

The Normal-Generalised Gamma-Pareto Process: A Novel Pure-Jump Lévy Process with Flexible Tail and Jump-Activity Properties

Fadhel Ayed^{*}, Juho Lee[†], and François Caron[‡]

Abstract. We propose a novel family of self-decomposable Lévy processes where one can control separately the tail behavior and the jump activity of the process, via two different parameters. Crucially, we show that one can sample exactly increments of this process, at any time scale; this allows the implementation of likelihood-free Markov chain Monte Carlo algorithms for (asymptotically) exact posterior inference. We use this novel process in Lévy-based stochastic volatility models to predict the returns of stock market data, and show that the proposed class of models leads to superior predictive performances compared to classical alternatives.

Keywords: stochastic volatility models, power-law, regular variation, Ornstein-Uhlenbeck, Bayesian inference, pseudo-marginal Markov chain Monte Carlo.

1 Introduction

Pure-jump Lévy processes are a flexible class of stochastic processes that have found a wide range of applications, including scalable Markov chain Monte Carlo (Şimşekli, 2017), tracking (Zhang and Paisley, 2018) or the analysis of phylogenetic traits (Landis, Schraiber and Liang, 2012; Landis and Schraiber, 2017). Finance is probably the main domain of application, as it is widely accepted that asset prices contain jumps, and such models have been used as building blocks of complex dynamic models of asset or option prices (Madan, Carr and Chang, 1998; Barndorff-Nielsen and Shephard, 2001; Carr et al., 2002; Cont and Tankov, 2004; Huang and Wu, 2004; Jing, Kong and Liu, 2012).

Let $(X_t)_{t \geq 0}$ be a real-valued pure-jump Lévy process. The process is said to have *heavy, power-law tails* if, for any $t, \Delta > 0$,

$$\Pr(|X_{t+\Delta} - X_t| > x) \stackrel{x \rightarrow \infty}{\sim} \Delta C_1 x^{-2\tau} \quad (1)$$

for some power-law exponent $\tau > 0$ and some constant $C_1 > 0$. That is, for large x , the survival function of the increments approximately behaves as a power function.

Many financial time series, such as historical asset returns, exhibit heavy-tails. Empirical evidence seems to indicate that the returns have nonetheless finite variance, hence

^{*}Department of Statistics, University of Oxford, fadhel.ayed@gmail.com

[†]AI Graduate School, KAIST, juholee@kaist.ac.kr

[‡]Department of Statistics, University of Oxford, caron@stats.ox.ac.uk

corresponding to a power-law exponent $\tau > 1$ (Cont and Tankov, 2004, Section 7.3). Starting from the early work of Mandelbrot (1963) with the stable distribution, various infinite-divisible distributions, closely related to Lévy processes, have been proposed to capture power-law tails. Examples include the student t (Blattberg and Gonedes, 1974) or Pareto (Champagnat et al., 2013) distributions; other models with (non power-law) semi-heavy tails such as the normal inverse Gaussian (Barndorff-Nielsen, 1997), generalized hyperbolic (Eberlein, Keller and Prause, 1998) and tempered stable distributions (Cont, Potters and Bouchaud, 1997; Carr et al., 2002) have also been proposed; see (Cont and Tankov, 2004, Section 7.3) for a review.

Another quantity of interest of the Lévy process is the Blumenthal-Gettoor (BG) index $\beta \in [0, 2]$, also known as fractional order. It is defined by

$$\beta = \inf \left\{ r > 0 \mid \sum_{i \geq 1} |J_i|^r \mathbb{1}_{\theta_i \leq 1} < \infty \right\} \quad (2)$$

where $(J_i, \theta_i)_{i \geq 1}$ is the set of jump sizes and jump times. The BG index measures the level of activity of the jumps: as the value of β increases, small jumps tend to become more and more frequent. It is also related to the smoothness properties of the time series (Cont and Tankov, 2004, Section 7.3) and therefore provides interpretable information on the process and its properties. A number of papers have proposed and analysed (model-free) estimators of this index (Aït-Sahalia and Jacod, 2009; Belomestny, 2010; Woerner, 2011; Belomestny and Panov, 2013). Some Lévy processes, such as the normal-tempered stable or tempered stable processes, can capture the whole range $[0, 2]$ via a tuning parameter; other processes, such as the variance gamma ($\beta = 0$), normal-inverse Gaussian, student t or generalised hyperbolic ($\beta = 1$), have a fixed BG index.

For a pure-jump Lévy process, a typical way to obtain a given power-law exponent and BG index is to assume the regular variation of the tail intensity of the Lévy measure ν characterising the Lévy process, such that

$$\int_{|s| \geq x} \nu(ds) \stackrel{x \rightarrow \infty}{\sim} x^{-2\tau} C_1 \quad \text{and} \quad \int_{|s| \geq x} \nu(ds) \stackrel{x \rightarrow 0}{\sim} x^{-\beta} \ell(1/x)$$

for some slowly varying function ℓ , that is such that $\lim_{t \rightarrow \infty} \ell(ct)/\ell(t) = 1$ for all $c > 0$. While many Lévy measures have been proposed in the literature, no tractable model is able to capture both the whole range of power-law exponent $\tau > 0$ and BG index $\beta \in [0, 2]$. Normal-tempered stable process for example (which includes as special case the variance-gamma and normal-inverse Gaussian) capture the whole range of the β index but have light tails. The class of generalised hyperbolic processes can capture heavy tails, but has a fixed BG index equal to 1. The normal-stable process can capture both heavy tails and different BG indices, but the same parameter controls both properties, and the process has infinite variance.

In this paper, we introduce a novel four-parameter pure-jump Lévy process, called *normal generalised gamma-Pareto* (NGGP) process, with the following properties.

- The model can capture power-law, heavy tails with a single parameter $\tau > 0$; for $\tau > 1$, the process has finite variance;

- Another parameter $\sigma \in (-\infty, 1)$ controls the BG index and therefore the activity of the jumps, with $\beta = \max(0, 2\sigma)$. The process is finite-activity for $\sigma < 0$; it is infinite-activity for $\sigma \geq 0$; it is of bounded variation if $\sigma < 1/2$ and of unbounded variation if $\sigma \in [1/2, 1)$;
- The other two parameters respectively are inverse scale and time scale parameters. More details on the interpretability of the parameters can be found in Section 2.2;
- One can sample exactly from the distribution of the increments, at any time scale; this enables the use of likelihood-free Markov chain Monte Carlo methods for inference;
- The distribution of the increments of the Lévy process is self-decomposable.

The NGGP process is obtained via Brownian subordination, using the subordinator introduced by Ayed, Lee and Caron (2019) for modeling power-law properties of text data. We derive a number of properties of the NGGP process and use the proposed model to predict the stock prices of some financial assets. We consider two Lévy based stochastic volatility models: an exponential Lévy model, and an Ornstein-Uhlenbeck based model Lévy-driven stochastic volatility model (Barndorff-Nielsen and Shephard, 2001). We show that, compared to other Lévy processes, the proposed model is both able to capture the heavy-tail and small-jump behaviours.

The article is organised as follows. In Section 2, we introduce the generalised gamma-Pareto subordinator, its properties, and the associated subordinated Brownian process. In Section 3 we describe two Lévy process based stochastic volatility models, and describe how to perform asymptotically exact posterior inference under our Lévy process with both models. In Section 4 we present experimental results on the modelling of stock prices and show that our model provides a very good fit to the data and good predictive performances compared to classical alternatives.

Notations We use the notation $a_n \stackrel{n \rightarrow \infty}{\sim} b_n$ for $\lim_{n \rightarrow \infty} a_n/b_n = 1$. For a random variable X , the notation $X \sim F$ indicates that X has distribution F . $\text{Gamma}(a, b)$ denotes the gamma distribution with shape parameter a and inverse scale parameter b . $\text{Poisson}(\lambda)$ denotes the standard Poisson distribution with rate λ .

2 The NGGP process

2.1 Generalised gamma subordinator

A generalised gamma (GG) subordinator (almost surely increasing Lévy process) $(Y_t)_{t \geq 0}$ has Lévy intensity (Hougaard, 1986; Aalen, 1992; Brix, 1999)

$$\rho_{\text{GG}}(w; \eta, \sigma, c) = \frac{\eta}{\Gamma(1 - \sigma)} w^{-1 - \sigma} e^{-cw}, \quad w > 0 \quad (3)$$

where $\eta > 0$ and $\sigma \in (-\infty, 1)$, $c > 0$ or $c = 0, \sigma \in (0, 1)$. The subordinator is finite-activity for $\sigma < 0$ and infinite-activity if $\sigma \in [0, 1)$. It admits as special cases the gamma process ($\sigma = 0$), inverse-Gaussian process ($\sigma = 1/2$) and stable process ($c = 0$). When

$\sigma > 0$, the process belongs to the general family of tempered stable processes introduced by Rosiński (2007) and some authors referred to this process simply as a tempered stable process (Barndorff-Nielsen and Shephard, 2002; Liang and Li, 2015). Ignoring the drift term, Y_t has Laplace transform

$$\mathbb{E}[e^{-\vartheta Y_t}] = \exp(-t\psi_{\text{GG}}(\vartheta; \eta, \sigma, c)) \quad (4)$$

where the Laplace exponent is given by

$$\psi_{\text{GG}}(\vartheta; \eta, \sigma, c) = \begin{cases} \frac{\vartheta}{\sigma} [(\vartheta + c)^\sigma - c^\sigma] & \sigma \neq 0 \\ \eta \log(1 + \vartheta/c) & \sigma = 0. \end{cases} \quad (5)$$

Y_t is said to have the generalised gamma distribution with parameters $(\eta t, \sigma, c)$, and we write $Y_t \sim \text{GG}(\eta t, \sigma, c)$. For $\sigma = 0$, $Y_t \sim \text{Gamma}(\eta t, c)$, while for $\sigma < 0$, Y_t is a compound Poisson-gamma distribution with

$$Y_t \stackrel{d}{=} \sum_{k=1}^{K_t} Y_{t,k}$$

where $K_t \sim \text{Poisson}(\eta t \frac{c^\sigma}{-\sigma})$ and $Y_{t,k} \sim \text{Gamma}(-\sigma, c)$ for $k = 1, \dots, K_t$. For $\sigma > 0$, Y_t is an exponentially tilted stable random variable, for which exact samplers exist (Devroye, 2009; Hofert, 2011).

