PROBABILITY
AND
MATHEMATICAL STATISTICS
Vol. 26, Fasc. 2 (2006), pp. 367–378

LÉVY PROCESSES AND SELF-DECOMPOSABILITY IN FINANCE

BY

N. H. BINGHAM (SHEFFIELD)

Abstract. The main theme of Urbanik's work was infinite divisibility and its ramifications. The aim of this memorial article is to trace the application of this theme in mathematical finance, one of the main growth areas in contemporary probability theory.

We begin in Section 1 with a discussion of the nature of prices. In particular, we focus on whether (or when) prices may be taken as continuous, with a view to using Lévy processes to model the case of prices with jumps. We turn in Section 2 to asset return distributions; prime candidates for modelling here include the normal, hyperbolic and Student t cases. In Section 3, we turn to distributions of type G, in particular, those in which the mixing law is not only infinitely divisible but also self-decomposable (i.e. in the class SD), which includes all three cases above. Then in Section 4 we turn to the dynamic counterpart of this, in which the law of class SD occurs as the limit law of a stochastic process of Ornstein–Uhlenbeck type, with Lévy driving noise. Finally, in Section 5 we discuss stochastic volatility models.

AMS Mathematics Subject Classification: Primary: 60E07, 60G51; Secondary: 91B28.

Key words and phrases: Infinite divisibility, self-decomposability, type G, elliptically contoured distributions, mathematical finance, portfolios, stochastic volatility, processes of Ornstein–Uhlenbeck type.

1. LÉVY PROCESSES IN FINANCE

The central theme of Urbanik's work in probability theory was infinite divisibility. This in turn is intimately linked with Lévy processes and their structure.

The history of Lévy processes in finance pre-dates, in a sense, that of Lévy processes themselves. As is well known, in his remarkable thesis, Louis Bachelier in 1900 (see [2]) was the first to use Brownian motion to model movements of stock prices, despite the fact that a proper mathematical basis for Brownian motion did not emerge until the work of Daniell in 1919–1921 and Wiener in 1923, which provided us with Wiener measure and the Wiener process (see

Shafer and Vovk [48], § 3.2, for Daniell's contribution). The Lévy–Khintchine formula, and the theory of Lévy processes, emerged through the work of several authors during the 1930s (see Lévy's obituary [35] for the detailed history). The use of Brownian motion as driving noise for modelling stock prices via geometric (or economic) Brownian motion was advocated by Samuelson from 1965 on [45], and this was the model used in the pioneering work by Black and Scholes [17] and of Merton [42] in 1973.

Now Brownian motion, and so geometric Brownian motion, is continuous, and so provides a model in which prices evolve continuously. Immediately, one has to stop, and consider this carefully. Are prices continuous? The answer, of course, is that it depends on the closeness with which prices are observed. In broad outline, prices do indeed evolve continuously, except under the influence of major economic shocks. In fine detail, prices jump. This is partly because prices are measured in terms of money, and money is quantized. More importantly, it is because prices are determined through trading – price is the level at which markets clear, or supply balances demand. Without trading, or in an illiquid market, one does not know how much an asset is worth. With trading, or in a liquid market, one does – not exactly, but approximately, or to within the interval within which price currently fluctuates under the influence of trading (or of the bid-ask spread needed to fund the maintenance of the market). Thus price is inherently dynamic.

Again, one needs to take into account the nature – principally, the size – of the economic agents involved, or the trades being made. Small investors, or minor economic agents, are *price takers and not price makers* – they are able to enter the market, trade in the volume they choose, and leave, without disturbing the market price. (Of course, this is only true to the approximation above: the very act of trading does shift price, if observed closely enough.) By contrast, large investors, or major economic agents, are *price makers*, because the size of trades they need to make inevitably shifts market prices (they thus lack the anonymity of the small trader, and this can be seriously damaging, especially when a large trader is forced to trade through publicly visible weakness).

