end{bmatrix} \operatorname{diag}(\varphi^*(t)) \begin{bmatrix} I & B \end{bmatrix}' \right\} \beta + \begin{bmatrix} I & B \end{bmatrix} w^\dagger(\varphi^*(t)),$$

where $w^\dagger$ is a $K + k$ dimensional vector of standard Brownian motions and φ^* is a $K + k$ dimensional vector chronometer

$$\varphi^*(t) = \int_0^t \varphi(u)\mathrm{d}u.$$

This implies $w^\dagger(\varphi^*(t))$ is a vector of independent SV models. Of course the interesting feature of this representation is that the dimension of the chronometer is higher than the log-price vector.

As in the univariate case, it is possible to introduce in the model (6.1) some leverage type effects. This can be done in several ways. A rather general possibility is described by

$$\mathrm{d}y^*(t) = (\mu + \Sigma(t)\beta)\mathrm{d}t + \Sigma^{1/2}(t)\mathrm{d}w(t) + \rho\mathrm{d}z(\lambda t) \tag{6.2}$$

where $\rho \in \mathbb{R}^{n \times (n+k)}$. The linear dependence of the covariance matrix and of the drift (and of possible leverage terms) on the factors ensures the analytical tractability of models of the type described above. In particular, the cumulant functional

$$K\left\{ 1 \ddagger \sum_{i=1}^K \int f_i(s)\mathrm{d}y_i^*(s) \right\} = \log \mathrm{E}\left[\exp\left\{ \sum_{i=1}^K \int f_i(s)\mathrm{d}y_i^*(s) \right\} \right]$$

with f_i deterministic "test functions" has a simple expression. These results are developed at length in Hubalek and Nicolato (2001) who also study the structure preserving equivalent martingale measures for this multivariate model in an analogous manner to the univariate case. They also deliver closed form solutions for derivatives written on several assets (such as spread options or exchange options).

7. Other issues

In this section we briefly discuss four other issues: (i) likelihood based inference for SV models, (ii) term structure of interest rates based on OU type processes, (iii) time changed Lévy processes, (iv) modelling and prediction of volatility surfaces generated off Black-Scholes implied volatilities. Each of these topics has been extensively worked upon recently, but we have only the space to briefly mention them here.

7.1. Inference

Our main work on inference for SV models has been based on the use of realised variance. This delivers fast and reasonably accurate estimators which are free from discretisation errors. This work has already been discussed in Section 4.2 of this paper.

Another approach to estimating these types of models is based on building Bayesian Markov chain Monte Carlo (MCMC) algorithms. These likelihood based methods are familiar in the context of SV models following the work of Jacquier, Polson, and Rossi (1994), Kim, Shephard, and Chib (1998) and Elerian, Chib, and Shephard (2001). However, in the context of OU based volatility models a new type of algorithm is necessary — see the discussion in Barndorff-Nielsen and Shephard (2001a).

Independent work by Fruhwirth-Schnatter and Soegner (2001) and Roberts, Papaspiliopoulos, and Dellaportas (2001) have successfully analysed SV models based on the Γ-OU process and superpositions of that process. The key to both approaches is to employ a reparameterisation of the model into jump times and jump heights, handling the model through a reversible jump MCMC algorithm.

7.2. Term structure of interest rates

In the bond market framework introduced by Heath, Jarrow, and Morton (1992) the forward rate $f(t, T)$ maturing at T is described by

$$\mathrm{d}f(t, T) = \alpha(t, T)\mathrm{d}t + \sigma(t, T)\mathrm{d}w(t) \tag{7.1}$$

where the drift and volatility coefficients $\alpha(t, T)$ and $\sigma(t, T)$ are processes satisfying some mild integrability conditions and smoothness in time to maturity and w is a (possibly d-dimensional) Brownian motion.

Several modifications have been proposed, subjecting this model class for interest rates to an evolution similar to that of the Samuelson-Black-Scholes model for asset prices. In particular, forward rates driven by a purely jumping Lévy process or by a more general semimartingale have been considered by Eberlein and Raible (1999), Björk, Kabanov, and Runggaldier (1997) and Björk, Di Masi, Kabanov, and Runggaldier (1997). As for stochastic volatility, it is typically introduced in (7.1) via the following factorization of the volatility coefficient

$$\sigma(t, T) = \sqrt{\tau(t)}h(T - t)$$

into a strictly positive stationary process τ depending only on calendar time t and a deterministic, strictly positive function h modelling the behaviour of volatilities across different time to maturities $T - t$. In Barndorff-Nielsen, Christiansen, and Nicolato (2001), the process τ is assumed to be independent of w and it is described by an OU type process as in (2.6). Under these assumptions, explicit term structure derivative formulas are obtained. However, the leverage effect seems to be particularly relevant in the case of bond markets.

