



Fast and general simulation of Lévy-driven Ornstein-Uhlenbeck processes for energy derivatives

Roberto Baviera ^a, Pietro Manzoni ^{b,*}

^a Politecnico di Milano, Department of Mathematics, Italy

^b University of Edinburgh, Business School, United Kingdom

ARTICLE INFO

JEL classification:

C63. Computational techniques
Simulation modeling
G13. Contingent pricing
futures pricing

Keywords:

Lévy-driven Ornstein-Uhlenbeck
Energy derivatives
Fourier methods
FFT
Monte Carlo

ABSTRACT

Lévy-driven Ornstein-Uhlenbeck (OU) processes represent a versatile class of stochastic processes that have garnered interest in the energy sector for their ability to capture typical features of market dynamics. However, in the current state of play, their Monte Carlo simulation is not straightforward for two main reasons: *i*) algorithms are only available for some specific processes within this class, and *ii*) they are often computationally expensive.

In this paper, we introduce a new simulation technique designed to address both challenges. It relies on the numerical inversion of the characteristic function, offering a general methodology applicable to all Lévy-driven OU processes. Moreover, leveraging FFT, the proposed methodology ensures fast and accurate simulations, providing a solid basis for the widespread adoption of these processes in the energy sector. Lastly, the algorithm allows explicit control of the numerical error.

We apply the proposed technique to the pricing of energy derivatives, comparing the results with the existing benchmarks. Our findings indicate that the proposed methodology is at least one order of magnitude faster than the existing algorithms, while maintaining an equivalent level of accuracy.

1. Introduction

Energy markets display rich dynamics characterised by mean reversion, jumps, and heavy tails, posing challenges to traditional modelling approaches. Classical stochastic models, such as plain Lévy and Gaussian Ornstein-Uhlenbeck (OU) processes, often fail to adequately capture these non-standard behaviours, motivating the exploration of non-Gaussian alternatives (see, e.g., [1]). Lévy-driven OU processes have emerged as a promising class of models capable of reproducing these key features. They are built upon the OU stochastic differential equation (SDE) – well-known for its mean-reverting behaviour – and incorporate Lévy processes, which allow for the presence of jumps and heavy-tailed distributions, enabling a comprehensive representation of energy market dynamics (see, e.g., [2]).

The objective of this study is to develop an efficient procedure for simulating the trajectories of Lévy-driven OU processes on a discrete time-grid. Such simulation is essential for Monte Carlo (MC) methods, in particular for pricing discretely-monitored derivatives – a class of instruments which is highly traded in energy markets (see, e.g., [3]). MC methods are employed by risk managers to derive information about the distribution of their portfolios and to gain a view on their overall risk profile. These portfolios often involve several thousands of discretely-monitored options; thus, it is crucial to rely on simulation schemes that are not only accurate but also fast, because the models need to run for different parameter settings and market conditions.

* Corresponding author.

E-mail addresses: roberto.baviera@polimi.it (R. Baviera), pmanzoni@ed.ac.uk (P. Manzoni).

In the literature, the simulation of Lévy-driven OU processes is realised via Exact Decomposition (ED) algorithms. These methods decompose a Lévy-driven OU into simpler subprocesses, for which simulation schemes exist (see, e.g., [4–16]). However, the ED schemes are not general: each one is specifically tailored to a particular Lévy-driven OU process, and as a result, they are only available for a very limited set of processes.

In this paper, we present a novel simulation approach that offers three key advantages: *i)* it is general, as it is applicable to the whole class of Lévy-driven OU processes; *ii)* it is fast, achieving runtimes at least one order of magnitude lower than existing ED schemes; *iii)* it ensures complete control over the numerical error, thereby providing a powerful and reliable simulation tool.

The proposed method is based on a double inversion procedure. Sample paths are simulated by inverting the Cumulative Distribution Function (CDF) of the process increments, leveraging the Probability Integral Transform. However, for Lévy-driven OU, the CDF of the increments is not explicitly available, but is known only through the characteristic function (CF). This requires an additional inversion – from the CF to the CDF – which makes path simulation non-trivial, as already documented in the literature on plain Lévy processes (see, e.g., [17,18]).

The numerical challenges persist, and are in fact compounded, when moving from the plain Lévy to the more general Lévy-driven OU framework, posing new difficulties in terms of both algorithmic stability and estimation of the numerical error. To overcome these issues effectively, in our approach we directly reconstruct the CDF of the increments from their complex-shifted CF.

The main contributions of this paper are threefold. First, we develop a *fast* and *general* simulation algorithm for Lévy-driven OU processes that provides explicit control over the numerical error. The stability of our method is achieved via a complex-plane shift – a technique enabled by new analytical results that we establish in this work – yielding accurate and efficient MC simulation.