The upshot of all this is that one needs to distinguish between different types of trading conditions, and model them differently. Under the 'normal scenario', one has the every-day movement of heavily-traded stocks under normal market conditions. Here, continuous price evolution (modelled by Brownian noise) may be suitable for many purposes – as in the Black–Scholes model, the benchmark model of mathematical finance. However, even here the model does not withstand close scrutiny, particularly over short time intervals. For example, it gives tails that are much thinner than those actually encountered. One has the attractive alternative of using a Lévy-based model instead (for textbook references, see [21], [47], [34]). Since the price movements one is attempting to model consist of 'jitter' – large numbers of very small movements taking place very rapidly – one has a wonderful modelling resource ready to hand. This is to use a Lévy measure which has *infinite mass* (in the neighbourhood of the origin, as it is finite elsewhere) – and thus produces infinitely many jumps in finite time.

The term used nowadays for infinite-mass Lévy processes is *infinite activity* (a term we learned from Professor Hélyette Geman). It provides the natural real-world context for infinite-mass Lévy processes. These processes, whose existence and path properties were laid bare by Lévy in the 1930s, once stood as prime examples of mathematical constructs which, while beautiful mathematically, seemed completely divorced from reality. Nowadays we are all used, as probabilists, to the extent to which our subject has been harnessed to serve the needs of mathematical finance, and of the financial services industry more generally. It is worth remarking that things move in the other direction too. Finance provides a setting in which some of our models in probability, previously regarded as arcane, idealized or as mathematics for its own sake, seem natural, realistic and inevitable.

The jumps in a Lévy process are very natural for modelling purposes in finance, and the first chapter in [21] gives a particularly good justification for them. By contrast, the independent increments assumption is less easy to defend. It is perfectly reasonable (at least to a first approximation) to treat tomorrow's price-sensitive information as independent of yesterday's, under normal market conditions. It is not reasonable during a sustained financial crisis. One's normal modelling assumptions thus break down, precisely when one needs them most – during a crisis. This is, of course, less an argument against Lévy models in finance than a recognition that one needs more than one model. At the very least, one needs a model for use during normal conditions, as above, *and* a model specifically designed for crises. Such models focus on *extremes*, rather than typical price movements: the relevant probability is *extreme value theory* (EVT), and the relevant finance is *quantitative risk management* (QRM). The application of EVT to QRM is very topical; a recent monograph account is [41], and a forthcoming one is [3].

We close this section by making some remarks on the interplay between economics and finance. Much of economics is concerned with how prices are arrived at. By contrast, in much of finance, one takes prices (of the underlying asset – the underlying) as given, and the focus is on questions – pricing, hedging and so on – concerning derivatives – things derived from the underlying. One may thus regard finance as a specialized part of economics, where prices are given. As the remarks above on agents being price takers or price makers show, this boundary between finance and economics is blurred rather than sharp. However, it is in the realm of infinite-activity Lévy-based models that this interface comes into focus.

It is interesting to compare the viewpoint above, in which we single out the infinite-activity case as crucial, to that expressed by Malliavin and Thalmaier ([37], § 8.1, p. 98) in their study of Malliavin calculus applied to mathematical finance. There, they restrict to the finite-activity case (which they call *finite type*), for technical convenience, as "this class of processes is qualitatively sufficient for the needs of mathematical finance". As they remark, the finite-activity case is dense in the general case.

2. RETURN DISTRIBUTIONS

In the Black–Scholes–Merton model, the benchmark model of mathematical finance, the price S_t of a stock evolves over time in such a way that the log-price $\log S_t$ is a Brownian motion (see e.g. [11], or any text on mathematical finance). Instead of log-prices, one may focus on *returns* over some return interval of length $\delta > 0$. These are the relative price changes

$$R_n := (S_{(n+1)\delta} - S_{n\delta})/S_{n\delta}.$$

Using the Taylor approximation $\log(1+x) \sim x$ for small x shows that working with returns is substantially equivalent to working with log-prices. Because of their great financial importance, return distributions have been studied in depth; see e.g. [32] for background.

