

Simulation Methods for Lévy-Driven CARMA Stochastic Volatility Models

Viktor Todorov
Duke University

George Tauchen
Duke University

This Draft: August 25, 2005 ^{*†}

Abstract

We develop simulation schemes for the new classes of non-Gaussian pure jump Lévy processes for stochastic volatility. We write the price and volatility processes as integrals against a vector Lévy process, which then makes series approximation methods directly applicable. These methods entail simulation of the Lévy increments and formation of weighted sums of the increments; they do not require a closed-form expression for a tail mass function nor specification of a copula function. We also present a new, and apparently quite flexible, bivariate mixture of gammas model for the driving Lévy process. Within this setup, it is quite straightforward to generate simulations from a Lévy-driven CARMA stochastic volatility model augmented by a pure-jump price component. Simulations reveal the wide range of different types of financial price processes that can be generated in this manner, including processes with persistent stochastic volatility, dynamic leverage, and jumps.

Keywords: Lévy process, simulation, stochastic volatility, diffusions, realized variance, quadratic variation.

JEL classification: G12, C51, C52.

*Author's Contact: Viktor Todorov: viktort@econ.duke.edu; George Tauchen: george.tauchen@duke.edu. Both at: Department of Economics, Duke University, Box 90097, Durham NC 27708.

†We thank members of the Duke Monday Econometrics and Finance Lunch Group for helpful comments along the way. We are especially grateful to Neil Shephard for detailed comments on a previous draft. Also, conversations with Tim Bollerslev were very important for the material in Subsection 6.3 that considers estimation issues in a very high frequency context.

1 Introduction

Modeling the evolution of a financial price series as forced by a stochastic volatility process has a long history in financial econometrics. In most models, the underlying driving process(es) are locally Gaussian — or possibly locally Gaussian with occasional rare jumps — and positivity of the volatility process is ensured by the functional form assumptions. Recently, Barndorff-Nielsen and Shephard (2001a,b) suggest a completely new class of models, termed non-Gaussian OU Models, where the driving process for a volatility factor is a pure-jump Lévy process with nonnegative increments; simple parametric sign restrictions ensure positivity. Brockwell (2001a,b) and Brockwell and Marquardt (2005) introduce a generalization to the Lévy-driven CARMA (continuous time autoregressive moving average) class of volatility models. These newer classes of models based on more general Lévy processes can be expected to supplant traditional Brownian-based processes in serious efforts to model the movements of financial price data at the very high frequency. Tauchen (2004) reviews the older classes of models and discusses some of the issues related to data and estimation methods for the newer classes of models.

Regardless of the estimation technique, it is clear that simulation will play a crucial role in the implementation of these newer classes of processes. With Bayesian methods, for example, simulation is used as part of the scheme to integrate out unobserved variables, including the unobserved values of the process between the sampling points. In frequentist likelihood-based approaches, simulation in conjunction with a cleverly chosen importance function, can, in certain problems, make evaluation of the likelihood practicable. In method-of-moments-based approaches, simulation is used to evaluate predicted moments under the model that are compared to sample moments via a chi-squared criterion.

In this paper, we develop and assess practical schemes to simulate from Lévy-driven models for financial price dynamics. In the models considered below, the volatility dynamics are governed by the Brockwell-style CARMA extension of the Barndorff-Nielsen and Shephard non-Gaussian OU setup. The returns process also contains a jump component, which is correlated with the jump innovations in volatility in order to accommodate the so-called leverage effect. We use simulation schemes based on series expansions that turn out to be considerably simpler to implement than schemes based on the tail mass function. The convenience of the series expansion is especially apparent in a bivariate (or multivariate) situation, because the need to determine the tail mass

functions and copula function is completely circumvented. Below, we introduce a two-dimensional mixture of gammas Lévy process that is extremely flexible while at the same time preserving positivity for the volatility Lévy increments and generating leverage type correlations between the volatility increments and price jump increments.

The remainder of the paper is organized as follows. Section 2 sets forth some notation and recalls some basic properties of the Lévy process. Section 3 sets out the series approximations to Lévy processes; it also shows how to adapt the series approximations to simulate Lévy-driven processes and assesses accuracy. Section 4 develops the flexible mixture of gammas process. Section 5 casts the material into a stochastic volatility framework. Section 6 contains examples that illustrate the flexibility of the simulation schemes and the realistic financial price dynamics that can be generated with judicious choice of the parameters; it also discusses challenges in estimation with such processes on very high frequency data. Section 7 contains the concluding remarks.

2 Lévy processes

In this section we review some basic facts associated with the Lévy processes, that will be used in the paper. For more details see Sato (1999) or Bertoin (1996).

Intuitively the Lévy process could be described as continuous time analogue to the random walk in discrete time. The following definition of Lévy process is taken from Sato (1999): A stochastic process $\{L_t : t \geq 0\}$ on $(\Omega, \mathcal{F}, \mathbb{P})$, taking values in \mathbb{R}^d is a Lévy process if the following conditions are satisfied

1. It has independent increments.
2. $L_0 = 0$ a.s.
3. The increments of the process are strictly stationary.
4. It is stochastically continuous.
5. There is $\Omega_0 \in \mathcal{F}$ with $P(\Omega_0) = 1$, such that that for every $\omega \in \Omega_0$ $L_t(\omega)$ is *càdlàg* (i.e it is right continuous with left limits).

The last condition says that the Lévy process has *càdlàg* version and it will be this version of the process which will be used throughout. If this condition is dropped the resulting process is the Lévy process in law.

There is an intimate link between the infinite divisible distributions and the Lévy processes in law. If $\{L_t : t \geq 0\}$ is a Lévy process in law, then the distribution of L_t for every $t \geq 0$ is infinitely divisible. The converse is also true for every infinitely

divisible distribution μ there exists Lévy process in law $L(t)$, such that the distribution of $L(1)$ is equal to μ . Therefore using the Lévy-Khintchine representation of the characteristic function of infinitely divisible distribution we can write the following for the characteristic function of a d -dimensional Lévy process $L(t)$

$$E[e^{i\varepsilon' L(t)}] = e^{t\ell(\varepsilon)}, \quad (1)$$

where

$$\ell(\varepsilon) = -\frac{1}{2}\varepsilon' \mathbf{A} \varepsilon + i\varepsilon' \mathbf{a} + \int_{\mathbb{R}_0^d} \left(e^{i\varepsilon' y} - 1 - i\varepsilon' y 1_{[|y| \leq 1]} \right) \nu(dy), \quad (2)$$

and \mathbf{A} is a symmetric positive semi-definite matrix. The measure ν , called the Lévy measure, on \mathbb{R}_0^d ($\mathbb{R}_0^d := \mathbb{R}^d \setminus \{0\}$), satisfies

$$\int_{\mathbb{R}_0^d} (|y|^2 \wedge 1) \nu(dy) < \infty. \quad (3)$$

The characteristic triplet $(\mathbf{A}, \mathbf{a}, \nu)$ of the measure μ completely determines the Lévy process L . The representation in (2) is not unique. There are many truncation functions which could be used besides $y 1_{[|y| \leq 1]}$ employed here; see Sato (1999) (p.38) for details.

Two specific cases of the Lévy process are:

1. $\nu = 0$. In this case the process reduces to Brownian motion and therefore has a continuous version.
2. $\mathbf{A} = \mathbf{0}$. In this case the process is pure jump.

Every other Lévy process is a combination of these two. The continuous part of every Lévy process is the Brownian motion, which has unbounded variation and quadratic variation proportional to time. The pure jump part of every Lévy process is of finite activity when $\nu(\mathbb{R}_0^d) < \infty$, and it is of infinite activity when $\nu(\mathbb{R}_0^d) = \infty$. The finitely active pure jump Lévy processes (also known as compound Poisson) jump at most a finite number of times in every finite interval almost surely, while the infinitely active jump processes jump an infinite number of times on finite intervals. Further, the set of infinitely active pure jump processes can be subdivided into those with finite variation or infinite variation. For a pure jump process to be of finite variation it is necessary and sufficient that $\int_{|y| \leq 1} |y| \nu(dy) < \infty$; see Sato (1999) (th.21.9). Intuitively, the pure jump processes of finite variation are characterized by the property that their trajectories are of finite length over finite intervals, almost surely, unlike those of infinite variation.

Since we use Lévy processes for modelling directly the stochastic volatility, we are interested in those that are increasing; i.e., for $d = 1$, $L_t(\omega)$ is an increasing function

of t . Such Lévy processes are called Lévy subordinators. For all Lévy subordinators $\nu[(-\infty, 0)] = 0$, and $\int_{0 < y \leq 1} y \nu(dy) < \infty$; see Sato (1999) (th.21.5).

3 Series Representation of Lévy Processes and Their Simulation

In this section we introduce the series representation of pure jump Lévy processes, which offers a convenient way for simulation of integrals with respect to them. In only certain cases is the transition density of the Lévy-process known in closed form (see for example Li, Wells, and Yu (2004) and references therein). The series representation of Lévy processes offers a particularly useful way of simulating Lévy processes, given the fact that in most cases the law of the increments of the Lévy process is not known in closed form. For its implementation, one needs to know only a shot noise decomposition of the Lévy measure of the process (to be explained below), without employing the transition density which will be often not known explicitly and rather expensive to evaluate numerically. Furthermore, working with the Lévy measure and its series representation is particularly useful for modeling and simulating dependence across multiple processes flexibly without imposing the extreme cases of perfectly dependent or independent arrival times. [Cont and Tankov (2004), Chapter 5, note the importance of working directly with the Lévy measure in the multivariate case.] Thus, the series-based methods developed here provide a way of exploring a much wider variety of interesting Lévy processes for financial applications.

3.1 General Theory

The fundamental result, which we make use of throughout the paper, is the generalized shot noise method for series representation of infinitely divisible distributions, introduced in Rosiński (2001). Here we state in a theorem, the implications for the series representation of Lévy processes. For the proof of this theorem, the reader is referred to Rosiński (2001).

Theorem 1 *Let $\{\Gamma_i\}_{i \geq 1}$ be a sequence of arrival times of a standard (unit intensity) Poisson process; i.e., $\{\Gamma_i\}_{i \geq 1}$ is the sequence of partial sums of standard exponential random variables. Also let $\{V_i\}_{i \geq 1}$ be a sequence of i.i.d. random variables with distribution F in a measurable space S and $\{U_i\}_{i \geq 1}$ be a sequence of i.i.d. uniform on $[0, 1]$ random variables such that $\{\Gamma_i\}_{i \geq 1}$, $\{V_i\}_{i \geq 1}$ and $\{U_i\}_{i \geq 1}$ are independent of each other. Let $H(r, v)$ be a measurable function $H : (0, \infty) \times S \rightarrow \mathbb{R}^d$.*

In addition make the following notation

$$\psi(r, B) = \mathbb{P}[H(r, V_i) \in B], \quad r > 0, \quad B \in \mathcal{B}(\mathbb{R}^d),$$

$$A(s) = \int_0^s \int_{|\mathbf{x}| \leq 1} \mathbf{x} \psi(r, d\mathbf{x}) dr, \quad \text{for } \forall s \geq 0,$$

and assume the measure ν can be decomposed as

$$\nu(B) = \int_0^\infty \psi(r, B) dr, \quad r > 0, \quad B \in \mathcal{B}(\mathbb{R}^d). \quad (4)$$

(a) If $\mathbf{a} := \lim_{s \rightarrow \infty} A(s)$ exists in \mathbb{R}^d and ν is a valid Lévy measure, then

$$\sum_{i=1}^{\infty} H(\Gamma_i, V_i) 1_{(U_i \leq t)} \quad (5)$$

converges almost surely, uniformly in t on $[0, 1]$ to a Lévy process with characteristic triplet $(\mathbf{0}, \mathbf{a}, \nu)$.