Second, we design a specialised algorithm for the simulation of an important subclass of Lévy-driven OU processes, namely those with finite-activity (FA), characterised by a discontinuous CDF. This algorithm introduces a Bernoulli-based scheme specifically tailored for simulating these processes.

Third, we conduct a thorough experimental comparison of our method against existing ED algorithms in the context of energy derivatives pricing. We demonstrate that our method achieves equivalent accuracy while delivering runtime at least one order of magnitude faster across different processes and parameter regimes. To our knowledge, this is the first work to successfully apply complex-integration techniques to the simulation of Lévy-driven OU processes.

The rest of the paper is organised as follows. [Section 2](#) provides an overview of the proposed simulation method, illustrating its core steps via a benchmark example. [Section 3](#) introduces Lévy-driven OU processes along with new key analytical properties. [Section 4](#) describes the simulation methodology, presenting our unified framework for the efficient simulation of Lévy-driven OU processes. [Section 5](#) establishes error bounds and discusses implementation details. [Section 6](#) presents applications to energy derivative pricing and provides performance comparison against alternative simulation techniques. Finally, [Section 7](#) summarises our findings.

2. Overview of the method

A Lévy-driven OU process is a mean-reverting process defined as the solution to the SDE

$$dX_t = -bX_t dt + dL_t, \quad (1)$$

where L_t is a Lévy process and $b > 0$ controls the mean-reversion speed (see, e.g., [19]). While their theoretical properties are well-established in the literature (see, e.g., [1]), achieving an accurate numerical implementation of Lévy-driven OU processes remains a major challenge, as their distributions over time are generally not accessible in closed form.

How, then, can these processes be simulated efficiently? The prevailing approach in the literature is to employ ED algorithms – schemes that represent a Lévy-driven OU as a sum of components which can be simulated individually. [Table 1](#) summarises the main ED algorithms, which primarily focus on two Lévy drivers: the Tempered Stable (TS) and the Normal Tempered Stable (NTS). Each referenced paper presents an ED algorithm designed for a specific process, often valid only within restricted parameter domains. Simulations typically rely on rejection sampling, and are therefore time-consuming and highly dependent on the choice of process parameters (see, e.g., [20]).

This paper introduces a fundamentally different approach that applies to all Lévy-driven OU processes, combining generality with a computational speed comparable to that of the Gaussian OU, i.e. the process driven by a Brownian Motion (BM). To illustrate our method concretely, it is helpful to first focus on this well-known benchmark case: when L_t is a BM, the solution to the SDE (1) over any interval $[r, t]$ of length $\Delta t := t - r$ reads

$$X_t = X_r e^{-b\Delta t} + \int_r^t e^{-b(t-s)} dL_s, \quad (2)$$

where the stochastic integral has a centred Gaussian law with variance $\sigma^2(\Delta t) := \frac{1-e^{-2b\Delta t}}{2b}$. Hence, simulating each increment of the process reduces to sampling a single Gaussian random variable (r.v.). This approach is simple and efficient, but it requires explicit knowledge of the increment distribution – a luxury we lack for non-Gaussian OU processes.

Table 1

Existing literature on ED algorithms for Lévy-driven OU simulation. The rows list the four processes considered in these studies (OU-TS, TS-OU, OU-NTS, and NTS-OU), rigorously defined in Appendix B; the columns classify them according to their jump properties. Finally, the symbol $\times$ indicates that no ED scheme exists in the literature. Our general method allows the simulation of any of the processes in the table.

Process	Finite-Activity	Infinite-Activity	
		Finite-Variation	Infinite-Variation
OU-TS	Sabino [12]	Qu et al. [10] Sabino and Cufaro Petroni [16]	$\times$
TS-OU	Sabino and Cufaro Petroni [15]	Zhang and Zhang [4,5] Kawai and Masuda [7] Qu et al. [10] Sabino and Cufaro Petroni [16]	$\times$
OU-NTS	$\times$	Sabino [11,14]	Sabino [14]
NTS-OU	Bianchi et al. [9] Sabino and Cufaro Petroni [15]	Sabino [13]	Sabino [13]

Our simulation method overcomes this limitation by requiring only the CF of the stochastic integral in (2), which can be easily derived for all Lévy-driven OU processes, as we discuss in the next section.¹ The simulation consists of two main steps. First, we compute the CDF $P(x)$ of the stochastic integral in (2) from its CF $\phi(u, \Delta t)$; the natural starting point for this operation is the basic formula:

$$P(x) = \frac{1}{2} - \frac{1}{2\pi} \int_{-\infty}^{+\infty} \frac{e^{-iux} \phi(u, \Delta t)}{iu} du, \quad (3)$$

valid when $\phi(u, \Delta t)$ is sufficiently regular (cf. [21,22]). Second, we draw a uniform r.v. $U \sim \mathcal{U}(0, 1)$ and obtain X_t by computing $P^{-1}(U)$ and adding the deterministic term $X_r e^{-b\Delta t}$.