The properties of return distributions depend (inter alia) on the length of the return interval δ. For long δ (of the order of a month, say), since the return over a month is the sum of the returns over the days of the month, and these may be taken independent (at least approximately), one has *aggregational Gaussianity*: the return is the sum of a sizeable number of approximately independent random variables, and so the central limit theorem applies; the returns are thus Gaussian, and one is back with the Black–Scholes–Merton model as described above. At the other extreme, one may have δ small, and be dealing with *high-frequency data* (*tick data*, with δ of the order of minutes or seconds, is common nowadays). For reasons involving *scaling* arguments (akin to those arising in physics) – see [38] – the return distributions in such cases have *heavy tails* (*Pareto* tails – decreasing like a power, or like a regularly varying function). This is in stark contrast to the ultra-thin tails in the Gaussian case above with δ large. As one might expect, for δ intermediate – daily returns, say – one obtains intermediate tail decay – typically, *semi-heavy* tails, in which the log-density decays linearly, rather than like a quadratic as in the Gaussian case or a logarithm in the Pareto case. This occurs in the *hyperbolic* model, for which see e.g. [12] and the references cited there.

Because the return over interval δ is the sum of the n returns over subintervals of length δ/n, and these are independent (to the order of accuracy considered here), return distributions are typically *infinitely divisible*. As the return interval δ varies, and indeed as the type of asset one is dealing with varies, the modelling flexibility provided by the Lévy–Khintchine formula becomes available; see e.g. [11], § 2.10–12, § 5.5. The field of *Lévy finance* is so important that several books are now devoted to it; see [47], [21], [34], cited above.

Recall that the infinitely divisible laws are the limit laws of triangular arrays – two-suffix arrays of independent random variables, individually negligible. It is plausible that one will still have sufficient modelling flexibility if one restricts this from two suffices to one – limits of independent *sequences*, suitably normed; the class of limit laws so obtained is the class of *self-decomposable* distributions, *SD*. This class *SD* has been found to serve very well, from the distributional or static point of view. Furthermore, it also serves from the *dynamic* point of view, when one considers *time series*. For, the defining property of self-decomposability is that, for each $c \in (0, 1)$, X should satisfy

$$X \stackrel{d}{=} cX + X_c$$

for some random variable X_c, where $\stackrel{d}{=}$ denotes equality in distribution and the variables on the right are independent. This relation has the form of an *autoregressive* scheme of order 1, thereby making available much of the machinery of time series (see e.g. [18]). For both these reasons, the class *SD* is a prime candidate for use in modelling asset return distributions.

Three prime examples are to hand:

1. *Normal distributions*. This is the Gaussian case of the Black–Scholes–Merton model, relevant to (say) monthly returns (the rule of thumb is that 16 trading days suffice for aggregational Gaussianity).

2. *Hyperbolic distributions*. Self-decomposability is due to Halgreen [27] in 1979. The log-density has linear asymptotes at $\pm\infty$, like the lower branch of a hyperbola (semi-heavy tails).

3. *Student t-distributions*. Infinite divisibility is due to Grosswald [26] in 1976; for self-decomposability see e.g. Jurek [29]. The density decays like $x^{-(\nu+1)}$, where ν is the degrees of freedom (heavy tails). Although the limit as $\nu \to \infty$ is Gaussian, this passage to the limit skips over the semi-heavy tails above.

Although the modern era in mathematical finance began in 1973 with Black, Scholes and Merton, mathematical finance itself goes back to 1952, with the work of Markowitz [40]. Markowitz left us two key insights:

1. *Look at risk and return together*, not separately. (Risk is measured by variances or covariances, return by means, hence *mean-variance theory*.)

2. *Diversify*. In order to protect oneself against the uncertainty inseparable from holding risky stock, one should hold a portfolio – a range of different assets – *balanced*, so that changes that harm some of our holdings will help others. This balance requires *negative correlations* in our holdings.

Thus by Markowitzian diversification, one should work in d dimensions, where d is the number of different assets we hold (d may be large, as the range of investment opportunities open to us is unlimited), and one should start with the d-vector μ of means and the $d \times d$ covariance matrix Σ, which should exhibit plentiful negative correlations.

The upshot of the above is that our asset returns will be modelled in d dimensions by a self-decomposable distribution, or process, with (μ, Σ) as a parameter.