(b) If ν is a valid Lévy measure, $|H(r, \nu)|$ is a nonincreasing in r and $c_i = A(i) - A(i-1)$, then

$$\sum_{i=1}^{\infty} (H(\Gamma_i, V_i) 1_{(U_i \leq t)} - tc_i) \quad (6)$$

converges almost surely, uniformly in t on the interval $[0, 1]$ to a Lévy process with characteristic triplet $(\mathbf{0}, \mathbf{0}, \nu)$.

The theorem gives a way to represent pure jump Lévy processes and integrals with respect to them. All we need to do is find a shot noise decomposition of the corresponding Lévy measure. Many of the existing methods for simulation of Lévy processes, such as the Inverse Lévy Measure method of Khintchine (Ferguson and Klass (1972)) and the standard way of simulating compound Poisson processes, are special cases of shot noise decomposition with particular choice for $H(r, \nu)$ and ν . More examples are in Rosiński (2001). Every Lévy process can have different shot noise decompositions of its Lévy measure. The convenience of every method will be determined by how easy it is to simulate from the distribution F and how fast we can evaluate $H(r, \nu)$. When the process is of finite variation the first condition in (a) is automatically satisfied. However notice that this condition could be also satisfied for processes of infinite variation: when their Lévy measure is symmetric (which is the case for type G Lévy processes-processes that could be presented as Brownian motions subordinated by nonnegative Lévy processes, see Rosiński (1991) for their

series representation). In case the first condition in part (a) of the Theorem fails it is still possible to derive series representation for the Lévy process, but in this case the sum needs to be centered appropriately as it is shown in part (b).

The shot noise decomposition of the Lévy measure in Theorem 1 could be used for deriving the series representation of integrals with respect to Lévy processes. In this paper we will be interested in the following integrals

$$X(t) = \int_0^t f(s-)dL(s), \quad (7)$$

where $f : \mathbb{R}^+ \rightarrow \mathbb{R}$ is a predictable and bounded deterministic function and $L(t)$ is a one dimensional Lévy process of finite variation with the following characteristic function

$$\mathbb{E}(e^{i\varepsilon L(t)}) = \exp \left(t \int_{\mathbb{R}_0} (e^{i\varepsilon y} - 1)\nu(dy) \right), \quad (8)$$

where ν satisfies equation (4). Using Theorem 1 we define the following approximations of the Lévy process $L(t)$ and the integral with respect to it $X(t)$

$$L^\tau(t) = \sum_{\Gamma_i \leq \tau} H(\Gamma_i, V_i)1_{(U_i \leq t)}, \quad (9)$$

$$X^\tau(t) = \sum_{\Gamma_i \leq \tau} f(U_i)H(\Gamma_i, V_i)1_{(U_i \leq t)}. \quad (10)$$

Using the result in Theorem 1, it is easy to see that the approximations $L^\tau(t)$ and $X^\tau(t)$ converge almost surely and uniformly in t on the interval $[0, 1]$ to $L(t)$ and $X(t)$ respectively.

The approximations $L^\tau(t)$ and $X^\tau(t)$ involve random number of terms (but on average they will be τ). This way of truncating the infinite series in (5) has the advantage that the approximation $L^\tau(t)$ is itself a Lévy process of finite activity with Lévy measure $\int_0^\tau \psi(r, \cdot)dr$.

In general the approximation given in equation (9) does not necessarily truncate the small jumps. However when $H(r, v)$ is nonincreasing in r , this "on average" will be true. That is the approximation error, which is itself a jump process will allocate less and less mass on bigger jumps.

The approximation error in the case when the truncation involves cutting exactly the small jumps of the Lévy process (which will be the true in the case of the Inverse Lévy Measure method) was analyzed in Asmussen and Rosiński (2001); see also the work of Wiktorsson (2002) for the integrals with respect to G-type Lévy processes.

3.2 Practical Implementation and Numerical Error

We now show the implementation of the above Lévy shot noise decomposition for the relatively simple case of a gamma process with Lévy measure given by

$$\nu(dy) = c \frac{e^{-\lambda y}}{y} I(y > 0) dy, \quad (11)$$

where λ is a scale parameter and c controls the overall intensity of the process, which is infinitely active. For the gamma process, a very convenient choice for the function $H(\Gamma, V)$ in the representation (5) is

$$H(\Gamma, V) = \frac{V}{\lambda} e^{-\Gamma/c}.$$

We illustrate how to approximate a realization $L(t)$, $t \in [0, 1]$, and the integral

$$X(t) = \int_0^t e^{-\rho s} dL(s), \quad t \in [0, 1],$$

where the kernel function in (7) is $f(s) = e^{-\rho s}$.

Table 1 shows an actual working **FORTRAN 90** code segment written to look like pseudo-code easily translatable into another language. For brevity, the table only shows the innermost part of the program that actually generates the approximations $L^\tau(t)$ and $X^\tau(t)$ in (9) and (10); the code segment is embedded in a larger main program available upon request. Keep in mind that both $L^\tau(t)$ and $X^\tau(t)$ are defined over all $t \in [0, 1]$, but on a computer it is only possible to evaluate each over a finite number of points. After execution of the code segment in Table 1, the arrays **levy** and **x** consist of $L^\tau(t_j)$, and $X^\tau(t_j)$, respectively, over the equi-spaced grid $t_j = j/N$, $j = 1, 2, \dots, N$, where N is the number of grid points, or bins. The process of computing the approximations simply consists of adding the shot noises (or weighted shot noises) to the appropriate “bins” corresponding to the times in $[0, 1]$ where the jumps occur.

Given values of c , λ , and the cutoff τ , along with various control parameters, the code segment works as follows. The first main loop in Table 1 generates a random number **nshot** of shot noises, **H(i)**, and integers, **bin(i)**, corresponding to the jump times, where **nshot** is the cutoff τ on average. The middle part of the code segment is just some checking that sufficient space was allocated and some initialization. The final loops accumulate the shot noises to form **levy(j)** = $L^\tau(t_j)$ and the weighted shot noises to form **x(j)** = $X^\tau(t_j)$, $j = 1, 2, \dots, N$.

A potential application is simulated method of moments estimation, so we assess

accuracy in that context. In particular, consider estimation via simulation of

$$\mathbb{E} \left[\int_0^1 X(s) ds \right] \tag{12}$$

$$\mathbb{E} \left[\int_0^1 X(s)^2 ds \right]. \tag{13}$$

The simulation-based estimator would be obtained by generating many replicates of

$$\frac{1}{N} \sum_{j=1}^N X^\tau(t_j) \tag{14}$$

$$\frac{1}{N} \sum_{j=1}^N [X^\tau(t_j)]^2 \tag{15}$$

using the code displayed in Table 1, and then averaging the replications. Figure 1 displays the relative error, (expected - simulation average)/expected, as a function of τ , for $\lambda = 1.0$, ρ set so that the half-life is 0.05, and two values $c = 5.0$ and $c = 15.0$. Closed form expressions are available for the expected values in (12) and (13) and we used a grid size of $N = 2000$ and 10,000 replications of (14) and (15). Figure 1 suggests that quite accurate approximations can be obtained with rather modest values of τ , although the minimum τ required for a given level of accuracy is higher for the more active process. The figure also indicates that the minimum required value of τ can in practice be determined by trying different values and looking for the smallest value for which the computations stabilize.

4 Mixture of Gammas Lévy-Processes

4.1 The One Dimensional Case

The one dimensional Lévy processes that we will be using in this paper are an extension of the gamma process (11). Our mixture of gammas process is a pure jump, infinitely activity and finite variation Lévy process, whose Lévy measure is given by

$$\begin{aligned} \nu(dy) = & \left(\frac{c_1 e^{-\lambda_1 |y|}}{|y|} + \frac{c_2 e^{-\lambda_2 |y|}}{|y|} + \dots + \frac{c_m e^{-\lambda_m |y|}}{|y|} \right) I(y < 0) dy \\ & + \left(\frac{c_{m+1} e^{-\lambda_{m+1} |y|}}{|y|} + \dots + \frac{c_{m+n} e^{-\lambda_{m+n} |y|}}{|y|} \right) I(y > 0) dy \end{aligned} \tag{16}$$

where c_1, c_2, \dots, c_{m+n} and $\lambda_1, \lambda_2, \dots, \lambda_{m+n}$ are positive real numbers. Note that in this case, unlike the gamma process, we allow for the possibility of negative jumps. This Lévy process is a superposition of the pure jump Lévy processes used in Carr, Geman, Madan, and Yor (2002, 2003) with tilting parameter set to zero.

For the mixture of gammas we can use as the shot noise in the series representation the function

$$H(\Gamma_i, V_i, J_i) = e^{-\Gamma_i/c} J_i V_i \quad (17)$$

where V_i are i.i.d. standard exponential, $c = c_1 + \dots + c_{m+n}$, Γ_i are the arrival times of standard Poisson process, and J_i are i.i.d random variables such that

$$J_i = \begin{cases} -\lambda_1^{-1} & \text{with probability } c_1/c \\ \vdots & \\ -\lambda_m^{-1} & \text{with probability } c_m/c \\ \lambda_{n+1}^{-1} & \text{with probability } c_{m+1}/c \\ \vdots & \\ \lambda_{m+n}^{-1} & \text{with probability } c_{m+n}/c . \end{cases}$$

Verification of this claim follows from the fact that the measure given in equation (16) is easily shown to be a valid Lévy measure (that is it satisfies the integrability condition (3)) and it is of finite variation. The rest of the proof follows directly from Theorem 1, part (a). Using the proposed shot noise (17) above and equations (9) and (10) we can simulate mixture of gammas and moving averages of them.

4.2 A Bivariate Mixture of Gammas

To accommodate the leverage effect in the stochastic volatility model, we will need a two dimensional pure jump Lévy process whose individual processes are linked. One process pertains to the price and the other to the volatility. In order to capture the possibility of common jumps and separate jumps, we consider the joint Lévy measure. Intuitively, if the two processes jump almost always together and their jumps are in constant ratio, this will manifest itself in the Lévy measure being concentrated predominantly on the set: $\{(x, y) : x = \text{constant} \times y\}$. This is an example of completely dependent Lévy processes such as those used in Barndorff-Nielsen and Shephard (2001b) or Carr and Wu (2004), among others. Note that the processes jumping in a fixed proportion is one of the many possible instances of complete dependence. Another one, for example, will be if the jumps of the one process are recovered from the jumps of the other process by raising them to the power three. If the processes comprising a two dimensional Lévy process never jump together, this will imply that the Lévy measure will be concentrated on the set $\{(x, y) : xy = 0\}$. In this case the two Lévy processes are independent.

An approach to capture the dependence would be to use the Lévy copula introduced by Tankov (2003) and Cont and Tankov (2004), and studied further by Barndorff-Nielsen and Lindner (2004). A possible advantage of this approach is that the dependence between the individual processes captured by the Lévy copula is disentangled from the marginal distributions. However, the simulation from Lévy processes linked by a Lévy copula is not so easy, since it employs the Inverse Lévy Measure method, and the inverse of the tail of the Lévy measure is known in closed form in very few cases.

In order to bridge between the two extremes of complete dependence and independence, we model directly the joint Lévy measure. Theorem 1 is already stated in multidimensional framework. All we need to do is to specify the Lévy measure in a way that allows us to capture parsimoniously the dependence structure.

We propose the following multidimensional Lévy measure ν

$$\nu(A) = \int_{\mathbb{R}_0^d} \int_0^\infty \int_0^\infty I_A(e^{-r}vz)e^{-v}drdv\mu(dz), \quad (18)$$

where the measure μ on \mathbb{R}_0^d satisfies

$$\int_{\mathbb{R}_0^d} |z|\mu(dz) < \infty. \quad (19)$$

It is easy to verify that when the condition in (19) is satisfied, the measure in (18) is a valid Lévy measure of finite variation, but infinite activity. It could be shown that this measure is the same as the tempered stable measure introduced in Rosiński (2002) for value of the tempering parameter zero. However note that the analysis in Rosiński (2002) is restricted only to the case of tempering parameter taking values in the interval $(0, 2)$ since only in this case the measure could be analyzed as tempering of a stable process, see the discussion in Rosiński (2002, 2004).