The construction and analysis of term structure models making use of volatility processes of the OU type and including (negative) correlation between bond prices and their volatilities is a topic of ongoing research.

7.3. Time changed Lévy processes

Recent contributions have focussed on introducing stochastic volatility in the context of exponential Lévy models. Given a Lévy process $z(t)$ and an independent positive process $\tau(t)$, Carr, Geman, Madan, and Yor (2001) and Winkel (2001) consider a time changed Lévy process

$$y^*(t) = z(\tau^*(t)), \tag{7.2}$$

where $\tau^*(t) = \int_0^t \tau(s)\mathrm{d}s$, while Eberlein, Kallsen, and Kristen (2001) examine

$$\mathrm{d}y^*(t) = \sqrt{\tau(t)}\mathrm{d}z(t). \tag{7.3}$$

Typical examples for the z process would be a zero mean variance gamma or *NIG* Lévy process. We will assume this to be the case here. Then we can write

$$z(t) = w(\chi^*(t)),$$

where χ^* is a subordinator and w is a Brownian motion independent of χ^* and τ. Therefore, both models (7.2) and (7.3) can again be thought of as time changed Brownian motions (2.2), since we express (7.2) as

$$y^*(t) \overset{\mathcal{L}}{=} w\left(\chi^*(\tau^*(t))\right)$$

and (7.3) as

$$y^*(t) \overset{\mathcal{L}}{=} w\left(\int_0^t \tau(s)\mathrm{d}\chi^*(s)\right).$$

Such a process will not, in general, have continuous sample paths.

An alternative approach, which has many of the features of the above process, is to model

$$y^*(t) \overset{\mathcal{L}}{=} \mu t + \beta\tau^*(t) + w_1\left(\tau^*(t)\right) + w_2\left(\chi^*(s)\right),$$

where w_1 and w_2 are independent standard Brownian motions which are also independent from τ^* and χ^*. Clearly this can be written as

$$y^*(t) \overset{\mathcal{L}}{=} \mu t + \beta\tau^*(t) + w\left(\tau^*(t) + \chi^*(s)\right),$$

which is both simple and tractable mathematically due to the linearity in the chronometer $\tau^*(s) + \chi^*(s)$.

7.4. Volatility surfaces

Based on a detailed and extensive empirical analysis Cont and da Fonséca (2001) have proposed an interesting dynamical model for the timewise behaviour of the surface of implied volatilities. Let $I_t(m,\tau)$ denote the implied volatility at time t, moneyness m[4] and time to maturity τ.

Then the authors represent the logarithm of the implied volatility by a Karhunen-Loéve decomposition into orthogonal terms

$$\log I_t(m,\tau) = X_\infty(m,\tau) + \sum_{k=1}^{\infty} x_k(t)f_k(m,\tau)$$

where the $f_k(m,\tau)$ are the eigenfunctions corresponding to the estimated covariance function $C(\xi_1,\xi_2)$, defined for points on the surface, and the coefficients $x_k(t)$ are modelled as independent OU processes. In this representation the functions $f_k(m,\tau)$ are deterministic and the randomness resides in the processes $x_k(t)$. In the actual empirical study, which was for options written on the SP500 index, the first two eigenmodes already captured most of the variation.

Hubalek and Tompkins (2000a) considered, among others, the SV model (3.1) with τ specified as a superposition of two Γ-OU processes. Choosing a particular structure preserving equivalent martingale measure, they compared the associated implied volatility surface with those observed

---

[4]m is defined as K/S_t where S_t denotes the price at time t of the underlying considered and K is the strike price of the option.

from options on four stock index futures markets (S&P500, FTSE100, DAX, NIKKEI225) and on three bond futures markets (US T-Bonds, German Bund, British Gilt). They found that smile surfaces arising from (3.1) could capture the shape (in particular the skew) of observed smile surfaces but systematic differences indicated that it is necessary to study the change of measure more thoroughly in order to calibrate the model to option prices.