3. DISTRIBUTIONS OF TYPE G

We turn now to a suitable subclass of SD_d, the class of self-decomposable distributions in d dimensions. A random d-vector Y, or its distribution ν, is said to be of *type* G (following Marcus [39] in 1987) if

$$Y = \sigma\varepsilon$$

in distribution, where σ, ε are independent, σ^2 is ID and $\varepsilon \sim N_d(0, \Sigma)$ is multivariate normal (multinormal). (Other definitions of type G in d dimensions are in use; see [7] for details and references.) Then Y has characteristic function (CF)

$$(G) \qquad \psi_Y(t) := E[\exp\{it^T Y\}]$$
$$= E\big[E[\exp\{t^T \sigma\varepsilon\}] \,|\, \sigma\big] = E[\exp\{-\tfrac{1}{2}\sigma^2 t^T \Sigma t\}] = \phi(\tfrac{1}{2}t^T \Sigma t),$$

where ϕ is the Laplace–Stieltjes transform (LST) of σ^2. Thus

$$X := \mu + Y$$

has CF

$$\psi_X(t) = \exp\{it^T \mu\}\, \phi(\tfrac{1}{2}t^T \Sigma t),$$

and so X is elliptically contoured, $X \sim EC_d(\mu, \Sigma, \phi)$ in the notation of [24], Definition 2.2. Also, as both the definition as an independent product and the above derivation of the CF show, Y is a normal variance mixture (NVM) (see [24], Chapter 2).

Suppose now that the law of σ^2 is not only infinitely divisible but also self-decomposable. That is, for each $c \in (0, 1)$,

$$(SD) \qquad \phi(s) = \phi(cs) \cdot \phi_c(s)$$

for some LST ϕ_c. Replace s by $\frac{1}{2}t^T \Sigma t$. As in the proof of (G), each of the three terms is the CF of a d-vector, which shows that Y and X are also SD. They are thus absolutely continuous ([46], Theorem 27.13). The density generator g of X thus exists ([13]; [24], § 2.2.3). Since the density of X is unimodal ([46], § 53) and is a function g of the quadratic form $(x-\mu)^T \Sigma^{-1} (x-\mu)$, g is decreasing. As σ^2 is SD, its Lévy measure is absolutely continuous, with density of the form $k(x)/x$ with k decreasing (see e.g. [46], Corollary 15.11).

From some points of view, the cumulant-generating functions are more convenient. Writing

$$\kappa_{\mathbf{X}}(\mathbf{t}) := \log \psi_{\mathbf{X}}(\mathbf{t}), \quad \kappa(\mathbf{t}) := \log \psi_{\mathbf{Y}}(\mathbf{t}), \quad K(s) := \log \phi(s),$$

we have

$$\kappa_{\mathbf{X}}(\mathbf{t}) = \mathbf{t}^T \mu + \kappa(\mathbf{t}), \quad \kappa(\mathbf{t}) = K(\tfrac{1}{2}\mathbf{t}^T \Sigma \mathbf{t}).$$

Distribution of type G with σ^2 SD are suitable for modelling asset return distributions in d dimensions; for background and details, see [13] and [14].

4. PROCESSES OF ORNSTEIN–UHLENBECK TYPE

We now introduce dynamics into the picture. Each SD ν is the limiting law, of Y_∞ say, of the process $Y = (Y_t\colon t \geqslant 0)$ of Ornstein–Uhlenbeck (OU) type given by the solution to the stochastic differential equation

$$\text{(OU)} \qquad dY_t = -cY_t + dZ_t,$$

where $c > 0$ and $Z = (Z_t\colon t \geqslant 0)$ is a Lévy process (the background driving Lévy process or BDLP) whose Lévy measure ν_0 satisfies the logarithmic integrability condition

$$\text{(log)} \qquad \int \log^+(|x|)\, d\nu_0(x) < \infty,$$

and conversely: each Lévy process satisfying (log) gives an SD law in this way. The stochastic representation

$$Y_\infty \stackrel{d}{=} \int_0^\infty e^{-ct}\, dZ_t$$

holds, and the cumulants are linked by

$$(\kappa) \qquad \kappa_{Y_\infty}(z) = \int_0^\infty \kappa_{Z_1}(ze^{-cs})\, ds$$

([46], § 17; [1]), or equivalently

$$\kappa_{Y_\infty}(z) = \int_0^z \kappa_{Z_1}(u)\, du/u, \quad \kappa_{Z_1}(z) = z\kappa'_{Y_\infty}(z).$$

Such processes provide a way to model asset returns dynamically: the distributional properties for fixed time are as above, and the process is stationary (either by starting in the stationary distribution, or by leaving equilibrium to be approached as time elapses). Of course, such processes are discontinuous except in the case where the BDLP is Brownian motion and the process of OU type is the classical Ornstein–Uhlenbeck process.