The strength of this approach lies in its broad applicability to Lévy-driven OU processes, provided that three technical considerations are properly addressed:

- (i) For a generic Lévy driver, the basic inversion formula in (3) is not well-posed, leading to instability in the computation of the CDF. A general formula for $P(x)$ is obtained by suitably shifting the integration path in the complex plane; we present the details of this procedure in Section 4.
- (ii) The CDF $P(x)$ must be evaluated on a fine grid $x_1, \dots, x_N$, which is computationally expensive if done point-by-point. In our method, we perform this task through the Fast Fourier Transform (FFT; [23]), which efficiently computes all points simultaneously and yields numerical estimates of the CDF $\hat{P}(x_1), \dots, \hat{P}(x_N)$.
- (iii) For fast sampling, efficient inversion of the CDF is required. We achieve this via cubic spline interpolation of the computed values $\hat{P}(x_1), \dots, \hat{P}(x_N)$, which reconstructs the CDF over its domain and allows immediate evaluation of $\hat{P}^{-1}(U)$ for any uniform U (see, e.g., [24]).

The Basic Algorithm below summarises the procedure.

Basic Algorithm for generating Lévy-driven OU processes

Require: X_r (Value at time r), t (Target time), $\Delta t := t - r$ (Simulation step)

- 1: Compute the CF $\phi(u, \Delta t)$ of the integral process $\int_r^t e^{-b(t-s)} dL_s$
- 2: Retrieve the approximated CDF $\hat{P}$ via FFT inversion of $\phi(u, \Delta t)$
- 3: Interpolate the obtained CDF using splines
- 4: Draw a uniform r.v. $U \sim \mathcal{U}(0, 1)$ and compute $\hat{P}^{-1}(U)$
- 5: Update $X_t = X_r e^{-b\Delta t} + \hat{P}^{-1}(U)$

In this section, we have offered an overview of the paper's content, considering a special example of Lévy-driven OU process: the Gaussian OU. The next section broadens the focus to the general case, wherein any Lévy process L_t can be chosen as the stochastic driver in (1), and discusses relevant properties of the CF of Lévy-driven OU processes. Readers primarily interested in the simulation methodology may proceed directly to Section 4, as the intervening section provides supporting theoretical results that can be consulted as needed.

¹ For example, in the considered Gaussian OU case, the CF takes the familiar form: $\phi(u, \Delta t) = e^{-\frac{1}{2}u^2 \sigma^2(\Delta t)}$.

3. The processes and their properties

Let $(\Omega, \mathcal{F}, (F_t)_t, \mathbb{P})$ be a filtered probability space, where F_t is the available information at time t , and let L_t be an F_t -adapted pure-jump Lévy process with Lévy density $\nu_L(x)$. For a specified initial condition X_0 , a Lévy-driven OU process is defined as the unique strong solution X_t of the SDE (1), given by

$$X_t = X_0 e^{-bt} + Z_t, \quad \text{with} \quad Z_t := \int_0^t e^{-b(t-s)} dL_s, \quad (4)$$

where the integral process Z_t is an additive process (see, e.g., [25]).

As discussed in Barndorff-Nielsen and Shephard [19], Lévy-driven OU processes can be constructed in two distinct ways, giving rise to two broad classes of processes: *OU-Lévy* and *Lévy-OU*. The former are constructed by specifying the Lévy driver L_t – and then deriving the resulting dynamics of X_t . The latter are constructed by specifying the stationary distribution of the solution X_t – from which the corresponding Lévy driver L_t is then determined.

This paper considers four prominent processes: OU-TS and OU-NTS (processes of OU-Lévy type, obtained using TS and NTS as Lévy drivers), TS-OU and NTS-OU (processes of Lévy-OU type, obtained using TS and NTS laws as stationary distributions). These are the processes listed in Table 1, for which exact ED schemes exist; detailed definitions are provided in Appendix B.

In this section, we recall the general expression for the CFs of both classes, specifying their relationships with either the Lévy driver (for OU-Lévy) or the stationary distribution (for Lévy-OU). Most importantly, we prove two new results on the regularity of these CFs, which are instrumental to our simulation method.