We specialize the measure μ to the following sums of atoms in \mathbb{R}_0^2

$$\mu = c_1\delta_{(\lambda_{11}^{-1}, 0)} + c_2\delta_{(0, \lambda_{22}^{-1})} + c_3\delta_{(\lambda_{13}^{-1}, \lambda_{23}^{-1})} + \dots + c_p\delta_{(\lambda_{1m}^{-1}, \lambda_{2m}^{-1})}, \quad (20)$$

and c_1, \dots, c_m are nonnegative numbers. This could be viewed as a generalization of the mixture of gamma processes to the bivariate case. The λ 's may be of either sign, as there is no way in the bivariate case to separate out the positive and negative jumps as in (16) above. Note that $\lambda_{2i}/\lambda_{1i}$ determines the ratio of the jumps.

We can enforce complete independence of the two individual processes with the measure

$$\mu = c_1\delta_{(\lambda_{11}^{-1}, 0)} + c_2\delta_{(0, \lambda_{22}^{-1})}. \quad (21)$$

In this case, the measure will be concentrated on the set $\{(x, y) : xy = 0\}$. The marginal distribution will be gamma processes. An example of two independent gamma processes are shown on the top panel of Figure 2.

The other extreme of complete dependence, where the jumps of the two individual processes are in fixed proportion, can be achieved with the measure

$$\mu = c\delta_{(\lambda_{13}^{-1}, \lambda_{23}^{-1})}. \quad (22)$$

In this case the marginal distributions are again gamma processes. This case is illustrated on the middle panel of Figure 2.

In general, we want to permit patterns of dependence between the two individual Lévy processes that is flexible and admits complete independence and dependence as special cases. Following the reasoning of the one dimensional case, we can specify a series representation in the two dimensional case where the Lévy measure is specified in equation (18). In particular, $H(\Gamma_i, V_i, J_i)$ is given by

$$H(\Gamma_i, V_i, J_i) = e^{-\Gamma_i/c} V_i J_i \quad (23)$$

where Γ_i are the arrival times of a standard Poisson process, V_i are i.i.d. standard exponential variables, $c = c_1 + c_2 + \dots + c_m$, and J_i are i.i.d. random vectors in \mathbb{R}^2 such that

$$J_i = \begin{cases} (\lambda_{11}^{-1}, 0) & \text{with probability } c_1/c \\ (0, \lambda_{22}^{-1}) & \text{with probability } c_2/c \\ \vdots & \\ (\lambda_{1m}^{-1}, \lambda_{2m}^{-1}) & \text{with probability } c_m/c. \end{cases} \quad (24)$$

This shot noise decomposition of the Lévy measure, combined with equations (9) and (10), gives us a way to simulate the two dimensional mixture of gammas proposed here as well as moving averages of it. The bottom panel of Figure 2 illustrates an example of this general case with two subordinators that can jump independently or in two fixed proportions. The measure μ in (18) need not be atomistic, and the setup is flexible enough to accommodate more general dependent structures of higher dimension than two.

5 Lévy driven Stochastic Volatility Models

5.1 The Generic Stochastic Volatility Model

In this subsection we define the generic Lévy driven stochastic volatility model. As it will be shown, this generic model nests most of the stochastic volatility models analyzed in the finance literature.

The scaled logarithm of the price at time t of the financial price will be denoted by $p(t)$; the scaling of the logarithm is often used so the increments represent a geometric return. The general Lévy driven stochastic volatility model is written as

$$dp(t) = \mu_p(t)dt + \sigma(t-)dL_{p1}(t) + dL_{p2}(t) \quad (25)$$

$$\sigma^2(t) = h(\mathbf{c}'\mathbf{X}(t)) \quad (26)$$

$$d\mathbf{X}(t) = \mathbf{a}(\mathbf{X}(t-), t)dt + \mathbf{b}(\mathbf{X}(t-), t)d\mathbf{L}_X(t) \quad (27)$$

where $L_{p1}(t)$ and $L_{p2}(t)$ are two independent Lévy processes, $L_{p2}(t)$ and $L_X(t)$ are potentially dependent Lévy processes, and $h : \mathbb{R} \rightarrow \mathbb{R}^+$. The system of equations (25)-(27) specifies a Markovian type SDE driven by Lévy process. Equation (26) specifies $\sigma^2(t)$ as driven by an n -dimensional vector of factors $X(t)$. The $(n \times 1)$ vector of constants \mathbf{c} defines the weights of the factors in the variance. The process $\sigma^2(t)$ is not the variance of the return process but only part of it. However, the other part, which is coming from the variance of the pure jump component, is not time varying. Throughout, we keep the name stochastic variance for $\sigma^2(t)$. The vector of factors $\mathbf{X}(t)$ driving the stochastic variance of the price process follow a Lévy-driven SDE, specified in equation (27), where $\mathbf{L}_X(t)$ is a k -dimensional Lévy process.

Most of the stochastic volatility models could be viewed as particular cases of this generic model. For example the model introduced by Barndorff-Nielsen and Shephard (2001a) can be obtained by setting $L_X(t)$ to be pure jump Lévy subordinator, $L_{p2}(t)$ to be a centered version of it, $L_{p1}(t)$ to be Brownian motion, and the factor $\mathbf{X}(t)$ to follow Ornstein-Uhlenbeck process.

The time changed Lévy driven stochastic volatility models for the risk neutral dynamics of Carr, Geman, Madan, and Yor (2003) are also embedded in the above, provided the time change is of a Brownian motion. The time change is continuous process, which is locally deterministic with activity rate following an SDE. The activity rate of the economy in these models plays the same role as the volatility here. The models of Carr, Geman, Madan, and Yor (2003) are different however when the time change is of a pure jump Lévy process.

Many of the jump-diffusion models in the finance literature can be nested in the generic stochastic volatility model (25)–(27). In these models, the jumps are rare events and are modeled as finitely active processes. The stochastic volatility is usually driven by Brownian motion, i.e., in the framework here $L_X(t)$ is a Brownian motion. The leverage effect is captured by specifying $L_{p1}(t)$ as another Brownian motion, which is correlated with the ones driving the factors of the volatility. Among others, the affine jump diffusion models of Duffie, Pan, and Singleton (2000) with time homogenous Lévy measure are nested in this framework.

There are three main features of the stock market returns that every stochastic volatility model should address: (1) volatility clustering and persistence with possible jumps in volatility, (2) jumps in the price, and (3) the leverage effect. Here we elaborate on how these features can be captured by the generic stochastic volatility model defined in equations (25)–(27).

1. **Volatility Persistence** The generic stochastic volatility model proposes a multi-factor structure for the spot variance. This enables one to produce different degrees of persistence in the spot volatility and the pure-jump Lévy component generates jumps in volatility.
2. **Jumps in the Price** We observe the price process only on discrete intervals, so it is natural to ask whether we can disentangle the jump part, if it exists, from the diffusion one. The answer to that question is affirmative, by making use of the nonparametric tests proposed by Barndorff-Nielsen and Shephard (2004, 2006). Using these statistics Andersen, Bollerslev, and Diebold (2005) and Huang and Tauchen (2006) find considerable evidence for jumps in exchange rate returns and asset returns respectively. Therefore we model the price process by allowing for jumps in the prices.
3. **Leverage Effect** The leverage effect pertains to a negative correlation between the volatility and the return process. There are different ways it can be captured in the generic stochastic volatility model. It should be kept in mind that the link could be done only through a link between the continuous parts in the price and the variance, through their jump parts or through a combination of those two, because the pure jump Lévy process is always orthogonal to the Brownian motion. One of the most frequently used ways is to let $L_{p1}(t)$ and $L_X(t)$ have Brownian Motions, which are correlated. That is, in this way we can produce the leverage effect through the diffusion parts of the variance and the price process. This

approach was followed by most of the jump-diffusion literature. An alternative modelling approach is to link the jumps in the price and that in the variance, that is, to let $L_{p2}(t)$ and the pure jump part of $L_X(t)$ be dependent pure jump processes. Eraker, Johannes, and Polson (2003) introduce leverage effect through a combination of these two modelling alternatives. Another approach is to capture the leverage effect by linking the pure jump components of $L_{p1}(t)$ and $L_X(t)$. This could also avoid the need for an additional jump component ($L_{p2}(t)$ here) for capturing the leverage effect; since $L_{p1}(t)$ is specified as a process containing jump part, we will have jumps in the price. This is similar in spirit to the approach, proposed in Carr and Wu (2004), where the pure jump Lévy process of the price and the instantaneous business activity rate are linked.

5.2 The Lévy-Driven CARMA Stochastic Volatility Model

In this subsection we adapt Barndorff-Nielsen and Shephard (2001a) and Brockwell (2001a,b) to introduce a Lévy-driven CARMA (continuous time autoregressive moving average) stochastic volatility model. The model is nested within the generic stochastic volatility introduced in the previous section, and it accommodates both jumps in the price process and the variance. In fact, the variance process is driven solely by jumps.

The price process is modelled as having a diffusion part and a pure jump component. The pure jump component is a Lévy process, which might be of infinite activity (even infinite variation) and can have positive and negative jumps. The diffusion part of the price displays time-varying variance, and the pure jump part is of constant variance.

The model is

$$dp(t) = \mu_p(t)dt + \sigma(t-)dW(t) + dL_p(t) \quad (28)$$

$$a(D)\sigma_t^2 = b(D)DL_\sigma(t) \quad (29)$$

where $W(t)$ is a standard Brownian Motion, $L_p(t)$ is the pure jump Lévy process in the price, $L_\sigma(t)$ is the Lévy process driving the stochastic variance, D is a differential operator, and $a(D)$ and $b(D)$ are given by

$$\begin{aligned} a(z) &= z^p + a_1 z^{p-1} + \dots + a_p \\ b(z) &= b_0 + b_1 z + \dots + b_q z^q, \quad q < p. \end{aligned} \quad (30)$$

In (29) $\sigma^2(t)$ is a Lévy-driven CARMA(p, q) process of Brockwell (2001a,b). Equivalently we can write

$$\sigma^2(t) = b'X(t), \quad (31)$$

where $X(t)$ is a solution to the SDE

$$dX(t) = \mathbf{A}X(t)dt + \mathbf{e}dL(t) \quad (32)$$

and where,

$$\mathbf{A} = \begin{pmatrix} 0 & 1 & 0 & \dots & 0 \\ 0 & 0 & 1 & \dots & 0 \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & 0 & \dots & 1 \\ -a_p & -a_{p-1} & -a_{p-2} & \dots & -a_1 \end{pmatrix}, \quad \mathbf{e} = \begin{pmatrix} 0 \\ 0 \\ \vdots \\ 0 \\ 1 \end{pmatrix}, \quad \mathbf{b} = \begin{pmatrix} b_0 \\ b_1 \\ \vdots \\ b_{p-2} \\ b_{p-1} \end{pmatrix}$$

The above shows that the Lévy driven stochastic volatility model is nested in the generic framework introduced in the previous section.