2.2 Generalised gamma-Pareto subordinator

Definition

Let $Z = (Z_t)_{t \geq 0}$ be a subordinator with no drift and Lévy intensity

$$\rho(w) = \frac{\eta}{c^\tau \Gamma(1 - \sigma)} w^{-1-\tau} [\gamma(\tau - \sigma + 1, cw) + (cw)^{\tau-\sigma} e^{-cw}], \quad w > 0 \quad (6)$$

where $\eta > 0$, $c > 0$, $\sigma \in (-\infty, 1)$, $\tau > 0$ and $\gamma(s, x) = \int_0^x t^{s-1} e^{-t} dt$ is the lower incomplete gamma function. For $\tau > \sigma$, using the identity (43) in Section A of the supplementary material, the Lévy intensity takes the simpler form

$$\rho(w) = \frac{\eta(\tau - \sigma)}{c^\tau \Gamma(1 - \sigma)} w^{-1-\tau} \gamma(\tau - \sigma, cw) \quad (7)$$

which is the form in which Ayed, Lee and Caron (2019) introduced the process, with a slightly different parameterisation. The Lévy intensity admits the following representation as a mixture of generalised gamma process

$$\rho(w) = \int_0^\infty u^{-1} \rho_{\text{GG}}\left(\frac{w}{u}; \frac{\eta(\tau - \sigma)}{\tau c^\sigma}, \sigma, c\right) f_U(u) du \quad (8)$$

where $f_U(u) = \tau u^{-1-\tau} \mathbb{1}_{u \geq 1}$ is the probability density function of a Pareto random variable $\text{Pareto}(\tau, 1)$ with support $[1, \infty)$ and power-law exponent $\tau > 0$, and ρ_{GG} is the

Lévy intensity of a GG subordinator, defined in Equation (3). We will thereafter refer to the subordinator with intensity (6) as a *Generalised Gamma-Pareto* (GGP) process. For $x > 0$, let

$$\bar{\rho}(x) = \int_x^\infty \rho(w)dw \quad (9)$$

be the tail Lévy intensity. For $t > 0$, we denote F_{Z_t} the cumulative distribution function of the random variable Z_t , with Laplace transform $\mathbb{E}[e^{-\vartheta Z_t}] = e^{-t\psi(\vartheta)}$, where ψ is the Laplace exponent which is given by, noting that $\gamma(\tau - \sigma + 1, cw) = w^{\tau - \sigma + 1} \int_0^c u^{\tau - \sigma} e^{-wu} du$,

$$\psi(\vartheta) = \int_0^\infty (1 - e^{-w\vartheta})\rho(w)dw = \frac{\eta}{c^\tau} \left[\frac{c^\tau}{\tau} - \int_0^c (u + \vartheta)^{\sigma-1} u^{\tau-\sigma} du + \frac{c^{\tau-\sigma}}{\sigma} ((\vartheta + c)^\sigma - c^\sigma) \right]. \quad (10)$$

Z_t is said to have GGP($t\eta, \sigma, \tau, c$) distribution.

Properties

We derive here a number of properties of the Lévy process Z and of the GGP distribution.

Positive stable process The positive stable process with Lévy intensity $\frac{\eta}{\Gamma(1-\sigma)} w^{-1-\sigma}$ is obtained as a special case when $c = 1$, $\sigma = \tau \in (0, 1)$.

Scaled GG process Let

$$Y_t = \sum_{i \geq 1} W_i \mathbf{1}_{\theta_i \leq t}$$

where $\{(W_i, \theta_i)\}_{i \geq 1}$ are the jump sizes and times of a GG subordinator. Then, for $\tau > \sigma$, the representation (8) implies that

$$Z_t \stackrel{d}{=} \sum_{i \geq 1} W_i U_i \mathbf{1}_{\theta_i \leq t}$$

where $U_i \sim \text{Pareto}(\tau, 1)$. The jump sizes of the GGP subordinator are obtained by scaling the jumps of a GG subordinator with independent Pareto random variables.

Moments and cumulants We have $\mathbb{E}[Z_t^m] < \infty$ for $m < \tau$ and $\mathbb{E}[Z_t^m] = \infty$ otherwise. For $1 \leq m < \tau$, the m th cumulant is given by

$$\kappa_m(Z_t) = t \int_0^\infty w^m \rho(w)dw = \frac{t\eta(\tau - \sigma)\Gamma(m - \sigma)}{c^m(\tau - m)\Gamma(1 - \sigma)}.$$

In particular,

$$\mathbb{E}[Z_t] = \frac{t\eta(\tau - \sigma)}{c(\tau - 1)}, \quad \text{var}(Z_t) = \frac{t\eta(\tau - \sigma)(1 - \sigma)}{c^2(\tau - 2)}.$$

where the expectation exists for $\tau > 1$ and the variance for $\tau > 2$.

Inverse scale parameter If the intensity ρ is of the form (6) for some parameters (η, σ, τ, c) , then $\rho(w/c)/c$ is also of the form (6) with parameters $(\eta, \sigma, \tau, 1)$. c is therefore an inverse scale parameter, and if $Z_t \sim \text{GGP}(t\eta, \sigma, \tau, c)$ then $cZ_t \sim \text{GGP}(t\eta, \sigma, \tau, 1)$.

Activity of the jumps and BG index The Lévy intensity (6) satisfies $\int_0^\infty \rho(w)dw = \infty$ if $\sigma \geq 0$ and the subordinator is therefore infinite-activity. If $\sigma < 0$, $\int_0^\infty \rho(w)dw < \infty$ and it is finite-activity. More precisely, as noted by Ayed, Lee and Caron (2019), the tail Lévy intensity is regularly varying at 0

$$\bar{\rho}(x) \stackrel{x \rightarrow 0}{\sim} \ell(1/x)x^{-\alpha} \quad (11)$$

where $\alpha = \max(0, \sigma)$ is the BG index, and the slowly varying function ℓ is defined by

$$\ell(t) = \begin{cases} \frac{\eta}{c^\sigma \sigma \Gamma(1-\sigma)} & \sigma > 0 \\ \eta \log(t) & \sigma = 0 \\ \frac{\eta(\tau-\sigma)}{-\sigma\tau} & \sigma < 0. \end{cases} \quad (12)$$

For any $x > 0$ and $t \geq 0$, notice that $\bar{\rho}(x) = \mathbb{E} \left[\sum_{i \geq 1} \mathbf{1}_{J_i \geq x} \mathbf{1}_{\theta_i \in [t, t+1]} \right]$, where $\{(J_i, \theta_i)\}_{i \geq 1}$ are the jump sizes and times of the GGP subordinator. Hence the BG index α controls the number of jumps above a certain threshold $x > 0$ per time unit. It also tunes a number of asymptotic properties of the Laplace exponent of the Lévy measure and of the cumulative distribution function and small time distribution of the increments, as described below.

It follows from the Abelian theorem (Gnedin, Hansen and Pitman, 2007, Proposition 17) that the Laplace exponent satisfies

$$\psi(\vartheta) \stackrel{\vartheta \rightarrow \infty}{\sim} \Gamma(1-\alpha)\vartheta^\alpha \ell(\vartheta). \quad (13)$$

For $\sigma = \alpha \in (0, 1)$, the cumulative distribution function F_{Z_t} satisfies (Bingham, Goldie and Teugels, 1989, Theorem 8.2.2. p. 341)

$$-\log F_{Z_t}(z) \stackrel{z \rightarrow 0}{\sim} (1-\alpha)\alpha^{\alpha/(1-\alpha)} \left(\frac{\eta t}{c^\alpha \alpha} \right)^{1/(1-\alpha)} z^{-\alpha/(1-\alpha)}. \quad (14)$$

Additionally, using (13), for small increments, we have, for all $\vartheta \geq 0$

$$\mathbb{E} \left[e^{-\vartheta \frac{cZ_t}{(\eta t/\alpha)^{1/\alpha}}} \right] \xrightarrow{t \rightarrow 0} e^{-\vartheta^\alpha}$$

hence $\frac{cZ_t}{(\eta t/\alpha)^{1/\alpha}}$ tends in distribution to a positive stable random variable with parameter $\alpha \in (0, 1)$ as $t \rightarrow 0$. For $\sigma < 0$, the cdf has a discontinuity at 0 with

$$F_{Z_t}(0) - F_{Z_t}(0_-) = \Pr(Z_t = 0) = e^{-t \frac{\eta(\tau-\sigma)}{-\sigma\tau}}.$$

Heavy tails and power-law behaviour As noted by Ayed, Lee and Caron (2019), the tail Lévy intensity is regularly varying at infinity, with power-law exponent τ . We have

$$\bar{\rho}(x) \stackrel{x \rightarrow \infty}{\sim} \frac{\eta \Gamma(\tau - \sigma + 1)}{\tau c^\tau \Gamma(1 - \sigma)} x^{-\tau}. \quad (15)$$

It follows from (Bingham, Goldie and Teugels, 1989, Theorem 8.2.1. page 341) that the survival function $1 - F_{Z_t}(z) = \Pr(Z_t > z)$ satisfies

$$1 - F_{Z_t}(z) \stackrel{z \rightarrow \infty}{\sim} \frac{\eta t \Gamma(\tau - \sigma + 1)}{\tau c^\tau \Gamma(1 - \sigma)} z^{-\tau} \quad (16)$$

and the increments have heavy, power-law tails with exponent $\tau > 0$.

Simulation of the increments First note that if $\sigma < 0$, the subordinator is a compound Poisson process with jump rate $\frac{\eta(\tau - \sigma)}{-\sigma\tau}$ and jumps being GBFRY distributed (see Section B in the supplementary material) with parameters $(-\sigma, \tau, c)$. We therefore have

$$Z_t \stackrel{d}{=} \sum_{j=1}^{K_t} G_{t,j} U_{t,j}$$

where $K_t \sim \text{Poisson}(t \frac{\eta(\tau - \sigma)}{-\sigma\tau})$, $G_{t,j} \sim \text{Gamma}(-\sigma, c)$ and $U_{t,j} \sim \text{Pareto}(\tau, 1)$, $j = 1, \dots, K_t$ are independent random variables. Consider now the case $\sigma \geq 0$. The Lévy measure admits the two-components mixture representation

$$\begin{aligned} \rho(w) &= \frac{\eta c^{-\sigma}}{\Gamma(1 - \sigma)} w^{-1 - \sigma} e^{-cw} + \frac{\eta}{c^\tau \Gamma(1 - \sigma)} w^{-1 - \tau} \gamma(\tau + 1 - \sigma, cw) \\ &= \frac{\eta c^{-\sigma}}{\Gamma(1 - \sigma)} w^{-1 - \sigma} e^{-cw} + \frac{\eta(\tau - (\sigma - 1))(1 - \sigma)}{c^\tau \Gamma(1 - (\sigma - 1))(\tau - (\sigma - 1))} w^{-1 - \tau} \gamma(\tau - (\sigma - 1), cw) \end{aligned} \quad (17)$$

The first component of the mixture representation (17) is the Lévy intensity of a GG subordinator with parameters $(\eta_1 = \eta/c^\sigma, \sigma_1 = \sigma, c_1 = c)$. The second component is the intensity of a GGP subordinator with parameters $(\eta_2 = \eta(1 - \sigma)/(\tau + 1 - \sigma), \sigma_2 = \sigma - 1, \tau_2 = \tau, c_2 = c)$; as $\sigma_2 = \sigma - 1 < 0$, this subordinator is a finite-activity compound Poisson process, and one can sample its increments as described above. We can therefore write

$$Z_t \stackrel{d}{=} Z_{t,1} + Z_{t,2} \quad (18)$$

where $Z_{t,1} \sim \text{GG}(\frac{\eta t}{c^\sigma}, \sigma, c)$ is an exponentially tilted stable random variable for which exact samplers exist (Devroye, 2009; Hofert, 2011), and

$$Z_{t,2} \stackrel{d}{=} \sum_{j=1}^{\tilde{K}_t} \tilde{G}_{t,j} \tilde{U}_{t,j}$$

where $\tilde{K}_t \sim \text{Poisson}(\frac{\eta t}{\tau})$, $\tilde{G}_{t,j} \sim \text{Gamma}(1 - \sigma, c)$ and $\tilde{U}_{t,j} \sim \text{Pareto}(\tau, 1)$.

Self-decomposability Self-decomposable distributions, a subclass of infinitely-divisible distributions, are closely related to stationary processes of Ornstein-Uhlenbeck type. Such models, described in Section 3.2, have been extensively used for the modeling of financial times series, see e.g. Barndorff-Nielsen and Shephard (2001).