The covariance is of course undefined except in the L_2 case, when the autocorrelation is as in the classical Ornstein–Uhlenbeck case:

$$\operatorname{corr}(Y_t, Y_{t+u}) = \exp(-c\,|u|).$$

By superposition of independent processes of OU type, correlations of the form

$$\sum_{i=1}^{m} w_i \exp(-c_i\,|u|)$$

can be obtained (for example, the case $m = 2$ allows one to handle both 'fast' and 'slow' effects, which is often useful). By passage to the limit (or by using independently scattered random measures and Lévy random fields), correlations decaying like a power rather than an exponential, and so giving *long-range dependence* (LRD), may be constructed. See [4] for background and details.

EXAMPLES. 1. *Student processes.* The Student $t = t(d, v, \Sigma)$ distribution in d dimensions with v degrees of freedom and covariance matrix Σ is defined by the density

$$f(\boldsymbol{x}) = \frac{1}{|\Sigma|^{1/2}} \frac{1}{\sqrt{(\pi v)^d}} \frac{\Gamma\left(\frac{1}{2}(v+d)\right)}{\Gamma\left(\frac{1}{2}v\right)} \left(1 + \frac{\boldsymbol{x}^T \Sigma^{-1} \boldsymbol{x}}{v}\right)^{-(v+d)/2}.$$

This is thus elliptically contoured, with density generator of the form

$$g(x) = c\left(1 + \frac{x^2}{v}\right)^{-(v+d)/2},$$

and as noted above it is also SD. Further, it is NVM: with the inverse gamma distribution $I\Gamma = I\Gamma(\alpha, \beta)$ defined for $\alpha, \beta > 0$ by the density

$$f(x) = \frac{\beta^\alpha}{\Gamma(\alpha)} x^{-\alpha-1} e^{-\beta/x} \qquad (x > 0),$$

this is the mixture of $N_d(0, \Sigma)$ with mixing law $I\Gamma(\frac{1}{2}v, \frac{1}{2}v)$. One may thus find a stationary Markov process Y with limiting law $t = t(d, v, \Sigma)$; see Heyde and Leonenko [28], Theorem 3.2. For $K_\lambda(x)$ the Bessel function of the third kind (Macdonald function)

$$K_\lambda(x) = \frac{1}{2}\int_0^\infty u^{\lambda-1} \exp\{-\tfrac{1}{2}x(u + 1/u)\}\,du,$$

the CF is given by

$$\exp\{i\boldsymbol{t}^T \mu\} \cdot K_{(v+d-1)/2}(\sqrt{v}\,u) \cdot (\sqrt{v}\,u)^{(v+d-1)/2} \cdot \frac{2^{1-(v+d-1)/2}}{\Gamma\left(\frac{1}{2}(v+d-1)\right)}, \qquad u := \boldsymbol{t}^T \Sigma \boldsymbol{t}.$$

See e.g. [28], (2.5)–(2.7) and Remark 2.2, and [33].

The Student t-distributions have Pareto tail-decay. They are thus useful for modelling stationary processes with heavy tails. Such tails may occur towards the high-frequency end of the data spectrum. For very high-frequency data, however, the elliptically contoured property may not hold, and details of market microstructure involving trading hours, lunch breaks and the like become important. For background, including an empirical study, see e.g. Bingham and Schmidt [16].

2. *Hyperbolic processes.* If instead of the inverse gamma distribution one uses the generalized inverse Gaussian (GIG) mixing laws, one obtains generalized hyperbolic (GH) distributions as the limit laws of the OU processes. Such GH laws have semi-heavy tails, and may occur for medium-frequency data (daily returns, for example); they were studied in [12], and were one of the motivating examples for [13]. A different dynamic version is contained in the work of Barndorff-Nielsen and Pérez-Abreu [6].