If the eigenvalues of \mathbf{A} , denoted ζ_j , $j = 1, 2, \dots, p$, have negative real parts, i.e., $\text{Re}(\zeta_j) < 0$ for $j = 1, 2, \dots, p$, then the CARMA(p, q) is a causal stationary process (Brockwell (2001b)). In this case, for t big enough the contribution of the starting value is negligible and

$$\sigma_t^2 = \int_0^t g(t-u)dL(u), \quad (33)$$

where $g(h)$ is the kernel for the corresponding CARMA process, and from Brockwell (2001b)

$$g(u) = \frac{1}{2\pi} \int_{-\infty}^{\infty} e^{it\lambda} \frac{b(i\lambda)}{a(i\lambda)} d\lambda. \quad (34)$$

The kernel function $g(u)$ determines the memory of the process $\sigma^2(t)$. It gives the weights with which the past observations enter the Lévy functional. Since we are modelling the spot variance of the return process by a CARMA process, we require its kernel to be nonnegative everywhere. The CARMA approach provides a rich variety of processes that generalize the Lévy-driven OU case (a CARMA(1, 0)) analyzed by Barndorff-Nielsen and Shephard (2001b) and also used by Nicolato and Venardos (2003) for options pricing.

Farther below in Subsection 6.2 we concentrate on a CARMA(2, 1) process for the stochastic variance, which we parameterize as

$$a(z) = (z - \rho_1)(z - \rho_2), \quad b(z) = 1 + b_1z,$$

for real $\rho_1 < 0$ and $\rho_2 < 0$, $\rho_1 \neq \rho_2$. The kernel is

$$g(h) = \frac{1 + b_1\rho_1}{\rho_1 - \rho_2} e^{\rho_1 h} + \frac{1 + b_1\rho_2}{\rho_2 - \rho_1} e^{\rho_2 h}, \quad h \geq 0 \quad (35)$$

A necessary and sufficient condition that guarantees nonnegativity of the kernel is $0 \leq b_1 \leq \max\{-1/\rho_1, -1/\rho_2\}$ as shown in the Appendix.

The kernel, given in equation (35) will be decreasing for $\forall h \geq 0$ if

$$b_1 \in \left[-\frac{1}{\rho_1 + \rho_2}, \max\left(-\frac{1}{\rho_1}, -\frac{1}{\rho_2}\right) \right],$$

while for

$$b_1 \in \left(0, -\frac{1}{\rho_1 + \rho_2} \right)$$

the kernel increases initially and reaches a maximum and decreases afterwards.

This Lévy-driven CARMA(2,1) setup accommodates the three key properties defined in Subsection 5.1 as follows:

- volatility persistence and ability of the volatility to move quickly

The studies of Chernov, Gallant, Ghysels, and Tauchen (2003) and Alizadeh, Brandt, and Diebold (2002) argue in favor of a two factor structure in the volatility process. One of the factors should slowly mean revert, allowing for the volatility persistence. The other factor on the other hand should be quickly mean reverting in order to allow the volatility to move quickly. Motivated by these two factor structure of the volatility, we propose here following Brockwell (2001b) a Lévy driven CARMA(2,1) model for the stochastic variance. One of the autoregressive roots is high in magnitude, corresponding to quick mean reversion, while the other one is low in magnitude and thus allowing for slow mean reversion. This structure of the autoregressive part of the CARMA(2,1) allows for pretty flexible autocorrelation in the spot variance process. Furthermore, since the driving Lévy process of the CARMA(2,1) process is a pure jump process, the Lévy driven CARMA(2,1) process could naturally produce the desired effect of quick moves in the variance as argued in Alizadeh, Brandt, and Diebold (2002). An alternative approach, which resembles the diffusion factor stochastic volatility models analyzed by Chernov, Gallant, Ghysels, and Tauchen (2003), is the superposition of Lévy driven Ornstein-Uhlenbeck processes as proposed by Barndorff-Nielsen and Shephard (2001a) and further evaluated empirically in the context of subordinated Levy processes without leverage by Barndorff-Nielsen and Shephard (2005a). This alternative allows more flexible autocorrelation structures for the spot variance over the exponential one implied by the single Ornstein-Uhlenbeck process in a similar way the CARMA(2,1) kernel does. The potential advantage of CARMA modelling is even more flexible autocorrelation structures and with only a single driving process.

- jumps in the price

The model proposed here, allows for jumps in the price process. In addition the price has a diffusion component in it. This makes this model being different from the models analyzed in Carr, Geman, Madan, and Yor (2003), where the price has no diffusion component in it. The modelling here resembles the jump-diffusion models. However unlike the jump-diffusion models, the jumps in the price need not be rare events, but can be infinitely many in any finite time interval.

- leverage effect

Since in our model the variance is driven solely by a pure jump process, $L_\sigma(t)$, the only way we can capture the leverage effect is by linking the two jump components $L_\sigma(t)$ and $L_p(t)$. That is the leverage effect in our model is captured by the jumps in the price process. Note that in our model the jumps in both the variance and in the price need not be rare events (as in the jump-diffusion models). There could be infinite number of small jumps in a given interval of time. Thus, the model here allows for a link between variance and returns, not only in the case of extreme events of big changes in the price, but also for small changes in the price and variance. A similar approach of modelling the leverage effect is taken in Barndorff-Nielsen and Shephard (2001a) and Carr, Geman, Madan, and Yor (2003). In these models however, the jump components in the price and in the variance are perfectly dependent. In terms of the notation adopted here in these models $L_p(t) = \text{constant} \times L_\sigma(t)$. This implies, that the jumps in the variance and in the price process arrive at the same time and are proportional. Here we relax the dependence structure between the jumps in the price and in the variance by using the two dimensional generalization of the mixture of gammas proposed in Subsection 4.2. As it was discussed in that section we allow for various degrees of dependence between the two jump processes: in some cases the jumps arrive at the same time, while in other cases they arrive in different times. Also in the case, the jumps arrive at the same time, the jump sizes can be in different proportions.

6 Examples

In this section we will demonstrate through several examples that the proposed stochastic volatility model can produce reasonable dynamics. In practice we observe, the price process in discrete intervals. For simplicity here we assume that the observational intervals are equally spaced. We denote with a the length of the observational time.

Throughout the paper our time is measured in number of trading days. In the examples below we work with half-hour returns and assuming eight hour trading day, we will have then $a = 1/16$.

6.1 Simulating from Lévy driven CARMA SV Models

The stochastic volatility model specified in equations (28)–(29), implies that the geometric return, $r_a(t)$, over the interval $(t - a, t]$ is

$$r_a(t) = p(t) - p(t - a) = \int_{t-a}^t \sigma(s-) dW(s) + L_p(t) - L_p(t - a), \quad (36)$$

where for now we assume the drift in (28) is zero. Using the fact that $\int_{t-a}^t \sigma(s-) dW(s)$ is Gaussian conditional on the pure jump processes in the price and the variance, and under the Lévy-driven model the variance process has no fixed time of discontinuity ($\sigma^2(s) = \sigma^2(s-)$ almost surely), we can write

$$p(t) - p(t - a) \stackrel{d}{=} Z_t \sqrt{\int_{t-a}^t \sigma^2(s) ds} + L_p(t) - L_p(t - a), \quad (37)$$

where $Z(t)$ is standard normal distribution independent of $L_p(t)$ and $L_\sigma(t)$ and $\int_{t-a}^t \sigma^2(s) ds$ is the integrated variance over the time interval a .

Applying the Fubini theorem to the general case (33) the integrated variance is

$$\begin{aligned} \int_{t-a}^t \sigma^2(s) ds &= \int_{t-a}^t \int_0^s g(s-u) dL(u) ds \\ &= \int_{t-a}^t \int_u^t g(s-u) ds dL(u) + \int_0^{t-a} \int_{t-a}^t g(s-u) ds dL(u) \\ &= \int_0^t g^*(t, u) dL(u) \end{aligned} \quad (38)$$

where the functional form of g^* can be obtained from that of g . In the case of CARMA(2,1)

$$g^*(t, u) = \begin{cases} (e^{\rho_1(t-u)} - e^{\rho_1(t-a-u)}) \frac{1+b_1\rho_1}{\rho_1(\rho_1-\rho_2)} + (e^{\rho_2(t-u)} - e^{\rho_2(t-a-u)}) \frac{1+b_1\rho_2}{\rho_2(\rho_2-\rho_1)} & \text{if } 0 < u < t - a \\ (e^{\rho_1(t-u)} - 1) \frac{1+b_1\rho_1}{\rho_1(\rho_1-\rho_2)} + (e^{\rho_2(t-u)} - 1) \frac{1+b_1\rho_2}{\rho_2(\rho_2-\rho_1)} & \text{if } t - a \leq u < t \end{cases} \quad (39)$$

The displayed equations in (38) can be viewed as an extension to the general case of the observation by Barndorff-Nielsen and Shephard (2001b,c) that in the non-Gaussian OU model the integrated variance is linear in the Lévy-innovations, and (39) gives the functional form for the CARMA(2,1) model. The proposed stochastic volatility model suggests a relatively easy way to simulate from it using the results in Section 3.

Suppose we want to simulate from the price process at intervals with length a . Conditional on the integrated variance and the increment from the pure jump component of the price, the price increment is Gaussian with variance equal to the integrated variance. Therefore once we develop a way to jointly simulate from the integrated variance and the pure jump component of the price, the simulation of the price increments is trivial. The generic simulation scheme is

1. Simulate jointly from the two pure jump Lévy processes $L_p(t)$ and $L_\sigma(t)$.
2. Generate the implied integrated variance by using (39).
3. Simulate the price increment by drawing from a normal distribution with mean zero and variance equal to the integrated variance computed in the preceding step and adding the increments of $L_p(t)$ that occur within the interval $(t - a, t]$.

Given the discussion in Section 3 we know that we are simulating approximately from the model specified above and exactly from a stochastic volatility model with pure jump Lévy processes $L_p(t)$ and $L_\sigma(t)$ substituted with high activity compound Poisson processes, which very closely resemble the infinite activity ones. Also note that the proposed scheme of simulation does not involve any Euler discretization. Finally, it generalizes in a direct way if there is a volatility risk premium, and the drift in (28) is an affine function of the spot variance, since the cumulated drift is thereby an affine function of the integrated variance.

6.2 The MG-CARMA(2,1) Model

To illustrate the properties of the proposed simulation schemes, we show the implied dynamics of an MG-CARMA(2, 1), which stands for a mixture of gammas Lévy-process that jointly drives a pure jump in the price and a CARMA(2, 1) stochastic volatility specification. We show characteristics of the simulated realizations for four different settings of the parameters as listed in Table 2. In all cases, the simulation was run for 1,000 days (four trading years) with a burn period of 250 days. We work with half hour returns, which correspond to sixteen equally spaced intervals in every trading day. Thus, in each of Figures 3–6 there are a total of 16,000 observations on half hour returns. The tolerance parameter τ is fixed at 20, implying on average 20 jumps are generated per trading day. As seen from Table 2, in each case $c = \sum_i c_i \approx 0.02$; from the analysis of Subsection 3.2 above, the value of $\tau = 20$ is well above that needed for a reasonable level of accuracy for this value of c and much larger values as well. This truncation entails discarding jumps of extremely small magnitude which, in the

presence of continuous component in the price process, are immaterial. In all four cases, we chose the parameter values so that the variance of the daily returns is equal to 1.00, which is about the variance of the daily S&P 500 Index return, and the drift in the price equation is always identically zero.

The parameters governing the two dimensional mixture of gammas process, which is the background driving Lévy-process (BDLP), are set as follows. The number of terms in the mixture of gammas in equation (20) is $m = 4$. The price and volatility jump independently if J_1 or J_2 occur, while they jump together but in different proportions if J_3 or J_4 occur.

For the parameters governing the CARMA(2,1) kernel, the choice was made as follows. The MA parameter b_1 is set equal to 0.001 in all four cases. Following the discussion in Section 5, ρ_2 has a very short half life 0.03 days for all cases, so it corresponds to the fast mean reverting autoregressive root. The other root, ρ_1 , is relatively slow mean reverting. In the first two cases, Figures 3 and 4, the half life is 50 days; in the other two, Figures 5 and 6, the half life is 10 days. The effects of the persistence of the root ρ_1 are quite noticeable in the dynamics of the spot variance. When $\rho_1 = 10$, a jump in the driving Lévy process of the CARMA(2,1) dies off much more quickly than when $\rho_1 = 50$. Looking at the spot variances in Figures 3–6 we can see that the MG-CARMA(2,1) model can produce the persistence of the variance and at the same time its ability to change quickly. This translates automatically in clustering of the volatility, which can be seen most clearly in the return dynamics in the bottom panels of Figures 3–6.