8. Conclusion

This paper has reviewed some of our recent on work on stochastic volatility. The main areas we have covered are: (i) subordination, (ii) OU based models, (iii) analytic option pricing, (iv) the asymptotics of realised variance and realised power variation, (v) building and analysing multivariate models. Some of these results hold generally for SV models, others are built out of the specific OU structure for volatility we have been developing in a number of the papers.

9. Acknowledgments

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, and F. X. Diebold (2002). Parametric and nonparametric measurement of volatility. In Y. Ait-Sahalia and L. P. Hansen (Eds.), *Handbook of Financial Econometrics*. Amsterdam: North Holland. Forthcoming.

Andersen, T. G., T. Bollerslev, F. X. Diebold, and H. Ebens (2001). The distribution of realised stock return volatility. *Journal of Financial Economics 61*, 43–76.

Andersen, T. G., T. Bollerslev, F. X. Diebold, and P. Labys (2001). The distribution of exchange rate volatility. *Journal of the American Statistical Association 96*, 42–55.

Barndorff-Nielsen, O. E., B. J. Christiansen, and E. Nicolato (2001). General stochastic volatility and explicit option prices in Heath-Jarrow-Morton term structure analysis. Unpublished paper, Department of Mathematical Sciences, Aarhus University.

Barndorff-Nielsen, O. E. and N. Shephard (2001f). How accurate is the asymptotic approximation to the distribution of realised volatility? In D. W. F. Andrews, J. L. Powell, P. A. Ruud, and J. H. Stock (Eds.), *Identification and Inference for Econometric Models*. Festschrift for Thomas J. Rothenberg. Forthcoming.

Barndorff-Nielsen, O. E. and N. Shephard (2001c). Integrated OU processes and non-Gaussian OU-based stochastic volatility. Unpublished paper: Nuffield College, Oxford, Submitted.

Barndorff-Nielsen, O. E. and N. Shephard (2001b). Modelling by Lévy processes for financial econometrics. In O. E. Barndorff-Nielsen, T. Mikosch, and S. Resnick (Eds.), *Lévy Processes – Theory and Applications*, pp. 283–318. Boston: Birkhäuser.

Barndorff-Nielsen, O. E. and N. Shephard (2001a). 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 (2001d). Normal modified stable processes. Unpublished paper: Nuffield College, Oxford.

Barndorff-Nielsen, O. E. and N. Shephard (2001e). Realised power variation and stochastic volatility. Unpublished paper: Nuffield College, Oxford.

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*. Forthcoming.

Bellamy, N. and M. Jeanblanc (2000). Incompleteness of markets driven by a mixed diffusion. *Finance and Stochastics 4*, 209–222.

Björk, T., G. Di Masi, Y. Kabanov, and W. Runggaldier (1997). Towards a general theory of bond markets. *Finance and Stochastics 1*, 141–174.

Björk, T., Y. Kabanov, and W. Runggaldier (1997). Bond market structure in the presence of marked point processes. *Mathematical Finance 7*, 211–239.

Black, F. (1976). Studies of stock price volatility changes. *Proceedings of the Business and Economic Statistics Section, American Statistical Association*, 177–181.

Carr, P., H. Geman, D. B. Madan, and M. Yor (2001). Stochastic volatility for Lévy processes. Unpublished paper: Robert H Smith School of Business, Maryland University.

Carr, P. and D. B. Madan (1998). Option valuation using the fast Fourier transform. Unpublished paper: Robert H Smith School of Business, University of Maryland.

Christensen, B. J. and N. R. Prabhala (1998). The relation between implied and realized volatility. *Journal of Financial Economics 37*, 125–150.

Cont, R. and J. da Fonséca (2001). Deformation of implied volatility surfaces: an empirical analysis. R.I. No. 466. Centre de Mathematiques Appliques, Ecole Polytechnique, Paris.

Cvitanic, J., H. Pham, and N. Touzi (1999). Super-replication in stochastic volatility models with portfolio constraints. *Journal of Applied Probability 36*, 523–545.

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.

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

Eberlein, E. and J. Jacod (1997). On the range of options prices. *Finance and Stochastics 1*, 131–140.

Eberlein, E., J. Kallsen, and J. Kristen (2001). Risk management based on stochastic volatility. FDM Preprint 72, University of Freiburg.

Eberlein, E. and S. Raible (1999). Term structure models driven by general Lévy processes. *Mathematical Finance 9*, 31–53.

El Karoui, N. and M.-C. Quenez (1995). Dynamic programming and pricing of contingent claims in an incomplete market. *SIAM Journal on Control and Optimization 33*, 29–66.