5. STOCHASTIC VOLATILITY

In the classical Black–Scholes–Merton model, the volatility of a stock is a parameter, the standard deviation of the return, measuring the sensitivity of the price to new information. In the Black–Scholes formula, the option price does not depend on the mean return, but does depend crucially on the volatility – which is unobserved, and has to be estimated, whether from past prices (historic volatility) or by inference from observed option prices (implied volatility). Because of its importance, volatility has been intensively studied – and this has revealed that volatility is not constant, but varies. Since the variability of volatility (or volatility of volatility, 'vol of vol') is difficult to account for in terms of what can be measured (asset price, strike price, time to maturity etc.), it is natural to take it as stochastic and use a *stochastic volatility* model (SV).

In the BNS model (Barndorff-Nielsen and Shephard [8]), one takes the log-price process x^* as in the Black–Scholes–Merton model,

$$dx^*(t) = \big(\mu + \beta\sigma^2(t)\big)\,dt + \sigma(t)\,dw(t),$$

where w is a Brownian motion driving the log-price process and the volatility process $\sigma^2(t)$ is assumed stationary and independent of w. It may thus be modelled as an OU-process,

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

where $z = \big(z(t)\big)_{t\geq 0}$ is a subordinator independent of w (as the increments of z are positive, the process σ^2 is also positive, as required by its interpretation as a volatility). Subject to the logarithmic integrability condition on z, the process σ^2 is well-defined as a stationary Markov process, whose equilibrium distribution is self-decomposable. One may approach the modelling either via this equilibrium distribution, or via the subordinator. A range of examples are considered in [8]; see also [9]. They also consider simulation, via series representations, and fit their model to various financial data sets. Pricing of financial derivatives is also discussed, and hedging (their model is arbitrage-free, so equivalent martingale measures exist, but as Lévy-based models are incomplete, they are not unique). Their approach has been influential, and is widely used.

The multivariate case is also considered in [8], § 6.4, § 6.5. In particular, *factor models* are considered (§ 6.5.2). Here, the dimensionality d of our number of assets is reduced to some lower effective dimensionality, q say, reflecting the fact that asset prices often move together, under the influence of a smaller number of driving mechanisms. For example, q may be the number of industrial sectors represented in our portfolio, and there may be symmetry within but not between sectors. See [14] and [15] for models of this kind.

For other approaches to stochastic volatility modelling, in one or higher dimensions and using Lévy or Ornstein–Uhlenbeck processes, see e.g. Barrucci et al. [10], Malliavin and Mancino ([36], or [37], Chapter 2), Geman et al. [25], Carr et al. [20], and Nicolato and Venardos [43]. For a comparison of COGARCH (continuous-time generalized autoregressive conditional heteroscedastic) and Orstein–Uhlenbeck approaches, see Klüppelberg et al. [31].

In the multidimensional case, one approach is to model the evolution of the stochastic volatility matrix Σ_t over time. Here, covariances and correlations may evolve with time, as happens with actual portfolios. One may harness for this purpose recent results in the theory of random matrices. For a recent study of this kind, see Philipov and Glickman [44], who use *Wishart processes* (Bru [19]), and work in discrete time. Wishart processes form a natural modelling tool in this area (and also in the theory of random matrices, an area of great current interest); for background, see [23] and [22]. As in [31], comparison between discrete and continuous time is very interesting, and is the subject of current work.

In Memoriam, Kazimierz Urbanik (1930–2005)

Like so many others, I was deeply influenced by Urbanik's work. Part of my thesis (1969) was influenced by his generalized convolutions, and this influence extended into several of my papers in the 1970s and 1980s. During this time, I had the pleasure of getting to know Urbanik, when he visited me in London. I always considered his work — and this aspect of my own work — as theoretical. Since my interests expanded into mathematical finance in the 1990s, it has been a constant source of interest to see how theoretical probability of this kind has found applications in fields unthought of in earlier times. It is for this reason that I chose this subject matter for my contribution to the Urbanik Memorial Volume.

REFERENCES

[1] T. Aoyama and M. Maejima, *Characterizations of subclasses of type G distributions in R^d by stochastic integral representations*, Bernoulli, to appear.