Finally the parameters in the stochastic volatility model were chosen in such a way that the pure jump component of the returns has a fixed proportion in the total variance of the returns: 0.3 in the cases corresponding to Figures 3, 5, and 0.1 in the cases corresponding to Figures 4, 6. Since the variance of the jump component of the returns is not varying, we would expect the clustering of the volatility to be less pronounced if the jumps contribute a higher proportion in the total variance of the returns. Indeed this is the case as can be easily confirmed by comparing Figures 3 with 4, and 5 with 6.

Figures 7–9 provide additional characteristics of the simulated returns corresponding to the case plotted in Figure 6.

Figure 7 shows the realized daily variance, the daily bipower variation of Barndorff-Nielsen and Shephard (2004), and the integrated daily variance. We define the bipower

variation as

$$\frac{\pi}{2} \sum_{j=1}^{16} |r_{t+ja}^a| |r_{t+(j-1)a}^a|, \quad (40)$$

where we work with half hour returns so $a = 1/16$, and the scale factor $\frac{\pi}{2}$ makes the measure directly comparable with the integrated variance. It is well known that the realized daily variance is not a consistent estimator of the integrated variance in the presence of jumps in the price process; the use of squared returns picks up the variation of the jump component along with that of the diffusive component. This is readily confirmed by comparing the top and bottom panels of Figure 7. In the presence of relatively big jumps during a given day, the realized variance is significantly higher than the integrated variance.

The bipower variation is generally (but not always) robust to jumps and provides a consistent estimator of the integrated variance as the sampling interval goes to zero. Barndorff-Nielsen, Shephard, and Winkel (2005) contains sufficient conditions for consistency, which are satisfied by our mixture of gammas. Comparison of the middle and the bottom panels of Figure 7 indicates that, as expected, the bipower variation does an excellent job of tracking the integrated variance despite the infinitely active character of the underlying BDLP.

The top panel of Figure 8 shows the difference between the realized variance and the bipower variation, which is a measure of the pure jump component of the price process. The second panel shows this difference divided by the realized variance, which is the measure of the relative share of jumps in total variance considered in Huang and Tauchen (2006). The underlying mixture of gammas is infinitely active, but in general with just a few large jumps and many small jumps. Comparing the top and middle panels of Figure 8 to the bottom panel of Figure 6 suggests that these jump measures are large on days with large jumps, but the measures are unable to separate the very small jumps from the Brownian component of the price. This contrast suggests that the simulation methods developed in this paper could be used for a far more extended analysis of the properties of the jump-detection tests of Barndorff-Nielsen and Shephard (2006) than that conducted by Huang and Tauchen (2006).

Finally, Figure 9 shows the autocorrelation in the absolute and the squared returns along with the cross correlation of the increment of the price jump process, $L_p(t) - L_p(t-a)$, and the integrated variance $\int_{t-a}^t \sigma_s^2 ds$. The plots extend for 160 lags, which corresponds to 10 trading days. In the top two panels the persistence in the absolute and squared returns is quite evident, even with ρ_1 having a half life of 10 days. The bottom panel shows the leverage effect, which in the this setup is generated by

the negative covariance between price jumps and variance increments in the bivariate gamma mixture model. The contemporaneous correlation is negative and then slowly dissipates in a manner consistent with that of observed data (Bollerslev, Litvinova, and Tauchen (2005), Litvinova (2004), Tauchen (2005)).

6.3 Discussion of Estimation

The continuous time MG-CARMA(2,1) model of the previous subsection can capture many of the known features of financial returns data. Nonetheless, taking such a model directly to very high frequency returns presents serious challenges pertaining to the data and estimation strategy.

As regards the data, the sampling interval cannot be too small, or the returns will be dominated by microstructure noise, which often limits the sampling interval to no finer than five minutes. Complications such as bid-ask bounce overwhelm the information in the data at the highest frequencies. Microstructure noise is an area of intense current research, and a full review is beyond the scope of this paper; a comprehensive recent survey of the issues is in Barndorff-Nielsen and Shephard (2005b). The noise issue notwithstanding, another problem is that return volatility shows a deterministic pattern over the day, with high volatility in the early part of the day, lower in the middle, and higher again towards the later part of the day. Also, the overnight return has a different variance than very short-term returns. Raw returns will therefore have to be adjusted for diurnal and overnight patterns. The result is an adjusted returns series that itself is not actually the return on any traded security and is not sampled as finely as one might think possible.

These data issues might be considered minor and addressable by data transformations, but the estimation problems associated with application to high very frequency returns are truly formidable. As is widely known, the challenge is that the model-implied transition density of returns given past returns is not available in simple closed form. Direct analytical likelihood-based methods are thus intractable. Except in the simplest situations, simulation can be expected to play a role somewhere in order to integrate out the unobserved volatility variable(s). Simulated likelihood in the manner of Durham and Gallant (2002) is applicable only in certain cases, because one generally lacks the joint density of the observed and unobserved variables, so there is no convenient way to define an importance function for efficient simulation. Indirect estimation techniques as discussed by Tauchen (1997) can potentially be applied in this context, but that approach requires a good statistical description of the high frequency returns,

and it is not clear if the extant models are adequate for this purpose.

MCMC in the style of Eraker, Johannes, and Polson (2003) and Li, Wells, and Yu (2004) can potentially be adapted for estimation of models like the MG-CARMA(2,1) but there are obstacles to overcome. With exceptions such as Eraker (2001) and Elerian, Chib, and Shephard (2001), MCMC applications typically assume the sampling interval is the same as a tick on the continuous time clock, which greatly reduces the number of unobserved variables to deal with. In effect, this reframes the task to that of estimating a discrete time stochastic volatility model. By doing so, it is not possible to explore the model's implications for the price series at intervals finer than the sampling interval, and there is discretization bias. Roberts, Papaspiliopoulos, and Delaportas (2004) avoid discretization bias in an MCMC context by using a shot noise decomposition as analyzed here, but only for OU models applied to low frequency data.

Another matter that pertains to any estimation strategy is that of compounding of specification error bias. Regardless of the estimation technique, small specification errors at the highest frequency can be expected to accumulate if the model's dynamics are spun out to the daily, weekly, or monthly intervals.

An alternative route to direct estimation on the high frequency returns is to work at the daily level, while retaining from the high frequency summary measures such as the realized variance and the bipower variation studied in Andersen, Bollerslev, and Diebold (2005) and others. An underlying continuous time model, e.g., the MG-CARMA(2,1) above, could be forced to confront the dynamics of the vector of the daily summary variables. One empirical effort in this style is Barndorff-Nielsen and Shephard (2005a), who use quasi-maximum-likelihood to generate semi-parametric estimates of continuous time volatility dynamics but do not generate estimates of the entire continuous time law of motion of the price process. In special cases, the law of motion might be estimable using likelihood-based techniques, but in general computation of the transition density of the daily summary measures is intractable. Simulated method of moments, or indirect estimation techniques, appear more directly applicable to the task, though that approach requires a good statistical description of the dynamics of the daily summary measures. A starting point in this direction is Bollerslev, Kretschmer, Pigorsch, and Tauchen (2005), which is an extensive statistical modeling effort of its own. In on-going work, the authors and other researchers are in the initial stages of implementing both method of moments-type estimators and simulated indirect estimation methods using the daily summary statistics.

Development of an appropriate estimation strategy that effectively uses the infor-

mation in the large high frequency data sets is topic of current research. Regardless of the strategy, however, it appears clear that simulation schemes from more general Lévy processes such as those above can be expected to play a central role.

7 Conclusion

We have developed simulation schemes for the new classes of non-Gaussian pure jump Lévy processes for stochastic volatility. We showed how to write the price and volatility processes as integrals against a vector Lévy process, which then makes the series approximation methods directly applicable. These methods entail simulation of the Lévy increments and formation of weighted sums of the increments; they do not require a closed-form expression for a tail mass function nor specification of a copula function. We have also presented a new, and apparently quite flexible, bivariate mixture of gammas model for the driving Lévy process. Within this setup, it is quite straightforward to generate simulations from a Lévy-driven CARMA stochastic volatility model augmented by a pure-jump price component. The simulations reveal the wide range of different types of financial price processes that can be generated in this manner. By appropriate choice of parameters, the resulting simulated price series displays many of the same features of observed price data, including persistent stochastic volatility, dynamic leverage, and jumps.

Appendix

A Necessary and Sufficient Condition for Non-negativity of the CARMA(2, 1) Kernel

The CARMA(2,1) kernel is given by

$$g(h) = \frac{1 + b_1\rho_1}{\rho_1 - \rho_2} e^{\rho_1 h} + \frac{1 + b_1\rho_2}{\rho_2 - \rho_1} e^{\rho_2 h}, \quad h \geq 0.$$

We want to find a condition on the parameters ρ_1 , ρ_2 and b_1 , which will guarantee the nonnegativity of the kernel for every $h \geq 0$. We assume $\rho_2 < \rho_1 < 0$ without loss of generality. Then this is equivalent to conditions on the parameters such that

$$(1 + b_1\rho_1)e^{\rho_1 h} \geq (1 + b_1\rho_2)e^{\rho_2 h}, \quad h \geq 0.$$

First note that this inequality will be violated for $b_1 < 0$ for values of h zero and sufficiently close to zero. Also the inequality is trivially satisfied for $b_1 = 0$. We concentrate on giving the necessary and sufficient condition for the nonnegativity of the CARMA(2,1) kernel in the case $b_1 > 0$.

The first derivative w.r.t. ρ of the function $f(\rho) = (1 + b_1\rho)e^{\rho h}$ is

$$f'(\rho) = e^{\rho h}(h + b_1\rho h + b_1).$$

We have three cases

1. $b_1\rho_1 + 1 > 0$ and $b_1\rho_2 + 1 \geq 0$. In this case $f(\rho)$ is increasing in the interval $[\rho_2, \rho_1]$ for every $h \geq 0$, which implies the positivity of the kernel.
2. $b_1\rho_1 + 1 \geq 0$ and $b_1\rho_2 + 1 < 0$. In this case the kernel is trivially nonnegative.
3. $b_1\rho_1 + 1 < 0$ and $b_1\rho_2 + 1 < 0$. In this case there exist values for $h \geq 0$ such that $f(\rho)$ is decreasing in the interval $[\rho_2, \rho_1]$ which will produce negative values for the kernel for those h .

Taken together, the results above imply the following necessary and sufficient condition for the nonnegativity of the CARMA(2,1) kernel: $0 \leq b_1 \leq \max\{-\frac{1}{\rho_1}, -\frac{1}{\rho_2}\}$.