Lévy-driven OU are naturally handled through their CFs. Throughout the paper, we adopt notation consistent with the reference book by Cont and Tankov [26]: given a process L_t , we denote its CF by $\phi_L(u, t)$, its log-characteristic function (LCF) by $\Psi_L(u, t)$ and, if L_t is Lévy, its characteristic exponent by $\psi_L(u)$. In the latter case, time-homogeneity implies that $\Psi_L(u, t) = t\psi_L(u)$.

3.1. OU-Lévy processes

Processes in the OU-Lévy class are constructed by specifying the dynamics of the Lévy driver L_t . Since the increments of L_t are stationary and independent, the LCF of the integral process Z_t in (4) reads

$$\Psi_Z(u, t) = \int_0^t \psi_L(ue^{-bs}) ds, \quad (5)$$

where $\psi_L(\cdot)$ is the characteristic exponent of the Lévy driver L_t (cf. [27, Theorem 1, p.411]). The CF of any OU-Lévy can thus be derived directly from the CF of its driver.

An important concept in the theory of CFs is the *analyticity strip*. Although the CF $\phi(u) = \mathbb{E}[e^{iuY}]$ of a r.v. Y is, in its standard formulation, defined for $u \in \mathbb{R}$, it is generally possible to extend its domain to complex arguments by analytic continuation. The analyticity strip is defined as the maximal horizontal strip in the complex plane, containing the real axis, where $\phi(u)$ is holomorphic (see, e.g., [28, 29]). We denote this strip by the open interval $(-p_+, p_-)$ with $p_-, p_+ > 0$, corresponding to the range of imaginary parts $\text{Im}(u)$ for which the CF remains analytic.

In the next proposition, we prove a new key property regarding the analyticity strip of OU-Lévy processes.

Proposition 1. *Let Z_t be the integral process of an OU-Lévy driven by L_t . Then, for every $t > 0$, the analyticity strip of $\phi_Z(u, t)$ coincides with that of $\phi_L(u)$.*

Proof. See Appendix A. $\square$

Hence, in the OU-Lévy framework, the integral process Z_t inherits the domain of analyticity of its Lévy driver L_t , implying that the CF $\phi_Z(u, t)$ can be analytically continued to the same strip as $\phi_L(u)$, while remaining free of singularities. This new result both establishes the validity of complex-integration techniques for OU-Lévy processes and determines the precise region in which they can be applied.²

We also recall a useful structural relationship from the literature on OU-Lévy processes: the Lévy density of Z_t can be expressed in terms of that of the driver L_t as

$$\nu_Z(x, t) = \frac{1}{bx} \int_x^{xe^{bt}} \nu_L(y) dy, \quad (6)$$

with $\nu_Z(0, t) = 0$ (see, e.g., [26, Proposition 15.1, p.482]). Hence, the cumulants³ $c_k(Z_t)$ of the process Z_t can be obtained from the cumulants of the driver at time $t = 1$, denoted $c_k(L_1)$:

$$c_k(Z_t) = \frac{1 - e^{-kbt}}{kb} c_k(L_1) \quad \forall k = 1, 2, \dots \quad (7)$$

² For the purposes of this work, we focus on processes whose analyticity strip is non-empty. This requirement coincides with requiring that the process has exponentially decaying tails (see, e.g., [30]), a standard assumption in the literature and satisfied by all processes of practical interest (see, e.g., [26]).

³ The cumulants $c_k(Y)$ of a r.v. Y with LCF $\Psi_Y(u)$ are defined for $k = 1, 2, \dots$ as $c_k(Y) := (-i)^k \frac{d^k}{du^k} \Psi_Y(u) \Big|_{u=0}$.

3.2. Lévy-OU processes

Processes in the Lévy-OU class are specified by the stationary distribution of the process X_t – denoted in the following as the r.v. X – rather than by directly choosing a Lévy driver L_t . However, not every stationary distribution is allowed: to ensure the existence of a compatible Lévy driver, X must be self-decomposable (cf. [31, Theorem 1, p.306]). For Lévy-OU processes, the LCF of Z_t in (4) is related to the characteristic exponent of X by the following relationship (cf. [19, Theorem 1, p.172]):

$$\Psi_Z(u, t) = \psi_X(u) - \psi_X(ue^{-bt}). \quad (8)$$

In the next proposition, we establish a new important property for Lévy-OU processes, which is analogous to that given in Proposition 1 for the OU-Lévy class.

Proposition 2. *Let Z_t be the integral process of a Lévy-OU with stationary distribution X . Then, for every $t > 0$, the analyticity strip of $\phi_Z(u, t)$ coincides with that of $\phi_X(u)$.*

Proof. See Appendix A. $\square$

This result ensures that – also in the Lévy-OU case – Z_t possesses a domain of analyticity, which is crucial for the application of complex-integration in our simulation. Furthermore, this domain is straightforward to identify, as it coincides precisely with the analyticity strip of the stationary distribution X .