Proposition 1. *The random variable $Z_t \text{GGP}(t\eta, \sigma, \tau, c)$ is self-decomposable if $\sigma \geq 0$. That is, for any $a \in (0, 1)$, there is $Z_t^{(a)}$ independent of Z_t such that*

$$Z_t \stackrel{d}{=} aZ_t + Z_t^{(a)}.$$

Proof. Let $k(w) = w\rho(w)$. Consider first that $\tau > \sigma \geq 0$. From equation (7), we have $k(w) \propto w^{-\sigma} \int_0^c u^{\tau-\sigma-1} e^{-wu} du$ which is non-increasing. Consider that $0 < \tau \leq \sigma$. From Equation (6), k takes the form $k(w) \propto w^{-\tau} g(w)$ where

$$g(w) = \gamma(\tau - \sigma + 1, cw) + (cw)^{\tau-\sigma} e^{-cw}.$$

As $g'(w) = (\tau - \sigma)c(cw)^{\tau-\sigma-1} e^{-cw} \leq 0$, it follows that k is monotone decreasing. Hence for all $\sigma \geq 0$, k is monotone decreasing; using Proposition 15.3 p.485 in (Cont and Tankov, 2004), we conclude that the process is therefore self-decomposable. \square

The self-decomposable random variable Z_1 admits the representation (Jurek, 2001)

$$Z_1 \stackrel{d}{=} \int_0^\infty e^{-s} d\tilde{Z}_s \quad (19)$$

where $(\tilde{Z}_s)_{s \geq 0}$ is termed the background driving Lévy process corresponding to the self-decomposable random variable Z_1 (Barndorff-Nielsen and Shephard, 2001, Section 2.2). \tilde{Z}_t has Lévy intensity

$$\tilde{\rho}(w) = -\rho(w) - w\rho'(w) = \frac{\eta\sigma}{c^\sigma \Gamma(1-\sigma)} w^{-1-\sigma} e^{-cw} + \frac{\eta\tau}{c^\tau \Gamma(1-\sigma)} w^{-1-\tau} \gamma(\tau - \sigma + 1, cw). \quad (20)$$

Importantly, for $\sigma = 0$, the background Lévy process is a finite-activity GGP process with intensity

$$\tilde{\rho}(w) = \frac{\eta\tau}{c^\tau} w^{-1-\tau} \gamma(\tau + 1, cw). \quad (21)$$

Interpretability of the parameters In summary, each of the four parameters governs a different property of the GGP process.

- $\eta > 0$ is a time-scaling parameter: if $(Z_t)_{t \geq 0}$ is a GGP process with parameters $(1, \sigma, \tau, c)$, then $(Z_{\eta t})_{t \geq 0}$ is a GGP process with parameters (η, σ, τ, c) ;
- $c > 0$ is an inverse-scale parameter: if (Z_t) is a GGP process with parameters $(\eta, \sigma, \tau, 1)$, then (Z_t/c) is a GGP process with parameters (η, σ, τ, c) ;
- $\sigma \in (-\infty, 1)$ tunes the activity of the jumps; the process is finite-activity if $\sigma < 0$ and infinite-activity otherwise, with corresponding BG index $\beta = \max(\sigma, 0)$;
- τ is the power-law exponent, controlling the tails of the distribution, with $\Pr(Z_t > z) \stackrel{z \rightarrow \infty}{\sim} C z^{-\tau}$ for some constant C .

2.3 Normal GGP process

Definition

Let $Z = (Z_t)_{t \geq 0}$ be a GGP subordinator with no drift and Lévy intensity given by (6). Let $B = (B_t)_{t \geq 0}$ be a Brownian motion on \mathbb{R} , independent from Z . The normal generalised gamma Pareto (NGGP) Lévy process, taking values in \mathbb{R} , is defined via Brownian subordination by

$$X_t = B_{Z_t}.$$

For any $t > 0$, Let F_{X_t} denote the cumulative distribution function of the random variable X_t , with characteristic function (Cont and Tankov, 2004, Section 4.2)

$$\mathbb{E}[e^{i\lambda X_t}] = e^{t\Psi(\lambda)}$$

where the characteristic exponent is given by (Cont and Tankov, 2004, Theorem 4.2)

$$\Psi(\vartheta) = -\psi(\vartheta^2/2) = \int_{\mathbb{R}} (e^{i\vartheta x} - 1)\nu(x)dx \quad (22)$$

where ψ is defined in Equation (10) and ν is a Lévy intensity on \mathbb{R} defined by

$$\nu(x) = \int_0^\infty \frac{1}{\sqrt{2\pi w}} e^{-\frac{x^2}{2w}} \rho(w)dw.$$

For any $x > 0$, let

$$\bar{\nu}(x) = \int_{|s|>x} \nu(s)ds$$

denote the expected number of jumps of absolute value larger than x in an unit-length interval. We also write $X_t \sim \text{NGGP}(\eta t, \sigma, \tau, c)$.

Properties

Most of the properties here follow from the properties of the subordinator. By construction, we have, for all t

$$X_t \stackrel{d}{=} \sqrt{Z_t} \epsilon_t \quad (23)$$

where $\epsilon_t \sim \mathcal{N}(0, 1)$. \sqrt{c} is therefore an inverse scale parameter, and if $X_t \sim \text{NGGP}(\eta t, \sigma, \tau, c)$ then $\sqrt{c}X_t \sim \text{NGGP}(\eta t, \sigma, \tau, 1)$.

Moments and cumulants Let $1 \leq m < 2\tau$. For m odd, the m 'th raw moment and cumulant of X_t satisfy

$$\mathbb{E}[X_t^m] = \kappa_m(X_t) = 0. \quad (24)$$

For $2 \leq m < 2\tau$, m even,

$$\kappa_m(X_t) = 2 \int_0^\infty \int_0^\infty \frac{x^m}{\sqrt{2\pi w}} e^{-\frac{x^2}{2w}} \rho(w)dw dx$$

$$\begin{aligned}
&= \frac{2^{m/2}}{\sqrt{\pi}} \kappa_{m/2}(Z_t) \Gamma((m+1)/2) \\
&= \frac{t\eta(\tau-\sigma)2^{m/2}\Gamma((m+1)/2)}{c^{m/2}(\tau-m/2)\sqrt{\pi}}
\end{aligned}$$

It follows, for $\tau > 1$, the Lévy process has finite variance with

$$\text{var}(X_t) = \mathbb{E}[Z_t] = \frac{t\eta(\tau-\sigma)}{c(\tau-1)} < \infty.$$

For $\tau > 2$, the excess kurtosis is finite and given by

$$\text{kurt}(X_t) = \frac{\kappa_4(X_t)}{\kappa_2(X_t)^2} = \frac{3\text{var}(Z_t)}{(\mathbb{E}[Z_t])^2} = \frac{3(\tau-1)^2}{t\eta(\tau-2)(\tau-\sigma)}.$$

Activity of the jumps and BG index The Lévy process X is infinite-activity if $\sigma \geq 0$ and finite-activity otherwise. Using Proposition 2 in the supplementary material C, the regular variation of $\bar{\rho}$ at 0 in Equation (11) implies the regular variation of $\bar{\nu}$ at 0

$$\bar{\nu}(x) \stackrel{x \rightarrow 0}{\sim} \frac{2^{\alpha+1}\Gamma(\alpha+1/2)}{\sqrt{\pi}} \ell(1/x^2) x^{-2\alpha}. \quad (25)$$

The BG index of X is therefore equal to $2\alpha = 2\max(0, \sigma) \in [0, 2)$. When $\sigma = \alpha > 0$, combining (22) with (13), we obtain the small time limit

$$\mathbb{E} \left[\exp \left(i\vartheta \frac{X_t \sqrt{2c}}{(\eta t)^{1/(2\alpha)}} \right) \right] \stackrel{t \rightarrow 0}{\rightarrow} e^{-|\vartheta|^{2\alpha}}$$

hence $\frac{X_t \sqrt{2c}}{(\eta t)^{1/(2\alpha)}}$ tends in distribution to a symmetric stable distribution with parameter $2\alpha \in (0, 2)$ when t tends to 0.

Heavy tails and power-law behaviour Using Proposition 2 in the Supplementary material (Ayed, Lee and Caron (2022)), the regular variation of $\bar{\rho}$ at infinity in Equation (15) implies the regular variation of $\bar{\nu}$ at infinity

$$\bar{\nu}(x) \stackrel{x \rightarrow \infty}{\sim} C_1 x^{-2\tau} \quad (26)$$

where

$$C_1 = \frac{2^{\tau+1}\Gamma(\tau+1/2)}{\sqrt{\pi}} \frac{\eta\Gamma(\tau-\sigma+1)}{\tau c^\tau \Gamma(1-\sigma)}. \quad (27)$$

Additionally, we have

$$\Pr(|X_t| > x) \stackrel{x \rightarrow \infty}{\sim} C_1 t x^{-2\tau}. \quad (28)$$

The increments have therefore heavy tails with power-law exponent 2τ .

Simulation of the increments As $X_t \stackrel{d}{=} \sqrt{Z_t} \epsilon_t$ one can simulate increments exactly by sampling Z_t from (18) and $\epsilon_t \sim \mathcal{N}(0, 1)$.

Self-decomposability The self-decomposability of X follows from the self-decomposability of the subordinator Z (Sato, 2001, Theorem 1).

Interpretability of the parameters In summary, each of the four parameters governs a different property of the NGGP process.

- $\eta > 0$ is a time-scaling parameter: if $(X_t)_{t \geq 0}$ is a NGGP process with parameters $(1, \sigma, \tau, c)$, then $(X_{\eta t})_{t \geq 0}$ is a NGGP process with parameters (η, σ, τ, c) ;
- $\sqrt{c} > 0$ is an inverse-scale parameter: if (X_t) is a NGGP process with parameters $(\eta, \sigma, \tau, 1)$, then (Z_t/\sqrt{c}) is a NGGP process with parameters (η, σ, τ, c) ;
- $\sigma \in (-\infty, 1)$ tunes the activity of the jumps; the process is finite-activity if $\sigma < 0$ and infinite-activity otherwise, with corresponding BG index $\beta = \max(2\sigma, 0)$;
- 2τ is the power-law exponent, controlling the tails of the distribution, with $\Pr(X_t > z) \stackrel{z \rightarrow \infty}{\sim} tC_1 z^{-2\tau}$ for some constant C_1 .

2.4 Generalisations

More general subordinators

One could consider more generally a Lévy intensity of the form

$$\rho(w) = \frac{\eta}{c^\sigma \Gamma(1-\sigma)} w^{-1-\sigma} \left(wc \int_0^1 u^{1-\sigma} h(u) e^{-ucw} du + h(1) e^{-cw} \right) \quad (29)$$

where $h : (0, 1] \rightarrow (0, \infty)$ is a differentiable function which satisfies

$$\int_0^1 h(u) du < \infty \quad \text{and} \quad h(u) \stackrel{u \rightarrow 0}{\sim} u^{\tau-1}. \quad (30)$$

For $\tau > \sigma$, we have $u^{1-\sigma} h(u) \rightarrow 0$ and (29) takes the alternative form

$$\rho(w) = \frac{\eta}{c^\sigma \Gamma(1-\sigma)} w^{-1-\sigma} \int_0^1 (u^{1-\sigma} h(u))' e^{-uwc} du. \quad (31)$$

The proposed subordinator (6) is obtained as a special case when $h(u) = u^{\tau-1}$.