Table 1: Working Code

```

* Parameters c, lambda, rho, and
* control values tau, N, seed, SIMMAX
* already initialized above.
*
* The function ran generates uniform random variables.
* -----
* SHOT NOISE GENERATION
* -----
    gamlag = 0.0d0
    i=0
    do while ( (gamlag .le. tau) .and. (i .le. SIMMAX) )
        i=i+1
        ur      = ran(seed)           ! generate uniform
        gam(i)  = gamlag - DLOG(ur)   ! increment is exp(1)
        gamlag  = gam(i)             ! save lagged value
        ur      = ran(seed)           ! another uniform
        v       = -DLOG(ur)          ! v is exp(1)
        H(i)    = (1.0/lambda)*DEXP(-gam(i)/c)*v ! ith shot noise
        u(i)    = ran(seed)           ! uniform for jump time
        bin(i)  = INT( 1.0d0 + u(i)*N ) ! bin number of jump
        nshot   = i                  ! number of shot noises
    enddo
* -----
*
* Check to see if we ran out of space before gam(i)>tau
    if ( (gam(nshot) .lt. tau) .and. (nshot .eq. SIMMAX) ) then
        write(*,'(a)') 'PROBLEM: SIMMAX too small for this seed,HALT'
        stop
    endif
*
* Initialize
    do j=1,N
        levy(j) = 0.0
        x(j)    = 0.0
        tbin(j) = DFLOAT(j)/DFLOAT(N)
    enddo
*
* -----
* DISTRIBUTION OF THE SHOT NOISES
* -----
    do i=1,nshot
        do j=bin(i),N
            levy(j) = levy(j) + H(i)
            x(j)    = x(j) + DEXP(-rho*(tbin(j) - u(i)))*H(i)
        enddo
    enddo
* -----

```

Table 2: Parameter Settings for MG-CARMA(2, 1) Stochastic Volatility Model

Parameter	Figure 3	Figure 4	Figure 5	Figure 6
Half life of ρ_1 (days)	50.0	50.0	10.0	10.0
Half life of ρ_2 (days)	0.0300	0.0300	0.0300	0.0300
b_1	0.0100	0.0100	0.0100	0.0100
c_1	0.0060	0.0060	0.0120	0.0120
c_2	0.0015	0.0015	0.0030	0.0030
c_3	0.00025	0.00025	0.0005	0.0005
c_4	0.0045	0.0045	0.0090	0.0090
λ_{11}	0.2600	0.0044	0.3600	0.6200
λ_{13}	-0.0430	-0.0740	-0.0600	-0.1000
λ_{14}	-0.2600	-0.0044	-0.3600	-0.6200
λ_{22}	0.0560	0.0440	0.0220	0.0170
λ_{23}	0.0700	0.0054	0.0028	0.0022
λ_{24}	0.0280	0.0220	0.0110	0.0087

Note: Shown above are the parameter choices for the CARMA(2,1) model (35) and those controlling the distribution of the J_i in (24).

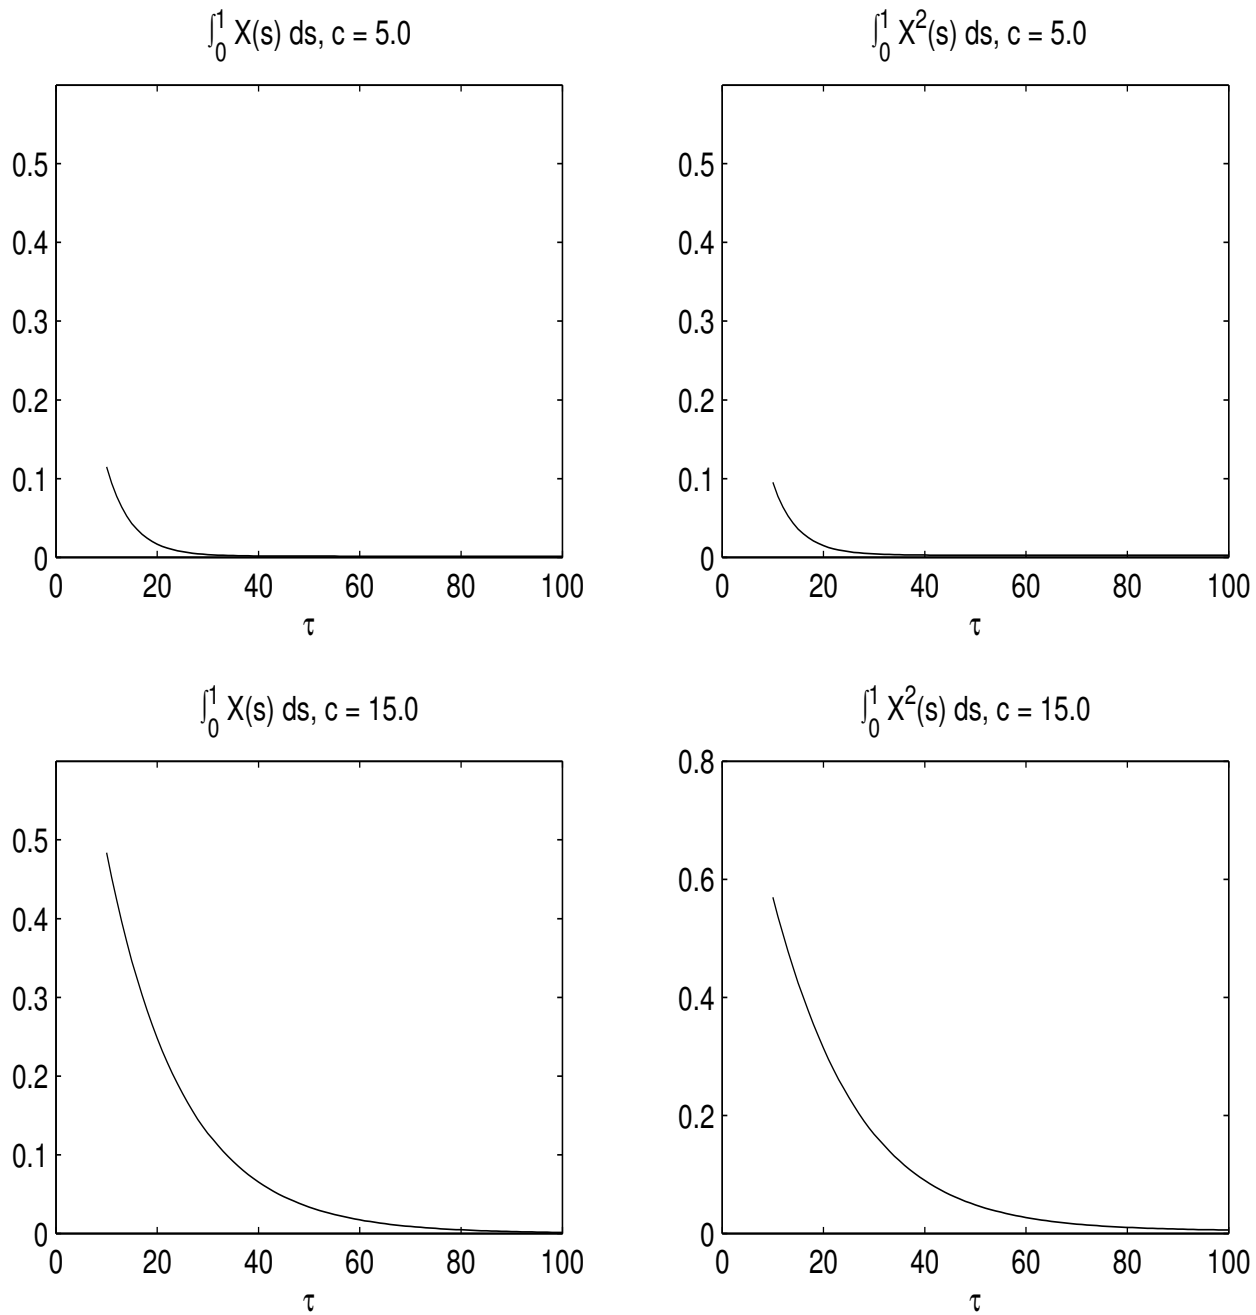


Figure 1: The panels show the relative approximation error as a function of τ in estimating $E \left[\int_0^1 X(s) ds \right]$ and $E \left[\int_0^1 X^2(s) ds \right]$ via simulation.

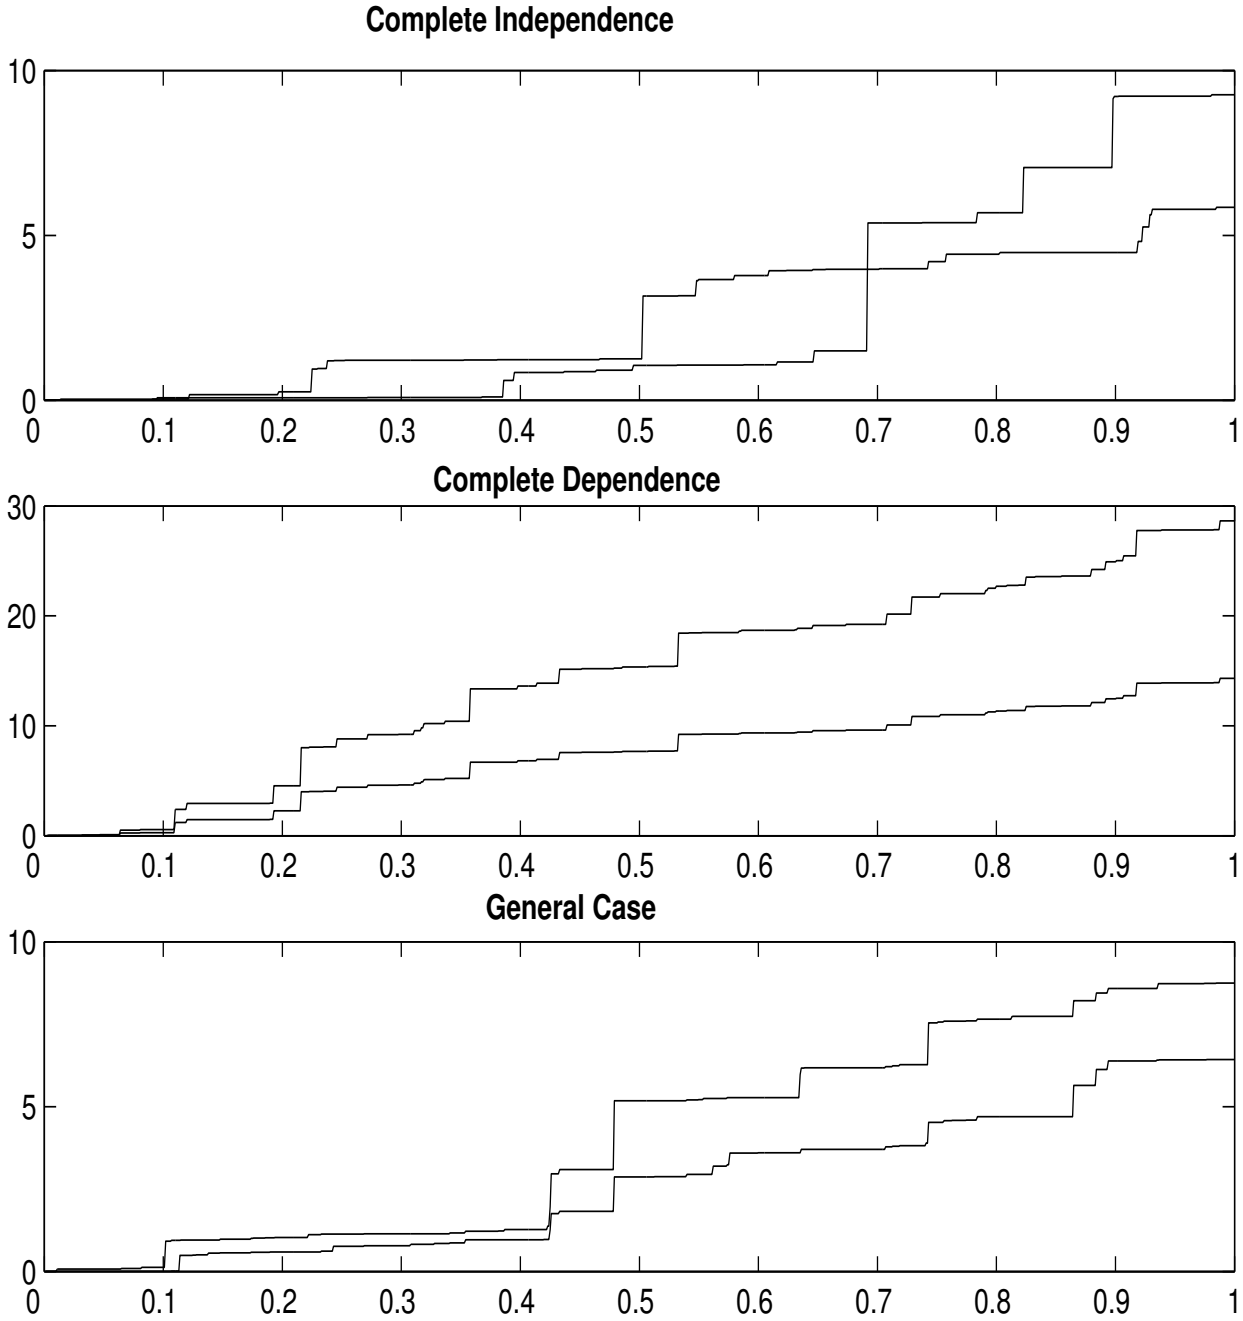


Figure 2: Different degrees of dependence between the two parts of a two dimensional mixture of gamma Lévy process. The top panel shows the complete independence case with parameters $c_1 = c_2 = 5$, $\lambda_{11} = \lambda_{22} = 0.5$; the second one illustrates the complete dependence case with parameters $c_1 = c_2 = 0$, $c_3 = 10$, $\lambda_{31} = 1$ and $\lambda_{32} = 2$; the third panel shows the general case with the following parameters $c_1 = c_2 = c_3 = c_4 = 2.5$, $\lambda_{11} = \lambda_{22} = 1$, $\lambda_{31} = \lambda_{42} = 1$ and $\lambda_{32} = \lambda_{41} = 0.5$.