In addition, we recall a known relationship between the Lévy density of Z_t and that of the stationary distribution X (see, e.g., [16, Proposition 1, p.2524]):

$$\nu_Z(x, t) = \nu_X(x) - \nu_X(xe^{bt})e^{bt}. \quad (9)$$

It follows that the cumulants $c_k(Z_t)$ of the process Z_t can be expressed in terms of those of X , as:

$$c_k(Z_t) = (1 - e^{-kbt}) c_k(X) \quad \forall k = 1, 2, \dots \quad (10)$$

In this section, we have reviewed the key properties of Lévy-driven OU processes and proved new results on their analyticity strips. The former are essential for deriving the LCF, Lévy density, and cumulants of any such process (in particular, explicit LCFs for the main processes are available in Appendix B); the latter are instrumental for the simulation algorithm presented in the next section.

4. The simulation algorithm

The simulation of the process X_t in (4) on the time-grid $t_0, t_1, \dots, t_Q$ (generally, the reset dates of the discretely-monitored derivative) leverages the exact solution of the SDE. Starting from the given initial condition $X_{t_0} = X_0$, we generate the sample paths by iterating, for $j = 0, \dots, Q - 1$, the equation:

$$X_{t_{j+1}} = e^{-b\Delta t_j} X_{t_j} + Z_{\Delta t_j} \quad \text{with} \quad Z_{\Delta t_j} := \int_{t_j}^{t_{j+1}} e^{-b(t_{j+1}-s)} dL_s \quad (11)$$

where $\Delta t_j := t_{j+1} - t_j$ is the length of the j -th time interval. At each step, $X_{t_{j+1}}$ is the sum of a mean-reversion term (known at time t_j) and an innovation $Z_{\Delta t_j}$, independent of the filtration $\mathcal{F}_{t_j}$. Notably, by time-homogeneity of L_t , the law of $Z_{\Delta t_j}$ depends only on Δt_j and not on the absolute time t_j .

The main challenge in the simulation is to effectively sample each r.v. $Z_{\Delta t}$ (considered hereinafter as the innovation over a generic time interval Δt), whose distribution is not available in closed form. In our approach, we proceed as follows. First, we compute the CF of $Z_{\Delta t}$, i.e.

$$\mathbb{E}[\exp\{iuZ_{\Delta t}\}] = \phi_Z(u, \Delta t), \quad (12)$$

which is readily available for any Lévy-driven OU process by using the relationships (5) and (8).

Then, we reconstruct numerically the CDF of $Z_{\Delta t}$ from the CF $\phi_Z(u, \Delta t)$, as outlined in Section 2. However, the basic inversion formula (3) cannot be applied to the CF of any general Lévy-driven OU, since the integrand function has a singularity at the origin. Following Lee [29, Theorem 5.1, p.12], we introduce a complex-plane shift $a \neq 0$ in the integration contour, and compute the CDF as:

$$P(x) = R_a - \frac{e^{-ax}}{2\pi} \int_{-\infty}^{+\infty} \frac{e^{-iux} \phi_Z(u - ia, \Delta t)}{i(u - ia)} du, \quad (13)$$

$$\text{where} \quad R_a = \begin{cases} 1 & \text{if } 0 < a < p_+ \\ 0 & \text{if } -p_- < a < 0 \end{cases} \quad (14)$$

and $(-p_+, p_-)$ denotes the analyticity strip of the CF $\phi_Z(u, \Delta t)$.

The incorporation of the complex shift $a \neq 0$ constitutes the key novelty of our approach, and distinguishes it from alternative methods based on the Hilbert transform (which corresponds to the case $a = 0$; see, e.g., [17,18] for applications to plain Lévy pro-

cesses). The shift removes the pole at the origin, a well-known source of numerical instability (as discussed for instance in [18,32]). Conversely, when computed inside the analyticity strip, the shifted integral (13) does not present any singularity.⁴

Moreover, when the CF $\phi_Z(\cdot, \Delta t)$ vanishes at infinity, the shifted integral can be truncated and efficiently evaluated via FFT-based integration. Since the integrand function is conjugate-symmetric along the integration path, the true CDF $P(x)$ is approximated with the Riemann sum

$$\hat{P}(x) := R_a - \frac{he^{-ax}}{\pi} \operatorname{Re} \left(\sum_{n=0}^{N-1} \frac{\phi_Z\left(\left(n + \frac{1}{2}\right)h - ia, \Delta t\right)}{i\left(n + \frac{1}{2}\right)h + a} e^{-ix\left(n + \frac{1}{2}\right)h} \right), \quad (15)$$

taken over N equispaced points with step size h (see, e.g., [29, p.14]). The sum is then computed via FFT, which retrieves the approximated CDF over an entire x -grid of N points.