The Lévy process is finite activity for $\sigma < 0$ and infinite-activity for $\sigma \geq 0$. Using Karamata's theorem for regularly varying functions, the tail Lévy intensity $\bar{\rho}$ of the Lévy intensity is regularly varying at 0 with BG index $\alpha = \max(0, \sigma)$ and at infinity with tail index τ . The Lévy intensity (29) takes the form of a sum of a compound Poisson intensity and a generalised gamma intensity. The compound Poisson intensity can be written as

$$\rho_1(w) = \frac{\eta}{c^{\sigma-1} \Gamma(1-\sigma)} \int_0^1 h(u) u(uw)^{-\sigma} e^{-uw} du$$

$$= \frac{\eta(1-\sigma) \int_0^1 h(u) du}{c^{\sigma-1}} \int_0^1 \frac{h(u)}{\int_0^1 h(v) dv} u \rho_{\text{GG}}(uw; \sigma-1, c) du$$

which is a mixture of (finite-activity) generalised gamma processes. It follows that if Z is a subordinator with Lévy intensity (29), we have $Z_t \stackrel{d}{=} Z_{t,1} + Z_{t,2}$ where $Z_{t,1} \sim \text{GG}(\frac{\eta h(1)}{c^\sigma}, \sigma, c)$ is an exponentially tilted stable random variable and

$$Z_{t,2} \stackrel{d}{=} \sum_{j=1}^{\tilde{K}_t} \tilde{G}_{t,j} \tilde{U}_{t,j}$$

where $\tilde{K}_t \sim \text{Poisson}(\eta t \int_0^1 h(u) du)$, $\tilde{G}_{t,j} \sim \text{Gamma}(1-\sigma, c)$ and $1/\tilde{U}_{t,j}$ have probability density function $\frac{h(u)}{\int_0^1 h(v) dv}$. Finally, if $u^{1-\sigma} h(u)$ is monotone increasing, with $\tau > \sigma$, then Equation (31) implies that the Lévy process is a tempered stable process if $\sigma > 0$, and a generalised gamma convolution if $\sigma = 0$; it is therefore self-decomposable.

Subordinated fractional Brownian motion

Many of the properties of the GGP process extend to the NGGP due to the self-similarity properties of Brownian motion. A stochastic process $(X_t)_{t \geq 0}$ with $X_0 = 0$ almost surely is self-similar if there exists an index $H > 0$ such that for all $a > 0$,

$$(X_{at})_{t \geq 0} \stackrel{d}{=} (a^H X_t)_{t \geq 0}.$$

The Brownian motion is self-similar with index $H = 1/2$. Another popular class of self-similar processes are fractional Brownian motions. A fractional Brownian motion $(F_t)_{t \geq 0}$ is a zero-mean Gaussian process with covariance

$$R(t, s) := \mathbb{E}[F_t F_s] = \frac{1}{2}(t^{2H} + s^{2H} - |t - s|^{2H})$$

Brownian motion is obtained as a special case for $H = 1/2$. The process is self-similar with index H and has stationary Gaussian increments. The increments are negatively correlated if $H < 1/2$, are independent if $H = 1/2$ and positively correlated if $H > 1/2$.

Most of the properties described in Section 2.3 can be similarly derived in the more general case where the Brownian motion is replaced by a fractional Brownian motion. Let (Z_t) be a GGP process with parameters (η, σ, τ, c) and (F_t) a fractional Brownian motion with index $H > \max(0, \sigma)/2$. The subordinated fractional Brownian process

$$X_t = F_{Z_t}$$

satisfies the self-similar property

$$X_t \stackrel{d}{=} Z_t^H \epsilon_t$$

where $\epsilon_t \sim \mathcal{N}(0, 1)$. It follows that

$$\Pr(|X_t| > x) \stackrel{x \rightarrow \infty}{\sim} C_1 t x^{-\tau/H} \quad (32)$$

where the constant C_1 is defined in Equation (27). Increments of the process at times t_1, \dots, t_n can be simulated exactly by first simulating $Z_{t_k} - Z_{t_{k-1}} \stackrel{d}{=} Z_{t_k - t_{k-1}}$ using (18), for $k = 1, \dots, n$ then, conditional on $(Z_{t_1}, \dots, Z_{t_n})$ simulate

$$X_1, \dots, X_n | Z_{t_1}, \dots, Z_{t_n} \sim \mathcal{N}(0, (R(Z_{t_i}, Z_{t_j}))_{1 \leq i, j \leq n}).$$

2.5 Comparison to other models and discussion

Comparison to Ayed et al. The (normalised) GGP process was introduced by Ayed, Lee and Caron (2019) as a prior for random probability measures with power-law properties, and applied to the modeling of word frequencies. Ayed, Lee and Caron (2019) introduced the form (7) which is only valid for $\tau > \sigma$. The alternative form (6) we introduce here allows to deal with the case $0 < \tau \leq \sigma$ as well; in particular, one obtains the stable process as a particular case. Ayed, Lee and Caron (2019) showed that the tail Lévy intensity of the GGP is regularly varying at 0 and infinity and deduced the asymptotic behaviour of large and small jumps. Here we derive a number of additional important properties of the process and of the distribution of the increments. We show that it is decomposable, and crucially, that one can sample exactly the increments at any time scale. Ayed, Lee and Caron (2019) used the name GBFRY process for the process, due to its form similar to the form of the GBFRY distribution (see Section B in the Supplementary material Ayed, Lee and Caron (2022)); however, as it is customary to give the same name to the process and to the distribution of the increments, which are not GBFRY distributed, we prefer here to use the name generalised gamma-Pareto.

Tempered stable process and generalised gamma convolutions If $\tau > \sigma > 0$, the subordinator falls in the general class of tempered stable processes, introduced by Rosiński (2007). Noting that $\gamma(\tau - \sigma, cw) = w^{\tau - \sigma + 1} \int_0^c u^{\tau - \sigma} e^{-wu} du$, the model (6) is indeed of the form $w^{-1 - \sigma} q(w)$ where the so-called tempering function q is given by

$$q(w) = \frac{\eta(\tau - \sigma)}{c^\tau \Gamma(1 - \sigma)} \int_0^c u^{\tau - \sigma - 1} e^{-wu} du. \quad (33)$$

By Bernstein's theorem, the function q is completely monotone.

For $\sigma = 0$, the subordinator belongs to the class of generalised gamma convolutions (Thorin, 1977; Bondesson, 1992; James, Roynette and Yor, 2008), of the form $w^{-1} \int_0^\infty e^{wu} U(du)$ with Thorin measure $U(du) = \frac{\eta\tau}{c^\tau} u^{\tau-1} 1_{u \in (0, c)} du$.

The subordinator, for any σ , also falls into the extended Thorin class described by Grigelionis (2007), see also the discussion in Section 1.8 in (James, Roynette and Yor, 2008).

Comparison to other models As mentioned in the introduction, a number of different Lévy processes have been proposed in the literature. While each process can capture some range of the different tail and jump behaviour, none of them is flexible enough to capture the whole range of tail and jump-activity indices. Variance gamma, normal

Model	Heavy tails	Finite 2nd moment	BG index	Tractable f_t for any t	Exact simulation from f_t for any t
VG	No	Yes	$\beta = 0$	Yes	Yes
NIG	No	Yes	$\beta = 1$	Yes	Yes
NGG	No	Yes	$\beta \in (0, 2)$	No	Yes
NS	Yes, $2\tau \in (0, 2)$	No	$\beta = 2\tau \in (0, 2)$	No	Yes
TS	No	Yes	$\beta \in [0, 2)$	No	Yes
St	Yes, $2\tau \in (0, \infty)$	Yes if $2\tau > 2$	$\beta = 1$	No	No
GH	Depends	Depends	$\beta = 1$	No	No
NGGP	Yes, $2\tau \in (0, \infty)$	Yes if $\tau > 1$	$\beta = 2 \max(0, \sigma) \in [0, 2)$	No	Yes

Table 1: Comparison between different Lévy processes. VG: Variance Gamma; NIG: Normal inverse Gaussian; NGG: Normal generalised gamma; NS: Normal Stable; TS: Tempered Stable; St: Student t; GH: Generalised hyperbolic.

inverse Gaussian, exponentially tilted stable and tempered stable process do not capture heavy tails; the normal stable process has infinite variance, and the same parameter tunes the activity of the jumps and the BG index; for generalised hyperbolic process, the BG index is fixed to 1.

A drawback of the proposed model is that, contrary to popular models such as the variance gamma or normal inverse Gaussian processes, the increments X_t do not have an analytical probability density function f_t . This is balanced however by the fact that one can sample exactly from the distribution of the increments, and one can therefore resort to likelihood-free methods for posterior inference, as described in the next section. Table 1 summarises the properties of the different models. Note that, as mentioned in (Cont and Tankov, 2004, Section 4.6), the generalised hyperbolic and student t are not closed under convolution, and so there is no analytic expression for f_t at any given time $t > 0$, which may be an issue if data are sampled irregularly.

Some interesting connections can be drawn with other classes of stochastic processes. If $\tau > \sigma$, due to the mixture form (8), the GGP distribution arises as the marginal distribution of a quantile clock process (see Theorem 3.1 by James and Zhang (2011)) with parameters (R, L) where R is a Pareto random variable and L a GG subordinator,

3 Lévy-driven stochastic volatility models

Let S_t denote the price of a financial asset, e.g. a market or a stock index, at time t . Denote $X_t = \log\left(\frac{S_t}{S_0}\right)$. Observations are obtained at fixed discrete times $t_0 = 0 < t_1 < t_2 < \dots < t_n$, and we write, for $k = 1, 2, \dots$

$$Y_k := \log \frac{S_{t_k}}{S_{t_{k-1}}} = X_{t_k} - X_{t_{k-1}} \quad (34)$$

the log-returns (or more shortly, called returns). Let $\Delta_k = t_k - t_{k-1}$ be the inter-arrival times between observations. We assume that

$$X_t = \mu_0 t + \mu_1 V_t^* + B_{V_t^*} \quad (35)$$

where μ_0 is the drift parameter, μ_1 is the risk premium, B_t is a Brownian motion, independent of the stochastic process V_t^* , which can be interpreted as the integrated stochastic volatility. For $k = 1, \dots, n$, let

$$\bar{V}_k = V_{t_k}^* - V_{t_{k-1}}^*$$

be the integrated stochastic volatility over the interval (t_{k-1}, t_k) . The observations (Y_1, \dots, Y_n) are conditionally independent given $(\bar{V}_1, \dots, \bar{V}_n)$, with

$$Y_k \mid \bar{V}_k \sim \mathcal{N}(\mu_0 \Delta_k + \mu_1 \bar{V}_k, \bar{V}_k).$$

We consider two different stochastic processes for the integrated volatility process (V_t^*) : a Lévy process and a Ornstein-Uhlenbeck based model.

3.1 Exponentiated Lévy process

Assume that $(V_t^*)_{t \geq 0}$ is a subordinator with no drift with Lévy intensity ρ parameterised by a vector ϕ . The integrated volatilities $(\bar{V}_1, \dots, \bar{V}_n)$ are therefore conditionally independent, with

$$\bar{V}_k \mid \phi \sim F_{V_{\Delta_k}^*} \quad (36)$$

where $F_{V_t^*}$ denotes the distribution of V_t^* , with Laplace transform

$$\int_0^\infty e^{-\lambda x} dF_{V_t^*}(x) = e^{-t \int_0^\infty (1 - e^{-\lambda w}) \rho(w) dw}.$$

If V_t^* is taken to be the GGP model with intensity (6), then $B_{V_t^*}$ is a NGGP Lévy process.