[2] L. Bachelier, *Théorie de la spéculation*, Ann. Sci. École Norm. Sup. 17 (1990), pp. 21–86. Reprinted in: *Louis Bachelier's Theory of speculation. The origins of modern finance*. Translated and with commentary by M. Davis and A. Etheridge (foreword by P. A. Samuelson), Princeton University Press, Princeton, NJ, 2006.

[3] A. A. Balkema and P. Embrechts, *High-risk Scenarios and Extremes: A Geometric Approach*, to appear.

[4] O. E. Barndorff-Nielsen, *Superpositions of Ornstein–Uhlenbeck type processes*, Theory Probab. Appl. 45 (2001), pp. 175–194.

[5] O. E. Barndorff-Nielsen, M. Maejima and K. Sato, *Some classes of multivariate infinitely divisible distributions admitting stochastic integral representations*, Bernoulli 12 (2006), pp. 1–34.

[6] O. E. Barndorff-Nielsen and V. Pérez-Abreu, *Stationary and self-similar processes driven by Lévy processes*, Stochastic Process. Appl. 84 (1999), pp. 357–369.

[7] O. E. Barndorff-Nielsen and V. Pérez-Abreu, *Extensions of type G and marginal infinite divisibility*, Theory Probab. Appl. 47 (2002), pp. 202–218.

[8] O. E. Barndorff-Nielsen and N. Shephard, *Ornstein–Uhlenbeck based models and some of their uses in financial econometrics (with discussion)*, J. Roy. Statist. Soc. Ser. B 63 (2001), pp. 197–241.

[9] O. E. Barndorff-Nielsen and N. Shephard, *Econometric analysis of realised volatility and its use in estimating stochastic volatility models*, J. Roy. Statist. Soc. Ser. B 64 (2002), pp. 253–280.

[10] E. Barrucci, P. Malliavin, M. E. Mancino, R. Renó and A. Thalmaier, *The price-volatility feedback rate: an implementable mathematical indicator of market stability*, Math. Finance 13 (2003), pp. 17–35.

[11] N. H. Bingham and R. Kiesel, *Risk-neutral Valuation: Pricing and Hedging of Financial Derivatives*, 2nd edition, Springer 2004 (1st edition 1998).

[12] N. H. Bingham and R. Kiesel, *Modelling asset returns with hyperbolic distributions*, pp. 1–20 in [32].

[13] N. H. Bingham and R. Kiesel, *Semi-parametric modelling in finance: Theoretical foundations*, Quant. Finance 2 (2002), pp. 241–250.

[14] N. H. Bingham, R. Kiesel and R. Schmidt, *Semi-parametric modelling in finance: Econometric applications*, Quant. Finance 3 (2003), pp. 426–441.

[15] N. H. Bingham, R. Kiesel and R. Schmidt, *Multivariate elliptic processes*, preprint, Faculty of Mathematics and Economics, University of Ulm.

[16] N. H. Bingham and R. Schmidt, *Interplay between distributional and temporal dependence. An empirical study with high-frequency asset returns*, pp. 69–90 in [30].

[17] F. Black and M. Scholes, *The pricing of options and corporate liabilities*, J. Political Economy 72 (1973), pp. 637–659.

[18] P. Brockwell and R. A. Davis, *Time Series: Theory and Methods*, 2nd edition, Springer, 1991 (1st edition 1987).

[19] M.-F. Bru, *Wishart processes*, J. Theoret. Probab. 4 (1991), pp. 725–751.

[20] P. Carr, H. Geman, D. B. Madan and M. Yor, *Stochastic volatility for Lévy processes*, Math. Finance 13 (2003), pp. 345–382.

[21] R. Cont and P. Tankov, *Financial Modelling with Jump Processes*, Chapman and Hall, 2004.

[22] Nizar Demni, *Laguerre process and generalized Hartman–Watson law*, preprint, Lab. des Probabilités, Univ. Paris–VI, 2005.

[23] C. Donati-Martin, Y. Doumerc, H. Matsumoto and M. Yor, *Some properties of the Wishart processes and matrix extensions of the Hartman–Watson laws*, Publ. Res. Inst. Math. Sci. 40 (2004), pp. 1385–1412.