For this methodology to apply to Lévy-driven OU processes, two technical conditions must be satisfied. First, the shift a must be selected within the analyticity strip of $\phi_Z(u, \Delta t)$. The new results in Propositions 1 and 2, which provide explicit characterisation of the analyticity strip for any Lévy-driven OU, are therefore essential for this step, as they identify the admissible range of shifts a for these processes.

Second, the CF $\phi_Z(u, \Delta t)$ must vanish as $|u| \rightarrow \infty$ to guarantee convergence of the truncated integral and the stability of the FFT approximation (15). A sufficient condition is that the distribution of $Z_{\Delta t}$ is absolutely continuous, i.e. it admits a Probability Density Function (as established by the Riemann-Lebesgue lemma; see, e.g., [33, p.10]). The following proposition addresses this requirement.

Proposition 3. *For any time interval $\Delta t > 0$, the distribution of $Z_{\Delta t}$ is either absolutely continuous, or it has exactly one probability atom (at zero) and is absolutely continuous elsewhere.*

Moreover, the latter case occurs if and only if $Z_{\Delta t}$ has finite-activity, i.e. if $\int_{\mathbb{R}} \nu_Z(x, \Delta t) dx < \infty$.

Proof. See Appendix A. $\square$

Hence, whenever $Z_{\Delta t}$ has infinite-activity (IA), the FFT-formulation (15) can be applied directly. On the contrary, when $Z_{\Delta t}$ has FA, the presence of a probability atom prevents the CF from decaying at infinity; for this case, we develop a modified simulation approach, which enables even the CF of these processes to satisfy the vanishing-at-infinity condition.

In the remainder of this section, we first provide all the details for implementing the simulation scheme in the IA case. Then, we present the novel scheme tailored for the FA case. Finally, we synthesise the two into the Fast and General Monte Carlo (FGMC) algorithm – a unified methodology for simulating any Lévy-driven OU process.

4.1. The infinite-activity case

Efficient simulation of Lévy-driven OU relies on fast and accurate evaluation of the inversion formula (13). For IA innovations, the simulation over each time interval $(r, t]$ of length Δt is performed as follows.

First, one computes the CF $\phi_Z(u, \Delta t)$ of the innovation $Z_{\Delta t}$ in (11). The key results are the relationships (5) and (8), which allow the determination of the CF $\phi_Z(u, t)$ for any Lévy-driven OU process. Then, the procedure consists of four steps:

1. Retrieve the CDF $\hat{P}$ on an x -grid, by computing (15) via FFT, after applying a suitable complex-plane shift a within the analyticity strip of ϕ_Z . Guidelines for selecting this shift are discussed in Section 5.3, in the context of error control.
2. Restrict the x -grid to the largest consecutive set of points $\{x_n\}$ for which the approximated CDF $\hat{P}$ is monotonically increasing and lies in $[0, 1]$. This step is crucial to ensure that $\hat{P}$ represents a valid CDF. The set begins with x_b and ends with x_e .
3. Reconstruct a continuous approximation of the CDF by performing an interpolation of the known values of $\hat{P}$ within the (restricted) range of the x -grid. The interpolation is performed using standard cubic splines with not-a-knot boundary conditions (see, e.g., [24]),⁵ which are well-suited to this setting since the CDF of $Z_{\Delta t}$ is continuous and, in most cases – including all IA Lévy-driven OU processes studied in the literature – at least C^1 -continuous.

Outside the range $[x_b, x_e]$, extrapolate the CDF using exponential tails,⁶ by setting

$$\hat{P}(x) = \begin{cases} \hat{P}(x_b) e^{\lambda_- (x - x_b)} & \text{for } x \leq x_b, \\ 1 - (1 - \hat{P}(x_e)) e^{-\lambda_+ (x - x_e)} & \text{for } x \geq x_e, \end{cases}$$

where the tail decay rates λ_- and λ_+ are computed as

$$\lambda_- := \frac{\ln \hat{P}(x_{b+1}) - \ln \hat{P}(x_b)}{x_{b+1} - x_b}, \quad \lambda_+ := \frac{\ln \hat{P}(x_e) - \ln \hat{P}(x_{e-1})}{x_e - x_{e-1}}.$$

⁴ Crucially, when computed within the analyticity strip, the CDF obtained via the shifted integral coincides with the original CDF, as the complex-plane shift modifies the integrand for numerical stability without affecting the result.