3.2 Ornstein-Uhlenbeck based stochastic volatility model

We also consider a non-Gaussian Ornstein-Uhlenbeck based model (Barndorff-Nielsen and Shephard, 2001) with

$$V_t^* = \int_0^t V_s ds$$

where the instantaneous stochastic volatility process $(V_t)_{t \geq 0}$ is stationary and satisfies

$$V_t = V_0 e^{-\lambda t} + \int_0^t e^{\lambda(s-t)} dZ_s$$

for some $\lambda > 0$ and some background driving Lévy process Z_t with Lévy measure $\tilde{\rho}$. Additionally, for any $t > 0$ the random variable $V_t \sim F$ is infinite-divisible and self-decomposable with Laplace transform

$$\mathbb{E}[e^{-\vartheta V_t}] = e^{-\int_0^\infty (1 - e^{-\vartheta w}) \rho(w) dw}$$

where $\tilde{\rho}$ and ρ are related by the expression

$$\tilde{\rho}(w) = -\rho(w) - w\rho'(w).$$

To define the model, one can either define the mean measure $\tilde{\rho}$ of the subordinator Z_t , or choose the stationary (self-decomposable) distribution F of V_t , hence ρ . In practice, the second approach is often chosen; examples include the gamma (Roberts, Papaspiliopoulos and Dellaportas, 2004; Griffin and Steel, 2006; Fröhlich-Schnatter and Sögner, 2009), generalized inverse Gaussian (Gander and Stephens, 2007), and exponentially tilted stable distributions (Gander and Stephens, 2007; Andrieu, Doucet and Holenstein, 2010) as marginals. The integrated stochastic volatilities over the interval (t_{k-1}, t_k) are obtained, for $k = 1, \dots, n$, by

$$\bar{V}_k = \int_{t_{k-1}}^{t_k} V_t dt = \lambda^{-1} (Z_{\lambda t_k} - V_{t_k} - (Z_{\lambda t_{k-1}} - V_{t_{k-1}})) \quad (37)$$

where $\begin{pmatrix} V_{t_k} \\ Z_{\lambda t_k} \end{pmatrix}$ follows a linear dynamic model with $Z_0 = 0$, $V_0 \sim F$, and for $k = 1, \dots, n$,

$$\begin{pmatrix} V_{t_k} \\ Z_{\lambda t_k} \end{pmatrix} = \begin{pmatrix} e^{-\lambda \Delta_k} V_{t_{k-1}} \\ Z_{\lambda t_{k-1}} \end{pmatrix} + \varepsilon_k, \quad \text{with } \varepsilon_k = \begin{pmatrix} e^{-\lambda \Delta_k} \int_0^{\Delta_k} e^{\lambda t} dZ_{\lambda t} \\ \int_0^{\Delta_k} dZ_{\lambda t} \end{pmatrix}. \quad (38)$$

Exact simulation of $(\bar{V}_1, \dots, \bar{V}_n)$ from the model defined by Equations (37-38) requires to be able to simulate from F and simulate the independent random variables $(\varepsilon_1, \dots, \varepsilon_n)$. We describe two models where exact simulation is possible.

Model with gamma marginal distribution

A classical choice (Barndorff-Nielsen and Shephard, 2001; Roberts, Papaspiliopoulos and Dellaportas, 2004; Griffin and Steel, 2006; Fröhlich-Schnatter and Sögner, 2009) is to take $F = \text{Gamma}(\eta, c)$ as marginal distribution for V_t . This corresponds to

$$\rho(w) = \eta w^{-1} e^{-cw}, \quad \tilde{\rho}(w) = \eta c e^{-cw}. \quad (39)$$

The background driving Lévy measure (Z_t) is therefore finite-activity, and one can sample exactly the state noise ε_k as follows.

1. Simulate $N_k \sim \text{Poisson}(\eta \lambda \Delta_k)$.
2. For $j = 1, \dots, N$, simulate $E_{kj} \sim \text{Exp}(c)$, $\theta_{kj} \sim U(0, \Delta_k)$.
3. Set $\varepsilon_k = \begin{pmatrix} e^{-\lambda \Delta_k} \sum_{j=1}^{N_k} e^{\lambda \theta_{kj}} E_{kj} \\ \sum_{j=1}^{N_k} E_{kj} \end{pmatrix}$.

Model with GGP marginal distribution

Let $\tau = \sigma > 0$ or $\tau > \sigma \geq 0$. As shown in Section 2.2, $\text{GGP}(\eta, \sigma, \tau, c)$ is self-decomposable. If V_t has marginal $F = \text{GGP}(\eta, \sigma, \tau, c)$ distribution, this corresponds

to ρ be defined by Equation (6), and the Lévy intensity $\tilde{\rho}$ of the background driving Lévy intensity (Z_t) is given by Equation (20). For $\sigma > 0$, (Z_t) is infinite-activity, and one needs to resort to numerical methods to approximately sample (ε_k) . This could be done by using the representation of the process as a sum of GG process and a finite activity process as described in Section 2.2, and using a truncated series representation for simulating the GG process.

We focus here on the case $\sigma = 0$, where $\tilde{\rho}$ simplifies to

$$\tilde{\rho}(w) = \frac{\eta\tau}{c^\tau} w^{-1-\tau} \gamma(\tau + 1, cw) \quad (40)$$

with $\int_0^\infty \tilde{\rho}(w)dw = \eta$, and the background Lévy process (Z_t) is therefore finite-activity. We can therefore simulate ε_k exactly as follows.

1. Simulate $N_k \sim \text{Poisson}(\eta\lambda\Delta_k)$
2. For $j = 1, \dots, N$, simulate $E_{kj} \sim \text{Exp}(c)$, $U_{kj} \sim \text{Pareto}(\tau, 1)$, $\theta_{kj} \sim U(0, \Delta_k)$
3. Set $\varepsilon_k = \left(\frac{e^{-\lambda\Delta_k} \sum_{j=1}^{N_k} e^{\lambda\theta_{kj}} E_{kj} U_{kj}}{\sum_{j=1}^{N_k} E_{kj} U_{kj}} \right)$
which is similar to the model with gamma marginals, with $E_{kj}U_{kj}$ in place of E_{kj} .

3.3 Posterior inference

Let ϕ denote the set of unknown parameters of both models. That is, ϕ includes the drift and risk premium parameters μ_0 and μ_1 , the parameters of the Lévy intensity and, for the Ornstein-Uhlenbeck based model, the discounting factor $\lambda > 0$. Let $\pi(\phi)$ be some prior density. We aim at approximating the posterior density $\pi(\phi \mid y_1, \dots, y_n)$. The marginal likelihood takes the form

$$p(y_1, \dots, y_n \mid \phi) = \int_{\mathbb{R}_+^n} \left[\prod_{k=1}^n p(y_k \mid \bar{v}_k, \mu_0, \mu_1) \right] d\mathbf{F}_n(\bar{v}_1, \dots, \bar{v}_n) \quad (41)$$

where

$$\mathbf{F}_n(\bar{v}_1, \dots, \bar{v}_n) = \Pr(\bar{V}_1 \leq \bar{v}_1, \dots, \bar{V}_n \leq \bar{v}_n \mid \phi)$$

denotes the joint cumulative distribution function of the integrated variances. In the exponentiated Lévy process, we have

$$\mathbf{F}_n(\bar{v}_1, \dots, \bar{v}_n) = \prod_{k=1}^n F_{V_{\Delta_k}}(\bar{v}_k).$$

If F_{V_t} does not admit a tractable probability density function, as it is the case for the proposed GGP model, neither \mathbf{F}_n nor $p(y_1, \dots, y_n \mid \phi)$ are tractable, preventing the implementation of a Metropolis-Hastings Markov chain Monte Carlo algorithm. The same applies for the Ornstein-Uhlenbeck model.

We therefore propose to use a pseudo-marginal Markov chain Monte Carlo (MCMC) algorithm (Beaumont, 2003; Andrieu and Roberts, 2009), which only requires to simulate from \mathbf{F}_n . The pseudo-marginal algorithm replaces the untractable marginal likelihood (41) by an unbiased estimator, yet admitting the posterior distribution of interest

as invariant distribution. Let q denote some proposal distribution for the parameters. At iteration i of the algorithm, we have

1. Sample $\phi^* \mid \phi^{(i-1)} \sim q(\cdot \mid \phi^{(i-1)})$
2. Compute an unbiased estimate $\widehat{p}(y_1, \dots, y_n \mid \phi^*)$
3. With probability

$$\min \left(1, \frac{\widehat{p}(y_1, \dots, y_n \mid \phi^*) \pi(\phi^*) q(\phi^{(i-1)} \mid \phi^*)}{\widehat{p}(y_1, \dots, y_n \mid \phi^{(i-1)}) \pi(\phi^{(i-1)}) q(\phi^* \mid \phi^{(i-1)})} \right)$$

set $\phi^{(i)} = \phi^*$ and $\widehat{p}(y_1, \dots, y_n \mid \phi^{(i)}) = \widehat{p}(y_1, \dots, y_n \mid \phi^*)$.

Otherwise, set $\phi^{(i)} = \phi^{(i-1)}$ and $\widehat{p}(y_1, \dots, y_n \mid \phi^{(i)}) = \widehat{p}(y_1, \dots, y_n \mid \phi^{(i-1)})$.

In the exponential Lévy model, an unbiased estimator can be obtained via Monte Carlo approximation

$$\widehat{p}(y_1, \dots, y_n \mid \phi) = \prod_{k=1}^n \frac{1}{n_p} \sum_{j=1}^{n_p} p(y_k \mid \bar{v}_k^{(j)}, \mu_0^{(j)}, \mu_1^{(j)})$$

where $\bar{v}_k^{(j)} \sim F_{V_{\Delta_k}}$ for $k = 1, \dots, n$ and $j = 1, \dots, n_p$, with n_p the number of Monte Carlo samples (called particles thereafter).

In the Ornstein-Uhlenbeck model, the marginal likelihood can be approximated with a (bootstrap) sequential Monte Carlo algorithm, a standard inference technique for this class of models (Andrieu, Doucet and Holenstein, 2010; Jasra et al., 2011; Chopin, Jacob and Papaspiliopoulos, 2013). The resulting algorithm is known in this case as a particle marginal Metropolis-Hastings algorithm (Andrieu, Doucet and Holenstein, 2010).

4 Experiments

Priors In all the experiments, the drift μ_0 and the premium μ_1 parameters are set to zero. For the GGP model, we assume that we are in the infinite-activity regime, with $\sigma \geq 0$, and with finite variance, hence $\tau > 1$. The priors are set as follows:

$$\eta \sim \text{Gamma}(0.1, 0.1), \quad c \sim \text{Gamma}(0.1, 0.1), \quad (\tau - 1) \sim \text{Gamma}(1, 1), \quad \sigma \sim \text{Unif}(0, 1).$$

The more informative prior for τ reflects the empirical evidence that, for many financial datasets, the power-law exponent (2τ for the NGGP) is in the range (2, 5) (Cont and Tankov, 2004, Section 7.1). For the Ornstein-Uhlenbeck model, we additionally set $\lambda \sim \text{Gamma}(0.1, 0.1)$.