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

Frey, R. and C. A. Sin (1999). Bounds on European option prices under stochastic volatility. *Mathematical Finance 9*, 97–116.

Fruhwirth-Schnatter, S. and L. Soegner (2001). Bayesian estimation of the Barndorff-Nielsen-Shephard stochastic volatility model with marginal gamma laws. Statistical research report, University of Vienna.

Heath, D., R. Jarrow, and A. Morton (1992). Bond pricing and the term structure of interest rates: A new methodology for contingent claim valuation. *Econometrica 60*, 77–105.

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.

Hubalek, F. and E. Nicolato (2001). On multivariate extensions of Lévy driven Ornstein-Uhlenbeck type stochastic volatility models and multi-asset options. Working Paper, University of Aarhus.

Hubalek, F. and R. Tompkins (2000a). Does it matter where jumps come from? underlying versus volatility processes. Working paper, Vienna University of Technology. (working paper).

Hubalek, F. and R. Tompkins (2000b). On explicit option pricing for stochastic volatility models driven by a superposition of Ornstein-Uhlenbeck type jump processes. Working paper, Vienna University of Technology. (working paper).

Hull, J. and A. White (1987). The pricing of options on assets with stochastic volatilities. *Journal of Finance 42*, 281–300.

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.

Jakubėnas, P. (1998). On option pricing in certain incomplete markets. *Proceedings of the Steklov Institute of Mathematics*. Forthcoming.

Kim, S., N. Shephard, and S. Chib (1998). Stochastic volatility: likelihood inference and comparison with ARCH models. *Review of Economic Studies 65*, 361–393.

Kramkov, D. O. (1996). Optional decomposition of supermartingales and hedging contingent claims in incomplete security markets. *Probability Theory and Related Fields 105*, 459–479.

Marcus, M. B. (1987). *ξ-Radial Processes and Random Fourier Series*, Volume 368. Memoirs of the American Mathematical Society.

Meddahi, N. and E. Renault (2002). Temporal aggregation of volatility models. *Journal of Econometrics*. Forthcoming.

Nelson, D. B. (1991). Conditional heteroskedasticity in asset pricing: a new approach. *Econometrica 59*, 347–370.

Nicolato, E. and E. Venardos (2001). Option pricing in stochastic volatility models of the Ornstein-Uhlenbeck type with a leverage effect. Unpublished paper: Dept. of Mathematical Sciences, Aarhus University.

Pitt, M. K. and N. Shephard (1999). Time varying covariances: a factor stochastic volatility approach (with discussion). In J. M. Bernardo, J. O. Berger, A. P. Dawid, and A. F. M. Smith (Eds.), *Bayesian Statistics 6*, pp. 547–570. Oxford: Oxford University Press.

Poterba, J. and L. Summers (1986). The persistence of volatility and stock market fluctuations. *American Economic Review 76*, 1124–1141.

Raible, S. (1998). Lévy processes in finance: Theory, numerics and empirical facts. PhD Thesis. Mathematics Faculty, Freiburg University.

Roberts, G., O. Papaspiliopoulos, and P. Dellaportas (2001). Bayesian inference for non-gaussian Ornstein-Uhlenbeck stochastic volatility processes. Submitted paper. Department of Mathematics and Statistics, Lancaster University.

Rosinski, J. (1991). On a class of infinitely divisible processes represented as mixtures of Gaussian processes. In S. Cambanis, G. Samorodnitsky, and M. S. Taqqu (Eds.), *Stable Processes and Related Topics*, pp. 27–41. Basel: Birkhäuser.

Rosinski, J. (2000). Series representations of Lévy processes from the perspective of point processes. In O. E. Barndorff-Nielsen, T. Mikosch, and S. Resnick (Eds.), *Lévy Processes – Theory and Applications*. Boston: Birkhäuser. Forthcoming.

Schwert, G. W. (1989). Why does stock market volatility change over time. *Journal of Finance 44*, 1115–1153.

Taylor, S. J. and X. Xu (1997). The incremental volatility information in one million foreign exchange quotations. *Journal of Empirical Finance 4*, 317–340.

Whittle, P. (1983). *Prediction and Regulation* (2 ed.). Oxford: Blackwell. 1st edition, 1963.

Winkel, M. (2001). The recovery problem for time-changed Lévy processes. MaPhySto Research Report 2001-37, Department of Mathematical Sciences, Aarhus University.