[24] K.-T. Fang, S. Kotz and K.-W. Ng, *Symmetric Multivariate and Related Distributions*, Chapman and Hall, 1990.

[25] H. Geman, D. B. Madan and M. Yor, *Stochastic volatility, jumps and hidden time-changes*, Finance Stoch. 6 (2002), pp. 63–90.
[26] E. Grosswald, *The Student t-distribution of any degrees of freedom is infinitely divisible*, Z. Wahrschein. Verw. Gebiete 36 (1976), pp. 103–109.
[27] C. Halgreen, *Self-decomposability of the generalized inverse Gaussian and hyperbolic distribution functions*, Z. Wahrschein. Verw. Gebiete 47 (1979), pp. 13–18.
[28] C. C. Heyde and N. N. Leonenko, *Student processes*, J. Appl. Probab. 37 (2005), pp. 342–365.
[29] J. Jurek, *Remarks on the self-decomposability and new examples*, Demonstratio Math. 34 (2001), pp. 241–250.
[30] Yu. Kabanov, R. Liptser and J. Stoyanov (Eds.), *From Stochastic Calculus to Mathematical Finance. The Shiryaev Festschrift*, Springer, 2006.
[31] C. Klüppelberg, A. Lindner and R. A. Maller, *Continuous-time volatility modelling: COGARCH versus Ornstein–Uhlenbeck models*, pp. 393–419 in [30].
[32] J. Knight and S. Satchell (Eds.), *Return Distributions in Finance*, Butterworth–Heinemann, 2001.
[33] S. Kotz and S. Nadarajah, *Multivariate t-distributions and their Applications*, Cambridge University Press, 2004.
[34] A. Kyprianou, W. Schoutens and P. Willmott, *Exotic Option Pricing and Advanced Lévy Models*, Wiley, 2005.
[35] M. Loève, *Paul Lévy (1886–1971), obituary*, Ann. Probab. 1 (1973), pp. 1–18.
[36] P. Malliavin and M. E. Mancino, *Fourier series method for measurement of multivariate volatilities*, Finance Stoch. 6 (2002), pp. 49–61.
[37] P. Malliavin and A. Thalmaier, *Stochastic Calculus of Variations in Mathematical Finance*, Springer, 2006.
[38] B. B. Mandelbrot, *Fractals and Scaling in Finance: Discontinuity, Concentration and Risk. Selecta Volume E*, Springer, 1997.
[39] M. B. Marcus, *ξ-radial Processes and Random Fourier Series*, Mem. Amer. Math. Soc. 368 (1987).
[40] H. Markowitz, *Portfolio selection*, J. Finance 7 (1952), pp. 77–91.
[41] A. McNeil, R. Frey and P. Embrechts, *Quantitative Risk Management: Concepts, Techniques, Tools*, Princeton University Press, 2005.
[42] R. F. Merton, *Theory of rational option pricing*, Bell J. Economics and Management Science 4 (1973), pp. 141–183 (reprinted as Chapter 8 in R. C. Merton, *Continuous-time Finance*, Blackwell, 1990).
[43] E. Nicolato and E. Venardos, *Option pricing in stochastic volatility models of Ornstein–Uhlenbeck type*, Math. Finance 13 (2003), pp. 445–466.
[44] A. Philipov and M. E. Glickman, *Multivariate stochastic volatility via Wishart processes*, preprint, Dept. Finance, American University, Vienna, VA, 2005.
[45] P. Samuelson, *Rational theory of warrant pricing*, Industrial Management Review 6 (1965), pp. 13–39.
[46] K. Sato, *Lévy Processes and Infinitely Divisible Distributions*, Cambridge University Press, 1999.
[47] W. Schoutens, *Lévy Processes in Finance: Pricing Financial Derivatives*, Wiley, 2003.
[48] G. Shafer and V. Vovk, *The sources of Kolmogorov's* Grundbegriffe, Statist. Sci. 21 (2006), pp. 70–98.

Department of Probability and Statistics
University of Sheffield
Sheffield S3
7RH, UK
nick.bingham@btinternet.com

Received on 29.8.2006