Software To fit both models, we use the Particles Library,¹ which allows to perform posterior inference in state-space models using particle MCMC algorithms. The code and datasets can be found on the anonymous github repository²

¹<https://github.com/nchopin/particles>

²<https://github.com/OxCSML-BayesNP/NGGP>

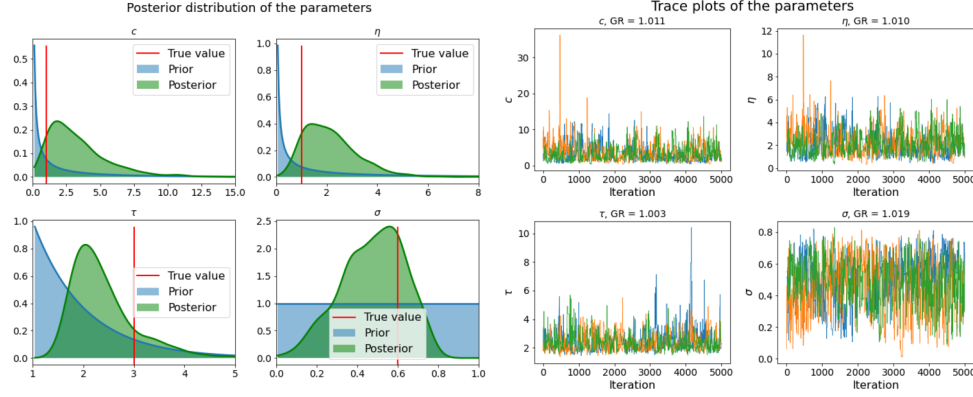


Figure 1: Histograms (left) and trace plots (right) of the posterior distributions of the parameters on the simulated data experiment. The blue line represents the value of the parameter used to generate the data. We also report the Gelman-Rubin scores to assess convergence of the chains (the lower the better, the empirical threshold for convergence is 1.1).

4.1 Exponentiated Lévy model

Simulated datasets

We first focus on the range of parameters that corresponds to processes with infinite activity ($\sigma \geq 0$) and finite variance ($\tau > 1$) as empirical evidence indicate that this is appropriate for financial applications (Cont and Tankov, 2004, Section 7.1). We generate a synthetic dataset of $n = 5\,000$ unit-spaced observations from the NGGP model, with parameters $\eta = 1$, $\sigma = 0.6$, $\tau = 3$ and $c = 1$. The priors are as described at the beginning of this section. We run three independent MCMC chains with $n_{mcmc} = 10\,000$ iterations each, of which 5000 iterations are used for burn-in. The number of particles to compute the marginal likelihood estimates is set to $n_p = 4\,000$. In Figure 1 we report histograms and trace plots of the posterior samples for each of the four parameters. Trace plots suggest the convergence of the MCMC algorithm.

To give a more complete picture, we also investigate whether the parameters can be recovered in the other three quadrants finite/infinite activity and finite/infinite variance. We consider for this priors with support in $(0, \infty)$ for τ and $(-\infty, 1)$ for σ . The estimated parameters and 95% credible intervals are reported in Table 2.

Further details and additional synthetic experiments, exploring the sensitivity to the choice of the prior, are provided in Supplementary material D.

Real-world datasets

Dataset We consider a dataset composed of the time-series of the stock prices of six large technology companies: Apple, Amazon, Facebook, Google, Microsoft and Netflix.

Regime	True values		Estimated values	
	σ	τ	σ	τ
finite-activity/finite-variance	-1	3	-1.00 (-1.23, -0.82)	3.03 (2.01, 5.50)
finite-activity/infinite-variance	-1	0.8	-1.00 (-1.22, -0.80)	0.77 (0.71, 0.84)
infinite-activity/infinite-variance	0.6	0.8	0.68 (0.40, 0.87)	0.75 (0.64, 0.87)

Table 2: Recovering the finite-activity/infinite-variance, finite-activity/finite-variance, and infinite-activity/infinite-variance settings. Posterior mean and 95% credible regions (in parenthesis) are reported.

The data are sampled every minute from the 10th of July 2019 until the 22nd of January 2020, with approximately 50 000 time points. We subsample 1 500 observations as training data to estimate the parameters of each model, and use the rest of the observations as test data.

Other models We compare the fit of the NGGP model to classical Lévy processes on the first dataset. The models compared are the normal-GG (also known as normal-tempered stable or normal-exponentially tilted stable model), with parameters η , σ and c ; two special cases of this model, the variance-gamma ($\sigma = 0$) and normal-inverse-Gaussian ($\sigma = 0.5$); the generalised hyperbolic (GH) model, with four parameters, and the student model, a special case of the GH model with two parameters. We use vague Gamma(0.1, 0.1) priors on all parameters, except for the parameter σ in the GG model, where a uniform prior on $[0, 1]$ is used as for the GGP, and for the degrees of freedoms ν_0 in the student t model, which tunes the power-law tail, where $(\nu_0/2 - 1) \sim \text{Gamma}(1, 1)$ to reflect the prior assumptions on the tails (as for τ in the NGGP). Note that we can compare here to the GH and student models as the observations are equally spaced.

Results We run 3 MCMC chains in parallel, with 5 000 iterations (2 500 burn-in) and 1 500 particles. The estimated parameters and 95% credible intervals for the parameters of the NGGP are reported in Table 3. The posterior mean for τ is around 2; this corresponds to a power-law exponent for X_t of around 4 which is in concordance with empirical observations (Cont and Tankov, 2004, Section 7.1). One exception is the Amazon stock, where τ is closer to 1, indicating a heavier tail. We first compare the models using the Kolmogorov-Smirnov (KS) statistics between the posterior predictive distribution and the empirical distribution of the test data. Results are reported in Table 4. The KS statistics is rather insensitive to the tail of the distribution, and the performances are similar for most models considered. To investigate the goodness-of-fit to the tails of the distribution, we compare the ranked empirical squared log-return to their posterior predictive distribution. Both GH, VG and NIG, which have exponentially decaying tails, provide similar results, and we only report the results of the GH. Results for Apple, Amazon, Facebook in Figure 2 (results for Google, Microsoft, Netflix are in the supplementary material section D). We can see that the NGGP model successfully captures the behaviour of tails for the different datasets, while the GH fails to provide accurate posterior predictive for some datasets such as Facebook. The NS model, which

Data	η	σ	τ	c
Apple	0.51, (0.13, 1.21)	0.44, (0.21, 0.61)	1.82, (1.11, 3.51)	1.14, (0.13, 4.26)
Amazon	0.64, (0.19, 1.40)	0.39, (0.14, 0.58)	1.18, (1.15, 3.59)	1.46, (0.22, 4.61)
Facebook	0.85, (0.34, 1.46)	0.25, (0.04, 0.51)	1.97, (1.25, 3.90)	1.86, (0.38, 4.63)
Google	0.18, (0.02, 0.71)	0.64, (0.50, 0.73)	1.93, (1.08, 4.16)	0.34, (0.01, 1.98)
Microsoft	0.27, (0.05, 0.81)	0.55, (0.39, 0.66)	1.98, (1.11, 4.35)	0.56, (0.05, 2.51)
Netflix	0.21, (0.07, 0.45)	0.54, (0.44, 0.63)	2.55, (1.25, 5.14)	0.29, (0.06, 0.90)

Table 3: Posterior mean and 95% credible interval for the four parameter of the NGGP model on the first dataset.

Data	NGGP	NGG	GH	NIG	NS	VG	Student
Apple	0.0194	0.0196	0.0194	0.0194	0.0196	0.0218	0.0196
Amazon	0.0087	0.0087	0.0085	0.0085	0.0159	0.0145	0.0092
Facebook	0.0181	0.0182	0.0182	0.0183	0.0245	0.1413	0.0181
Google	0.0205	0.0209	0.0197	0.0193	0.0237	0.0848	0.0200
Microsoft	0.0285	0.0285	0.0287	0.0286	0.0286	0.1567	0.0289
Netflix	0.0079	0.0080	0.0080	0.0084	0.0098	0.0162	0.0079
Mean	0.0172	0.0173	0.0171	0.0171	0.0204	0.0726	0.0173

Table 4: Kolmogorov-Smirnov distance between the empirical distribution of the test and the posterior predictive for different models on the first tech companies dataset (the smaller the better).

has the same parameter to capture the jump-activity and the tail behaviour, underestimates the value of the tail exponent, and gives a poor fit. The student t model tends to provide poor credible intervals, possibly due to the lack of flexibility of this two-parameter model.

4.2 Ornstein-Uhlenbeck based model

We now consider the Ornstein-Uhlenbeck based stochastic volatility model with NGGP marginal with $\sigma = 0$. As discussed in Section 3.2, in this case, the simulation of the state noise ε can be done exactly and exact posterior inference is possible. We compare the model with NGGP to the model with normal-gamma marginal (NG), described in Section 3, and demonstrate that NGGP better captures the heavy-tails with minimal computational overhead compared to NG.

Simulated dataset

We first demonstrate that our posterior sampler based on particle MCMC could successfully recover the true parameters on a simulated dataset. We simulate data from the Lévy-driven stochastic volatility model with NGGP marginal for $n = 2000$ time-steps and parameters $\eta = 5, c = 1, \tau = 3.0$ and $\lambda = 0.1$. We run three independent par-

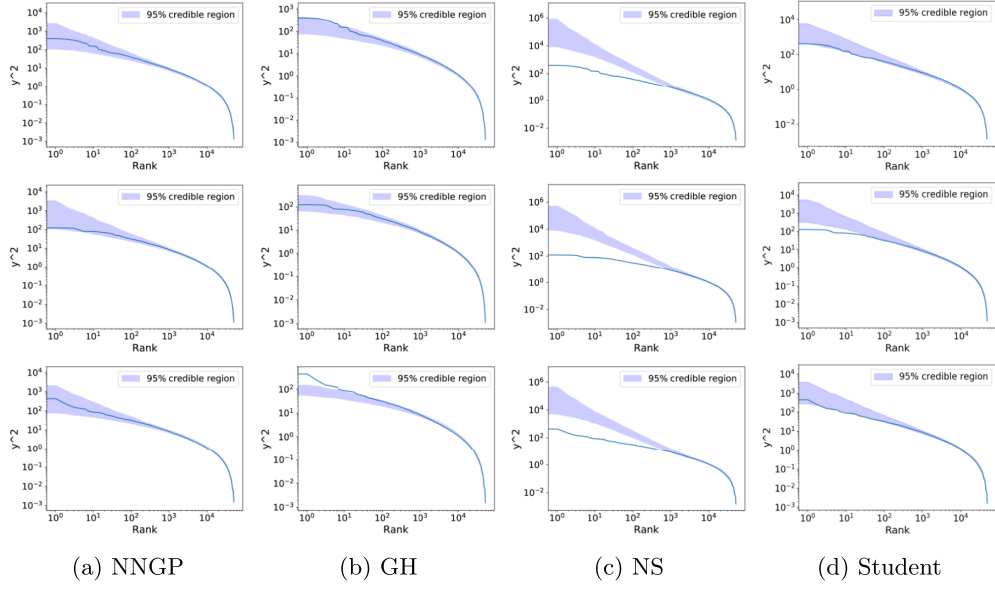


Figure 2: Ranked squared increments on the tech companies dataset. From top to bottom row: Apple, Amazon, Facebook. The line represents the ranked y^2 in the test dataset; the shaded area represent the 95% credible region.