Figure 3: Simulated Half Hour Realizations from the MG-CARMA(2,1) stochastic volatility model with parameter setting specified in Table 2. The top panel shows the spot variance; the second illustrates the Lévy subordinator driving the variance; the third panel shows the pure jump part of the price process and the bottom one shows the half hour price change.




Figure 4: Simulated Half Hour Realizations from the MG-CARMA(2,1) stochastic volatility model with parameter setting specified in Table 2. The top panel shows the spot variance; the second illustrates the Lévy subordinator driving the variance; the third panel shows the pure jump part of the price process and the bottom one shows the half hour price change.




Figure 5: Simulated Half Hour Realizations from the MG-CARMA(2,1) stochastic volatility model with parameter setting specified in Table 2. The top panel shows the spot variance; the second illustrates the Lévy subordinator driving the variance; the third panel shows the pure jump part of the price process and the bottom one shows the half hour price change.




Figure 6: Simulated Half Hour Realizations from the MG-CARMA(2,1) stochastic volatility model with parameter setting specified in Table 2. The top panel shows the spot variance; the second illustrates the Lévy subordinator driving the variance; the third panel shows the pure jump part of the price process and the bottom one shows the half hour price change.

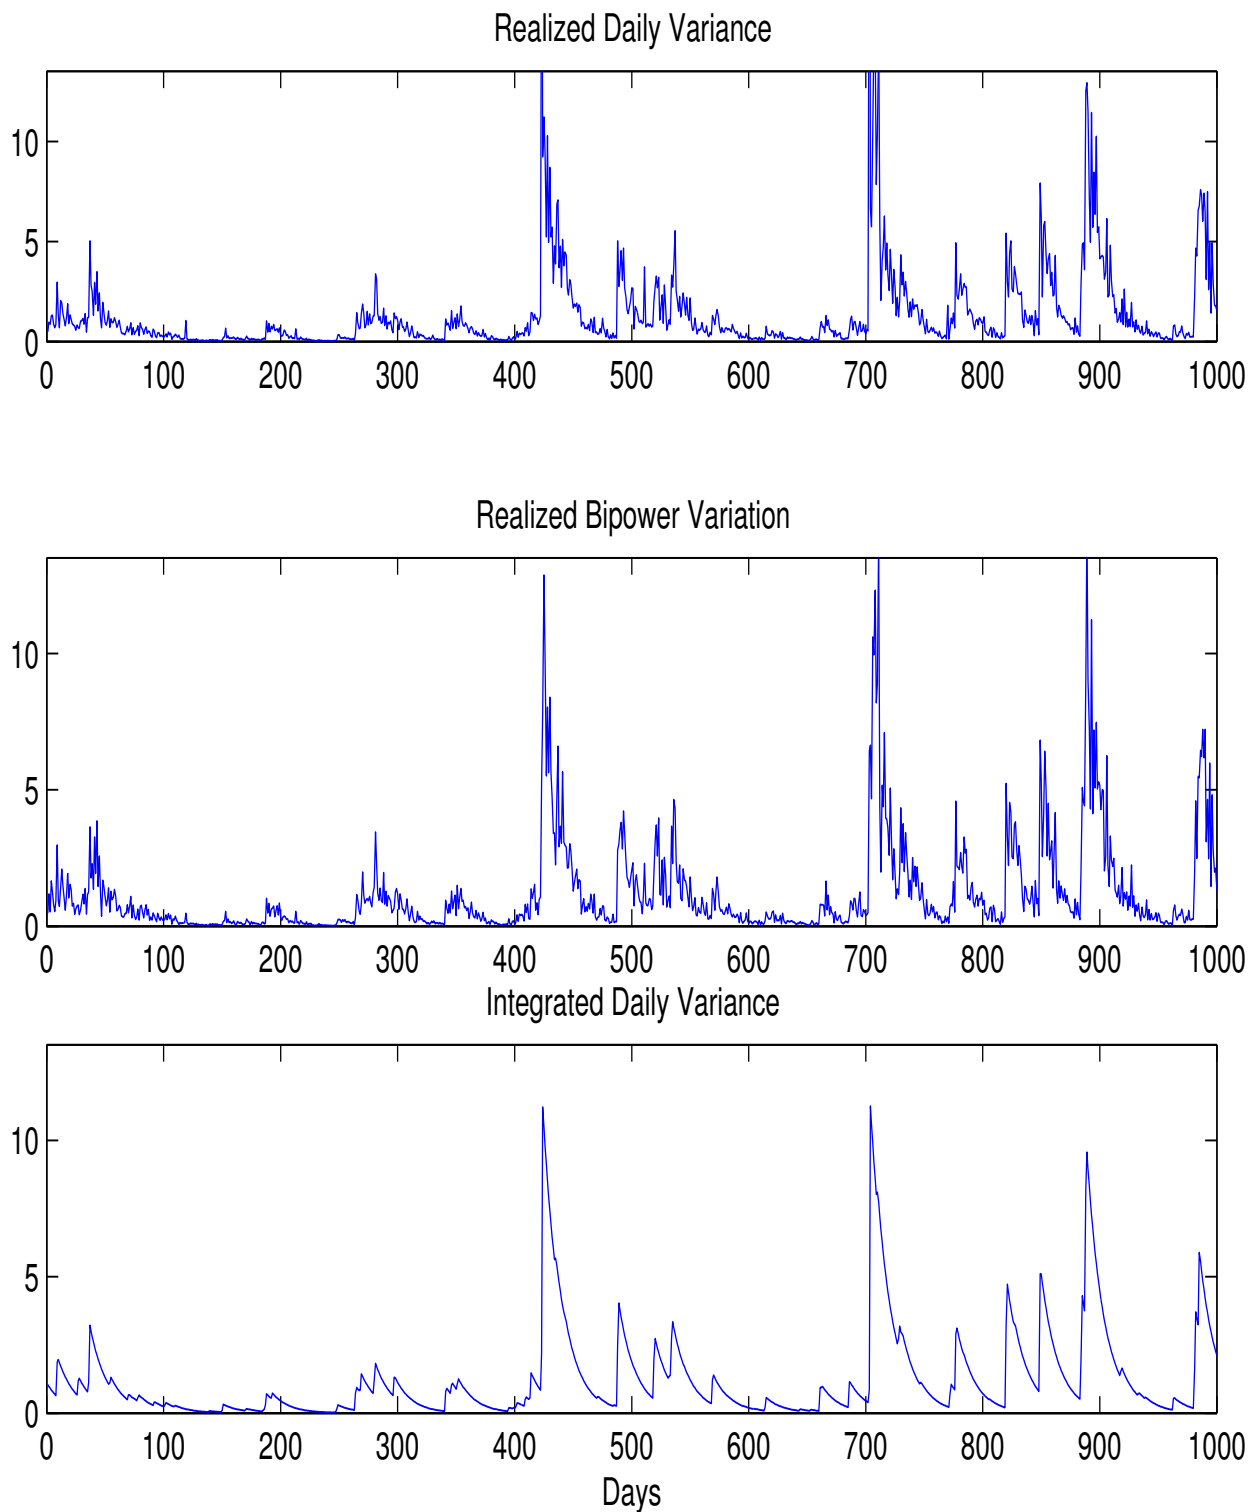


Figure 7: Summary statistics for the simulated MG-CARMA(2,1) stochastic volatility model, corresponding to Figure 6 with parameter settings specified in Table 2. The top panel shows the realized daily variance, the middle one shows the bipower variation and the third panel shows the integrated daily variance.