⁵ This is often the default implementation of cubic splines (e.g., MATLAB's `spline` and Python's `scipy.interpolate.CubicSpline` routines).

⁶ The exponential decay of the CDF follows from the existence of the analyticity strip $(-p_+, p_-)$. Extensive numerical experiments have shown that, in general, the proposed extrapolation approach is the most effective; alternative methodologies include, e.g., matching the theoretical asymptotic decay: $\hat{P}(x_b) e^{\rho_- (x - x_b)}$ for $x \leq x_b$ and $1 - (1 - \hat{P}(x_e)) e^{-\rho_+ (x - x_e)}$ for $x \geq x_e$.

4. Draw a uniform r.v. $U \sim \mathcal{U}(0, 1)$ and generate the innovation as $Z_{\Delta t} = \hat{P}^{-1}(U)$.

Finally, the new value of the process X_t is computed by summing the mean-reversion term $e^{-b\Delta t} X_r$ and the innovation term $Z_{\Delta t}$, as specified by the update rule (11).

The proposed scheme – outlined in Section 2 in its basic structure – is detailed in Algorithm 1 below.

Algorithm 1 Generation of $Z_{\Delta t}$ for Lévy-driven OU processes (infinite-activity).

Require: $\phi_Z(u, \Delta t)$ (CF of the innovation), Δt (Simulation step)

- 1: Retrieve the CDF $\hat{P}$ by FFT inversion of $\phi_Z(u, \Delta t)$, using (15) with a suitable complex-shift a
 - 2: Restrict the x -grid between x_b and x_e
 - 3: Interpolate the CDF using cubic splines and apply exponential extrapolation
 - 4: Draw a uniform r.v. $U \sim \mathcal{U}(0, 1)$ and compute $Z_{\Delta t}$ by CDF inversion
-

To obtain a complete trajectory over a time-grid $t_0, t_1, \dots, t_Q$, it is sufficient to iterate this procedure for every time interval. Furthermore, when the time-grid is equally spaced, steps 1 to 3 have to be performed only once: in such a case, the law of the innovation $Z_{\Delta t}$ is the same for every time interval, allowing for further improvement of the algorithmic efficiency.

4.2. The finite-activity case

The applicability of Algorithm 1 requires that the CF $\phi_Z(\cdot, \Delta t)$ vanishes at infinity. This condition is not satisfied by FA processes, since their distributions naturally contain a probability atom (cf. Proposition 3). In this subsection, we introduce a new algorithm for simulating Lévy-driven OU with FA.

By definition, an additive process Z_t has FA if its Lévy density $v_Z(x, t)$ integrates to a finite quantity

$$\Lambda(t) := \int_{-\infty}^{+\infty} v_Z(x, t) dx < \infty \quad (16)$$

(see, e.g., [25, Section 9]). In this case, for any fixed time t , the distribution of Z_t admits a representation in law as a compound Poisson r.v. with time-dependent parameters:

$$Z_t \stackrel{\text{law}}{\sim} \sum_{k=1}^{K_t} J_{t,k},$$

where K_t is a Poisson r.v. with intensity parameter $\Lambda(t)$, and $\{J_{t,k}\}_k$ are i.i.d. absolutely continuous jumps with density $v_Z(x, t)/\Lambda(t)$. Equivalently, the LCF takes the form

$$\Psi_Z(u, t) = \Lambda(t)(\phi_J(u, t) - 1), \quad \text{with} \quad \phi_J(u, t) := \frac{1}{\Lambda(t)} \int_{-\infty}^{+\infty} e^{iux} v_Z(x, t) dx. \quad (17)$$

Since K_t follows a Poisson distribution, there is always a positive probability that no jumps occur, namely $\mathbb{P}(K_t = 0) = e^{-\Lambda(t)}$. Therefore, for every $t > 0$ the distribution of Z_t has a probability atom at zero, which prevents the CF $\phi_Z(u, t)$ from vanishing at infinity (see, e.g., [34, Section 2.2]).