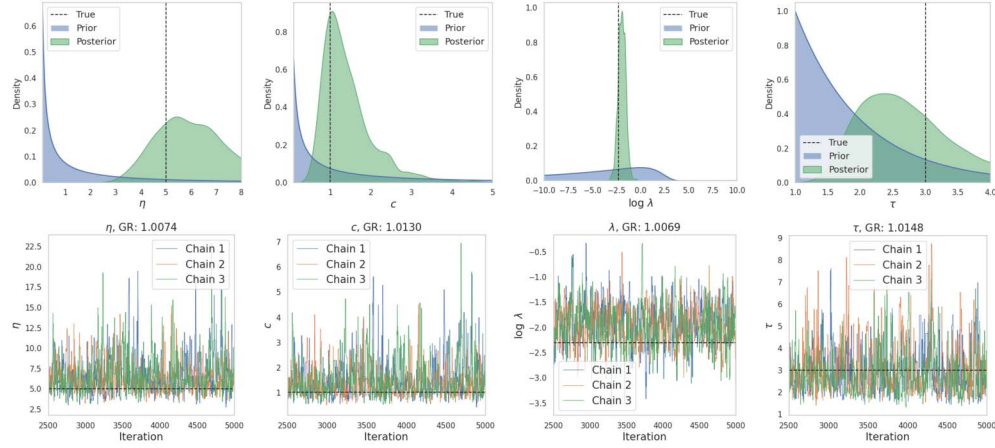


Figure 3: Posterior samples of the parameters on simulated data from OU-based stochastic volatility model with NNGP marginal.

ticle MCMC chains with 5000 iterations (2500 burn-in) and 3000 particles. Figure 3 shows that our sampler successfully recovers the parameters. Trace plots suggest the convergence of the sampler. Figure 4 shows the posterior estimate and credible interval




Figure 4: Posterior mean (solid red line) and 95% credible intervals (shaded area) of the integrated volatility. True volatility is in dashed green line.

for the integrated volatility, together with the true value. Additional simulation results are reported in the supplementary material E, for data generated with a smaller value $\tau = 1.5$. We also assess the sensitivity to the choice of prior, by reporting the posterior distributions under a different prior distribution for τ .

Real-world datasets

Dataset The dataset is obtained from the Realized library.³ We collected 14 daily stock data from 05-11-2007 to 07-10-2011 (around the time of subprime mortgage crisis), and fitted the Lévy driven stochastic volatility models on daily log return values. The data is accompanied with the estimates of the integrated variances \bar{v}_k for each day, obtained with an estimator based on higher-frequency data; we use these values as ground-truth of the integrated variance, and note it \bar{v}_k^{true} .

Evaluation metrics For $k = 1, \dots, n$, let $\bar{V}_k^{(1)}, \dots, \bar{V}_k^{(n_s)}$ denote the posterior samples of the integrated variance over the k th interval, where n_s is the number of MCMC iterations after burn-in. For $k = 1, \dots, n$ and any $\bar{v}_k \geq 0$, let $\hat{G}_k(\bar{v}_k) = \frac{1}{n_s} \sum_{i=1}^{n_s} \mathbb{1}_{\bar{V}_k^{(i)} \leq \bar{v}_k}$ be the Monte Carlo approximation of the posterior cumulative distribution function of the integrated variance $G_k(\bar{v}_k) = \Pr(\bar{V}_k \leq \bar{v}_k \mid y_1, \dots, y_n)$. For $k = 1, \dots, n$, denote

$$\zeta_k = 1 - \hat{G}_k(\bar{v}_k^{\text{true}}) = \frac{1}{n_s} \sum_{i=1}^{n_s} \mathbb{1}_{\bar{V}_k^{(i)} \geq \bar{v}_k^{\text{true}}}.$$

In order to assess the goodness-of-fit of the model, we calculate the KS statistics between the empirical distribution of $(\zeta_1, \dots, \zeta_n)$ and the distribution of a uniform random variable on $[0, 1]$. We also compare the fit of the model for different loss functions. Let L be a loss function. If $L(x, y) = (x - y)^2$ is the ℓ_2 loss, the Bayes estimator is the posterior mean. In case of $\ell_{1,\alpha}$ loss defined as

$$L(x, y) = \begin{cases} x - y & \text{if } x \geq y \\ \frac{1-\alpha}{\alpha} |x - y| & \text{if } x < y \end{cases},$$

³<https://realized.oxford-man.ox.ac.uk>

the Bayes estimator is the α -quantile $G_k^{-1}(\alpha)$ of G_k . We assess the fit of each model by computing the average loss

$$\frac{1}{n} \sum_{k=1}^n L(\bar{v}_k^{\text{true}}, \hat{v}_k^L)$$

where \hat{v}_k^L is the Bayes estimator under the loss L .

Results For every stock data, we run three independent chains of particle MCMC with 5 000 iterations (2 500 burn-in) and 1 000 particles. The estimated parameters and credible intervals are given in Table 5 (more datasets in Supplementary material F). The comparisons between the fits of the two models in terms of KS statistics and empirical loss, for different loss functions, are reported in Table 6. The model with NGGP marginal outperformed the one with NG marginal for all used metrics. Especially, since NGGP

	NG			NGGP			
	η	λ	c	η	λ	c	τ
AEX	1.43 (0.71, 2.50)	0.02 (0.01, 0.04)	1.41 (0.66, 2.56)	2.58 (1.17, 4.47)	0.03 (0.01, 0.05)	9.34 (3.00, 23.04)	1.49 (1.05, 2.37)
AORD	1.97 (0.92, 3.42)	0.02 (0.01, 0.03)	2.06 (0.88, 3.60)	3.62 (1.52, 6.87)	0.03 (0.01, 0.04)	11.26 (2.89, 29.19)	1.65 (1.09, 3.05)
DJI	1.28 (0.68, 2.05)	0.02 (0.01, 0.03)	1.26 (0.65, 2.10)	2.19 (1.10, 3.69)	0.03 (0.01, 0.04)	8.53 (2.30, 19.95)	1.38 (1.02, 2.25)
FTSE	1.34 (0.68, 2.22)	0.02 (0.01, 0.02)	1.46 (0.73, 2.37)	3.19 (1.35, 6.67)	0.02 (0.01, 0.04)	14.36 (3.55, 38.67)	1.36 (1.04, 2.08)
GSPTSE	1.23 (0.59, 2.13)	0.01 (0.01, 0.02)	1.28 (0.57, 2.31)	2.31 (0.92, 4.26)	0.02 (0.01, 0.03)	10.33 (2.42, 25.48)	1.42 (1.02, 2.29)

Table 5: Posterior mean and 95% credible intervals of the parameters of the NG and NGGP marginals for the different indices.

Data	KS($(\zeta_k), U(0,1)$)		ℓ_2		$\ell_{1,0.5}$		$\ell_{1,0.95}$		$\ell_{1,0.99}$	
	NG	NGGP	NG	NGGP	NG	NGGP	NG	NGGP	NG	NGGP
AEX	0.237	0.200	0.920	0.950	0.398	0.396	0.127	0.113	0.074	0.053
AORD	0.531	0.511	0.688	0.680	0.465	0.453	0.117	0.101	0.065	0.047
DJI	0.371	0.341	1.859	1.677	0.476	0.456	0.162	0.138	0.107	0.077
FTSE	0.269	0.241	2.590	2.510	0.479	0.480	0.186	0.159	0.134	0.092
GSPTSE	0.450	0.432	13.993	13.656	0.615	0.612	0.277	0.259	0.226	0.198
HSI	0.351	0.335	1.081	1.056	0.426	0.411	0.133	0.126	0.086	0.072
IBEX	0.264	0.245	0.824	0.788	0.422	0.413	0.122	0.106	0.067	0.045
IXIC	0.433	0.421	0.849	0.898	0.412	0.418	0.102	0.092	0.055	0.042
KS11	0.237	0.177	1.740	1.207	0.405	0.358	0.151	0.095	0.096	0.039
MXX	0.580	0.553	1.030	1.142	0.523	0.518	0.095	0.087	0.047	0.036
N225	0.283	0.230	0.674	0.807	0.360	0.362	0.087	0.073	0.045	0.030
RUT	0.570	0.392	1.217	1.192	0.449	0.454	0.079	0.072	0.031	0.025
SPX	0.388	0.337	1.318	1.317	0.440	0.435	0.131	0.115	0.082	0.058
SSMI	0.276	0.259	1.420	1.292	0.438	0.437	0.160	0.135	0.098	0.062
Mean	0.374	0.344	2.157	2.084	0.451	0.443	0.138	0.119	0.087	0.063
Mean	0.372	0.345	2.157	2.084	0.451	0.443	0.138	0.119	0.087	0.063

Table 6: Comparison of the fit of the NG and NGGP models under different metrics.




Figure 5: Posterior estimates (solid line) and 95% credible intervals of the integrated volatility under the NG (left) and NGGP (right) models for the AORD index. The true integrated volatility is represented by a green dashed line.

better captures the heavy-tails, the performance gap becomes more significant for the metrics emphasizing the heavy-tail regime $(\ell_{1,0.95}, \ell_{1,0.99})$. This is well highlighted in Figure 5, which gives the estimated volatility and credible intervals under both models for the AORD stock indices. NGGP in general better captures “spikes” in the log-return values while NG often fails to get credible interval with good coverage.

Comparison to ARMA-GARCH

We compare the OU-based stochastic volatility model to ARMA (1,1)-GARCH (1,1) on data from the Oxford Realized library. For this, we used 2,000 log-return values from the Oxford dataset, and split them into 1,100 training time steps and 900 test time steps. We fit the OU-based model with NG and NGGP marginals with our sampler. We report in Table 7 the marginal log-likelihood on the training data for the three models (more datasets are left to Appendix E). NGGP outperforms NG but performs slightly worse than ARMA-GARCH. We further compare the models in terms of prediction. We generate one-step predictions for the 900 test time-steps, compute Value at Risk (VaR) values for each time step prediction, and counted the fraction of actual test data less than or equal to the negative of VaR values. This is to see whether estimated VaR values fit the test data well by checking

$$\Pr(-Y \leq \text{VaR}_\alpha) = 1 - \alpha \iff \Pr(Y > -\text{VaR}_\alpha) = 1 - \alpha \iff \Pr(Y \leq -\text{VaR}_\alpha) = \alpha,$$

so that the fraction of test data less than or equal to $-\text{VaR}_\alpha$ being closer to α means better prediction. For all models, we collected posterior samples, conducted prediction for each posterior samples using corresponding model parameters and state estimates, computed empirical CDFs using those samples, and computed VaR values. As summarised in Table 8 (more datasets are left to supplementary material, section E), the results of the different methods are comparable.

Supplementary Material

Supplementary material for “The Normal-Generalised Gamma-Pareto process: A novel pure-jump Lévy process with flexible tail and jump-activity properties” (DOI: [10.1214/22-BA1343SUPP](https://doi.org/10.1214/22-BA1343SUPP); .pdf). Complementary theoretical and empirical results.

Data	NG	NGGP	ARMA-GARCH
AEX	-1466.021	-1465.811	-1459.768
AORD	-1495.441	-1494.830	-1492.148
DJI	-1442.606	-1438.013	-1424.026
FTSE	-1448.833	-1445.721	-1437.881
GSPTSE	-1455.563	-1454.125	-1445.307

Table 7: Comparison of the marginal log-likelihood values of OU-NG, OU-NGGP and ARMA-GARCH models on data from the Realized-library.

	NG	NGGP	ARMA -GARCH	NG	NGGP	ARMA -GARCH
Data	$\alpha = 0.95$			$\alpha = 0.99$		
AEX	0.962	0.958	0.969	0.993	0.993	0.992
AORD	0.962	0.960	0.970	0.996	0.993	0.991
DJI	0.956	0.959	0.968	0.996	0.994	0.989
FTSE	0.956	0.953	0.950	0.994	0.993	0.983
GSPTSE	0.967	0.968	0.978	0.997	1.000	0.996

Table 8: The results of VaR test on predicted sequence. Values closer to α mean better VaR prediction.