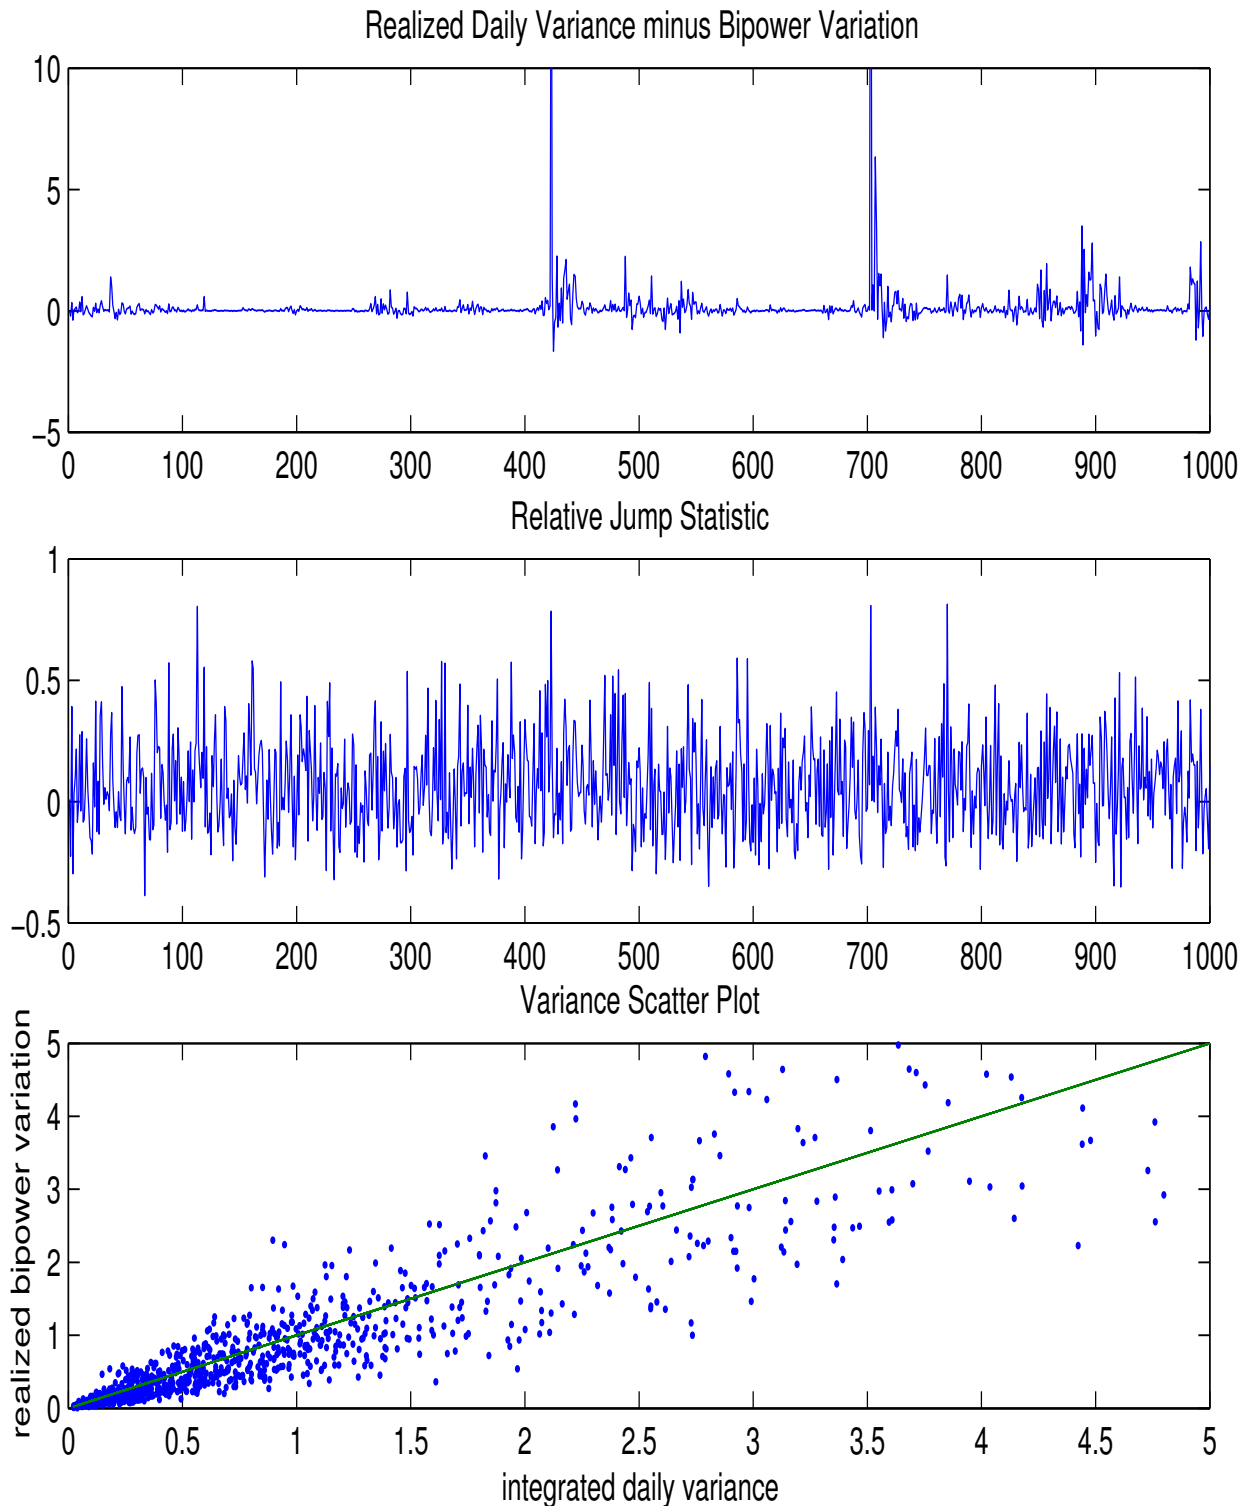


Figure 8: Summary statistics for the simulated MG-CARMA(2,1) stochastic volatility model, corresponding to Figure 6 with parameter settings specified in Table 2. The top panel shows the difference between the realized variance and the bipower variation; the middle one is a plot of the relative jump statistic; the bottom panel plots the bipower variation against the integrated daily variance.




Figure 9: Summary statistics for the simulated MG-CARMA(2,1) stochastic volatility model, corresponding to Figure 6 with parameter settings specified in Table 2. The top panel shows the autocorrelation in the absolute half hour returns; the second panel shows the autocorrelation in the squared half hour returns; the third shows the serial cross correlation between the increment in the price jump process and lags and leads of the integrated variance. In all three panels, the maximum number of lags corresponds to two weeks of trading days.

References

- Alizadeh, S., M. W. Brandt, and F. Diebold (2002). Range-Based Estimation of Stochastic Volatility Models. *Journal of Finance* 57, 1047–1091.
- Andersen, T., T. Bollerslev, and F. Diebold (2005). Some Like it Smooth, and Some Like it Rough: Disentangling Continuous and Jump Components in Measuring, Modelling and Forecasting Asset Return Volatility. Working paper, Duke University.
- Asmussen, S. and J. Rosiński (2001). Approximation of Small Jumps of Lévy Processes with a View Toward Simulation. *Journal of Applied Probability* 38, 482–493.
- Barndorff-Nielsen, O. E. and A. Lindner (2004). Some Aspects of Lévy Copulas. Research Report (Submitted.).
- Barndorff-Nielsen, O. E. and N. Shephard (2001a). Modelling by Lévy Processes for Financial Econometrics. In O. E. Barndorff-Nielsen, T. Mikosch, and S. Resnick (Eds.), *Lévy Processes-Theory and Applications*. Boston: Birkhäuser.
- Barndorff-Nielsen, O. E. and N. Shephard (2001b). Non- Gaussian Ornstein-Uhlenbeck-based Models and Some of Their Applications in Financial Economics. *Journal of the Royal Statistical Society: Series B* 63, 167–241.
- Barndorff-Nielsen, O. E. and N. Shephard (2001c). Normal Modified Stable Processes. *Theory of Probability and Mathematical Statistics* 65, 1–19.
- Barndorff-Nielsen, O. E. and N. Shephard (2004). Power and Bipower Variation with Stochastic Volatility and Jumps. *Journal of Financial Econometrics* 2, 1–37.
- Barndorff-Nielsen, O. E. and N. Shephard (2005a). Impact of Jumps on Returns and Realised Variances: Econometric Analysis of Time-Deformed Lévy Processes. *Journal of Econometrics*, forthcoming.
- Barndorff-Nielsen, O. E. and N. Shephard (2005b). Variation, Jumps, Market Frictions and High Frequency Data in Financial Econometrics. manuscript, Nuffield College, Oxford.
- Barndorff-Nielsen, O. E. and N. Shephard (2006). Econometrics of Testing for Jumps in Financial Economics using Bipower Variation. *Journal of Financial Econometrics*, forthcoming.

- Barndorff-Nielsen, O. E., N. Shephard, and M. Winkel (2005). Limit Theorems for Multipower Variation in the Presence of Jumps in Financial Econometrics. manuscript, Nuffield College, Oxford.
- Bertoin, J. (1996). *Lévy Processes*. Cambridge, UK: Cambridge University Press.
- Bollerslev, T., U. Kretschmer, C. Pigorsch, and G. Tauchen (2005). A Discrete-Time Model for Daily S&P500 Returns and Realized Variations: Jumps and Leverage Effects. Working paper, Duke University.
- Bollerslev, T., J. Litvinova, and G. Tauchen (2005). Volatility Asymmetry in High Frequency Data. Working paper (preliminary), Duke University.
- Brockwell, P. (2001a). Lévy -Driven CARMA Processes. *Ann.Inst.Statist.Math* 53, 113–124.
- Brockwell, P. (2001b). Continuous-Time ARMA Processes. In D. Shanbhag and C. Rao (Eds.), *Handbook of Statistics*, Volume 19. North-Holland.
- Brockwell, P. and T. Marquardt (2005). Lévy-Driven and Fractionally Integrated ARMA Processes with Continuous Time Parameter. *Statistica Sinica* 15, 477–494.
- Carr, P., H. Geman, D. Madan, and M. Yor (2002). The Fine Structure of Asset Returns: An Empirical Investigation. *Journal of Business* 75, 305–332.
- Carr, P., H. Geman, D. Madan, and M. Yor (2003). Stochastic Volatility for Lévy Processes. *Mathematical Finance* 13, 345–382.
- Carr, P. and L. Wu (2004). Time-changed Lévy Processes and Option Pricing. *Journal of Financial Economics* 71, 113–141.
- Chernov, M., R. Gallant, E. Ghysels, and G. Tauchen (2003). Alternative Models for Stock Price Dynamics. *Journal of Econometrics* 116, 225–257.
- Cont, R. and P. Tankov (2004). *Financial Modelling With Jump Processes*. London: Chapman & Hall.
- Duffie, D., J. Pan, and K. Singleton (2000). Transform Analysis and Asset Pricing for Affine Jump-Diffusions. *Econometrica* 68, 1343–1376.

- Durham, G. and A. R. Gallant (2002). Numerical Techniques for Maximum Likelihood Estimation of Continuous-Time Diffusion Processes. *Journal of Business and Economic Statistics* 20, 297–316.
- Elerian, O., S. Chib, and N. Shephard (2001). Likelihood Inference for Discretely Observed Non-Linear Diffusions. *Econometrica* 69, 959–993.
- Eraker, B. (2001). MCMC Analysis of Diffusion Models with Application to Finance. *Journal of Business and Economic Statistics* 19, 177–191.
- Eraker, B., M. Johannes, and N. Polson (2003). The Impact of Jumps in Volatility and Returns. *Journal of Finance* 58, 1269–1300.
- Ferguson, T. and M. Klass (1972). A representation of independent increment processes without gaussian components. *Annals of Mathematical Statistics* 43, 1634–1643.
- Huang, X. and G. Tauchen (2006). The Relative Contributions of Jumps to Total Variance. *Journal of Financial Econometrics*, forthcoming.
- Li, H., M. Wells, and L. Yu (2004). A MCMC Analysis of Time-Changed Lévy Processes of Stock Returns Dynamcis. Working paper, Cornell University.
- Litvinova, J. (2004). Volatility Asymmetry in High Frequency Data. PhD Thesis, Duke University.
- Nicolato, E. and E. Venardos (2003). Option Pricing in Stochastic Volatility Models of the Ornstein-Uhlenbeck Type. *Mathematical Finance* 13, 445–466.
- Roberts, G., O. Papaspiliopoulos, and P. Dellaportas (2004). Bayesian Inference for Non-Gaussian Ornstein Uhlenbeck Stochastic Volatility Processes. *Journal of the Royal Statistical Society: Series B* 66, 369–393.
- Rosiński, J. (1991). On a Class of Infinitely Divisible Processes Represented as Mixtures of Gaussian Processes. In S. Cambanis, G. Samorodnitsky, and T. S. Taqqu (Eds.), *Stable Processes and Related Topics*. Boston: Birkhäuser.
- Rosiński, J. (2001). 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.

- Rosiński, J. (2002). Tempered Stable Processes. In O. E. Barndorff-Nielsen (Ed.), *Second MaPhySto Conference on Lévy Processes: Theory and Application*.
- Rosiński, J. (2004). Tempering Stable Processes. Working paper, University of Tennessee.
- Sato, K. (1999). *Lévy Processes and Infinitely Divisible Distributions*. Cambridge, UK: Cambridge University Press.
- Tankov, P. (2003). Dependence Structure of Lévy Processes with Applications in Risk Management. Rapport Interne 502, CMAP, Ecole Polytechnique.
- Tauchen, G. (1997). New Minimum Chi-square Methods in Empirical Finance. In D. Kreps and K. Wallis (Eds.), *Advances in Econometrics, Seventh World Congress*, pp. 279–317. Cambridge UK: Cambridge University Press.
- Tauchen, G. (2004). Recent Developments in Stochastic Volatility: Statistical Modelling and General Equilibrium Analysis. Working paper, Duke University.
- Tauchen, G. (2005). Stochastic Volatility in General Equilibrium. Working paper, Duke University.
- Wiktorsson, M. (2002). Simulation of Stochastic Integrals with respect to Lévy Processes of type G. *Stochastic Processes and their Applications* 101, 113–125.