References

- AALLEN, O. (1992). Modelling heterogeneity in survival analysis by the compound Poisson distribution. *The Annals of Applied Probability* 951–972. [MR1189425](#). 3
- AÏT-SAHALIA, Y. and JACOD, J. (2009). Estimating the degree of activity of jumps in high frequency data. *The Annals of Statistics* **37** 2202–2244. [MR2543690](#). doi: <https://doi.org/10.1214/08-AOS640>. 2
- ANDRIEU, C., DOUCET, A. and HOLENSTEIN, R. (2010). Particle Markov chain Monte Carlo methods. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* **72** 269–342. [MR2758115](#). doi: <https://doi.org/10.1111/j.1467-9868.2009.00736.x>. 16, 18
- ANDRIEU, C. and ROBERTS, G. O. (2009). The pseudo-marginal approach for efficient Monte Carlo computations. *The Annals of Statistics* **37** 697–725. [MR2502648](#). doi: <https://doi.org/10.1214/07-AOS574>. 17
- AYED, F., LEE, J. and CARON, F. (2019). Beyond the Chinese Restaurant and Pitman-Yor processes: Statistical Models with double power-law behavior. In *International Conference on Machine Learning* 395–404. 3, 4, 6, 7, 13
- AYED, F., LEE, J. and CARON, F. (2022). “Supplementary Material for “The Normal-Generalised Gamma-Pareto Process: A novel Pure-Jump Lévy Process with Flexible Tail and Jump-Activity Properties”.” *Bayesian Analysis*. doi: <https://doi.org/10.1214/22-BA1343SUPP>. 10, 13

- BARNDORFF-NIELSEN, O. E. (1997). Processes of normal inverse Gaussian type. *Finance and stochastics* **2** 41–68. MR1804664. doi: <https://doi.org/10.1007/s007800050032>. 2
- BARNDORFF-NIELSEN, O. E. and SHEPHARD, N. (2001). Non-Gaussian Ornstein–Uhlenbeck-based models and some of their uses in financial economics. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* **63** 167–241. MR1841412. doi: <https://doi.org/10.1111/1467-9868.00282>. 1, 3, 8, 15, 16
- BARNDORFF-NIELSEN, O. and SHEPHARD, N. (2002). Normal modified stable processes. *Theory of Probability and Mathematical Statistics* **65** 1–20. MR1936123. 4
- BEAUMONT, M. A. (2003). Estimation of population growth or decline in genetically monitored populations. *Genetics* **164** 1139–1160. 17
- BELOMESTNY, D. (2010). Spectral estimation of the fractional order of a Lévy process. *The Annals of Statistics* **38** 317–351. MR2589324. doi: <https://doi.org/10.1214/09-AOS715>. 2
- BELOMESTNY, D. and PANOV, V. (2013). Estimation of the activity of jumps in time-changed Lévy models. *Electronic Journal of Statistics* **7** 2970–3003. MR3151759. doi: <https://doi.org/10.1214/13-EJS870>. 2
- BINGHAM, N. H., GOLDIE, C. M. and TEUGELS, J. L. (1989). *Regular variation* **27**. Cambridge university press. MR0898871. doi: <https://doi.org/10.1017/CB09780511721434>. 6, 7
- BLATTBERG, R. C. and GONEDES, N. J. (1974). A comparison of the stable and student distributions as statistical models for stock prices. *The journal of business* **47** 244–280. 2
- BONDESSON, L. (1992). *Generalized gamma convolutions and related classes of distributions and densities* **76**. Lecture notes in Statistics. Springer Verlag. MR1224674. doi: <https://doi.org/10.1007/978-1-4612-2948-3>. 13
- BRIX, A. (1999). Generalized gamma measures and shot-noise Cox processes. *Advances in Applied Probability* **31** 929–953. MR1747450. doi: <https://doi.org/10.1239/aap/1029955251>. 3
- CARR, P., GEMAN, H., MADAN, D. B. and YOR, M. (2002). The fine structure of asset returns: An empirical investigation. *The Journal of Business* **75** 305–332. 1, 2
- CHAMPAGNAT, N., DEACONU, M., LEJAY, A. E., NAVET, N. and BOUKHEROUAA, S. (2013). An empirical analysis of heavy-tails behavior of financial data: The case for power laws. 2
- CHOPIN, N., JACOB, P. E. and PAPASPILOPOULOS, O. (2013). SMC²: an efficient algorithm for sequential analysis of state space models. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* **75** 397–426. MR3065473. doi: <https://doi.org/10.1111/j.1467-9868.2012.01046.x>. 18

- CONT, R., POTTERS, M. and BOUCHAUD, J. P. (1997). Scaling in stock market data: stable laws and beyond. In *Scale invariance and beyond* 75–85. Springer. 2
- CONT, R. and TANKOV, P. (2004). *Financial modelling with jump processes* 2. CRC press. 1, 2, 8, 9, 14, 18, 19, 20
- ŞİMŞEKLI, U. (2017). Fractional Langevin Monte Carlo: Exploring Lévy driven stochastic differential equations for Markov Chain Monte Carlo. In *Proceedings of the 34th International Conference on Machine Learning-Volume 70* 3200–3209. JMLR. org. 1
- DEVROYE, L. (2009). Random variate generation for exponentially and polynomially tilted stable distributions. *ACM Transactions on Modeling and Computer Simulation (TOMACS)* 19 18. doi: <https://doi.org/10.1145/1596519.1596523>. 4, 7
- EBERLEIN, E., KELLER, U. and PRAUSE, K. (1998). New insights into smile, mispricing, and value at risk: The hyperbolic model. *The Journal of Business* 71 371–405. doi: <https://doi.org/10.1086/209749>. 2
- FRÜHWIRTH-SCHNATTER, S. and SÖGNER, L. (2009). Bayesian estimation of stochastic volatility models based on OU processes with marginal Gamma law. *Annals of the Institute of Statistical Mathematics* 61 159–179. MR2481032. doi: <https://doi.org/10.1007/s10463-007-0130-8>. 16
- GANDER, M. and STEPHENS, D. A. (2007). Stochastic volatility modelling in continuous time with general marginal distributions: Inference, prediction and model selection. *Journal of Statistical Planning and Inference* 137 3068–3081. MR2364151. doi: <https://doi.org/10.1016/j.jspi.2006.07.015>. 16
- GNEDIN, A., HANSEN, B. and PITMAN, J. (2007). Notes on the occupancy problem with infinitely many boxes: general asymptotics and power laws. *Probability surveys* 4 146–171. MR2318403. doi: <https://doi.org/10.1214/07-PS092>. 6
- GRIFFIN, J. E. and STEEL, M. F. J. (2006). Inference with non-Gaussian Ornstein–Uhlenbeck processes for stochastic volatility. *Journal of Econometrics* 134 605–644. MR2328421. doi: <https://doi.org/10.1016/j.jeconom.2005.07.007>. 16
- GRIGELIONIS, B. (2007). Extended Thorin classes and stochastic integrals. *Lithuanian Mathematical Journal* 47 406–411. MR2392715. doi: <https://doi.org/10.1007/s10986-007-0028-4>. 13
- HOFERT, M. (2011). Sampling exponentially tilted stable distributions. *ACM Transactions on Modeling and Computer Simulation (TOMACS)* 22 3. MR2955859. doi: <https://doi.org/10.1145/2043635.2043638>. 4, 7
- HOUGAARD, P. (1986). Survival models for heterogeneous populations derived from stable distributions. *Biometrika* 73 387–396. MR0855898. doi: <https://doi.org/10.1093/biomet/73.2.387>. 3
- HUANG, J. Z. and WU, L. (2004). Specification analysis of option pricing models based on time-changed Lévy processes. *The Journal of Finance* 59 1405–1439. 1

- JAMES, L. F., ROYNETTE, B. and YOR, M. (2008). Generalized gamma convolutions, Dirichlet means, Thorin measures, with explicit examples. *Probability Surveys* **5** 346–415. [MR2476736](#). doi: <https://doi.org/10.1214/07-PS118>. 13
- JAMES, L. F. and ZHANG, Z. (2011). Quantile clocks. *The Annals of Applied Probability* **21** 1627–1662. [MR2884047](#). doi: <https://doi.org/10.1214/10-AAP752>. 14
- JASRA, A., STEPHENS, D. A., DOUCET, A. and TSAGARIS, T. (2011). Inference for Lévy-driven stochastic volatility models via adaptive sequential Monte Carlo. *Scandinavian Journal of Statistics* **38** 1–22. [MR2760137](#). doi: <https://doi.org/10.1111/j.1467-9469.2010.00723.x>. 18
- JING, B. Y., KONG, X. B. and LIU, Z. (2012). Modeling high-frequency financial data by pure jump processes. *The Annals of Statistics* **40** 759–784. [MR2933665](#). doi: <https://doi.org/10.1214/12-AOS977>. 1
- JUREK, Z. J. (2001). Remarks on the selfdecomposability and new examples. *Demonstratio Mathematica* **34** 29–38. [MR1833180](#). 8
- LANDIS, M. J., SCHRAIBER, J. G. and LIANG, M. (2012). Phylogenetic analysis using Lévy processes: finding jumps in the evolution of continuous traits. *Systematic biology* **62** 193–204. 1
- LANDIS, M. J. and SCHRAIBER, J. G. (2017). Pulsed evolution shaped modern vertebrate body sizes. *Proceedings of the National Academy of Sciences* **114** 13224–13229. 1
- LIANG, C. and LI, S. (2015). Option pricing and hedging in incomplete market driven by Normal Tempered Stable process with stochastic volatility. *Journal of mathematical Analysis and Applications* **423** 701–719. [MR3273203](#). doi: <https://doi.org/10.1016/j.jmaa.2014.10.033>. 4
- MADAN, D. B., CARR, P. and CHANG, E. (1998). The variance gamma process and option pricing. *Review of Finance* **2** 79–105. doi: <https://doi.org/10.1023/A:1009703431535>. 1
- MANDELBROT, B. (1963). THE VARIATION OF CERTAIN SPECULATIVE PRICES. *The Journal of Business* **36** 394–419. 2
- ROBERTS, G. O., PAPASPILIOPOULOS, O. and DELLAPORTAS, P. (2004). Bayesian inference for non-Gaussian Ornstein–Uhlenbeck stochastic volatility processes. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* **66** 369–393. [MR2062382](#). doi: <https://doi.org/10.1111/j.1369-7412.2004.05139.x>. 16
- ROSIŃSKI, J. (2007). Tempering stable processes. *Stochastic processes and their applications* **117** 677–707. [MR2327834](#). doi: <https://doi.org/10.1016/j.spa.2006.10.003>. 4, 13
- SATO, K. I. (2001). Subordination and self-decomposability. *Statistics & probability letters* **54** 317–324. [MR1857946](#). doi: [https://doi.org/10.1016/S0167-7152\(01\)00110-9](https://doi.org/10.1016/S0167-7152(01)00110-9). 11

- THORIN, O. (1977). On the infinite divisibility of the lognormal distribution. *Scandinavian Actuarial Journal* **1977** 121–148. [MR0552135](#). doi: <https://doi.org/10.1080/03461238.1977.10405635>. 13
- WOERNER, J. (2011). Analyzing the fine structure of continuous time stochastic processes. In *Seminar on Stochastic Analysis, Random Fields and Applications VI* 473–492. Springer. [MR2857040](#). doi: https://doi.org/10.1007/978-3-0348-0021-1_26. 2
- ZHANG, A. and PAISLEY, J. (2018). Deep Bayesian Nonparametric Tracking. In *International Conference on Machine Learning* 5828–5836. 1

Acknowledgments

The authors thank Cian Naik, Lancelot James and Matthias Winkel for useful feedback on an earlier version of this article.