The key idea of our simulation procedure for FA Lévy-driven OU is to isolate this probability atom and focus on the simulation of the residual absolutely continuous part. By separating the zero-jumps case ($k = 0$) from the at-least-one-jump case ($k > 0$), we decompose the CF of Z_t as:

$$\phi_Z(u, t) = \mathbb{E}[e^{iuZ_t}] = e^{-\Lambda(t)} \sum_{k=0}^{\infty} \frac{\Lambda(t)^k}{k!} \phi_J(u, t)^k = e^{-\Lambda(t)} + \frac{e^{\Lambda(t)\phi_J(u, t)} - 1}{e^{\Lambda(t)} - 1} (1 - e^{-\Lambda(t)}).$$

Hence, for any fixed time t , the distribution of Z_t can be expressed as a mixture of two components:

$$Z_t \stackrel{\text{law}}{\sim} \begin{cases} Z_t^{\mathbb{C}} & \text{if } B_t = 1, \\ 0 & \text{if } B_t = 0. \end{cases} \quad (18)$$

where B_t is a Bernoulli r.v. with success probability $1 - e^{-\Lambda(t)}$, indicating whether at least one jump occurs, and $Z_t^{\mathbb{C}}$ represents the total sum of the jumps $\{J_{t,k}\}_k$, conditional on at least one jump occurring. In particular, the CF of $Z_t^{\mathbb{C}}$ – which we refer to as the *conditional CF* – reads

$$\phi_Z^{\mathbb{C}}(u, t) := \mathbb{E}[e^{iuZ_t} | B_t = 1] = \frac{e^{\Lambda(t)\phi_J(u, t)} - 1}{e^{\Lambda(t)} - 1}. \quad (19)$$

Crucially, $\phi_Z^{\mathbb{C}}(u, t)$ vanishes at infinity,⁷ ensuring the applicability of the FFT-based inversion.

⁷ This follows because $\phi_J(u, t) \rightarrow 0$ as $|u| \rightarrow \infty$ by the Riemann-Lebesgue lemma, since $v_Z(x, t)/\Lambda(t)$ is a valid PDF.

Thus, for any Lévy-driven OU process with FA, the innovation $Z_{\Delta t}$ is simulated in two steps. First, we sample the Bernoulli r.v. $B_{\Delta t}$, determining whether at least one jump occurs. Then, if $B_{\Delta t} = 1$, we simulate the total jump magnitude $Z_{\Delta t}^c$ via [Algorithm 1](#). The procedure is summarised in [Algorithm 2](#).

Algorithm 2 Generation of $Z_{\Delta t}$ for Lévy-driven OU processes (finite-activity).

Require: $\phi_Z(u, \Delta t)$ (CF of the innovation), Δt (Simulation step)

- 1: Draw a Bernoulli r.v. $B_{\Delta t}$ with success probability $1 - e^{-\Lambda(\Delta t)}$
 - 2: **if** $B_{\Delta t} = 1$ **then**
 - 3: Compute the conditional CF $\phi_Z^c(u, \Delta t)$ as in (19)
 - 4: Simulate $Z_{\Delta t}$ from the conditional CF $\phi_Z^c(u, \Delta t)$, using [Algorithm 1](#)
 - 5: **else**
 - 6: No jumps: $Z_{\Delta t} = 0$
-

It is worth emphasising that in most cases this method can be substantially faster than the standard approach used for simulating compound Poisson r.v.s – commonly appearing in ED schemes for FA processes – which involves first simulating the number of jumps K_t , and then simulating the size of each of the K_t jumps individually. The computational cost of this second operation can be relevant when many jumps occur. In contrast, the proposed algorithm requires sampling at most two uniform r.v.s, since all jumps are simulated at once by numerically reconstructing their “aggregate” CDF.

4.3. Fast and general MC for Lévy-driven OU processes

The FGMC algorithm is a general methodology that unifies the IA and FA schemes, allowing the simulation of any Lévy-driven OU on a discrete time-grid. It is naturally tailored for pricing discretely-monitored derivatives, as it enables the simulation of the process directly at the relevant monitoring dates. The algorithm is schematically represented in [Fig. 1](#).

The FGMC requires only three parameters of the numerical scheme (hereinafter, *numerical parameters*) associated with the FFT integration in (15), which are: the number of points N in the x -grid (conventionally chosen equal to $N = 2^M$, with $M \in \mathbb{N}$), the integration step h used in the Fourier space, and the complex shift a .⁸ As we show in the next section, these parameters directly control the numerical error of the method.

While users can select the values of these three numerical parameters based on their specific accuracy requirements, [Section 5.3](#) provides guidelines to support this choice. In particular, we establish selection rules for h and a , leaving $M := \log_2 N$ as the only free parameter that controls the resolution of the CDF reconstruction.




Fig. 1. Flowchart of the FGMC algorithm. This FFT-based algorithm can simulate any Lévy-driven OU process, both with finite- and infinite-activity.

⁸ As standard in the FFT literature (see, e.g., [35]), we consider a symmetric x -grid, with $x_1 = -x_N$.

Table 2

The three numerical parameters for the proposed FGMC algorithm.

Numerical parameter	Description
M	Integer number equal to $\log_2 N$, with N number of FFT points
h	Step size in the Fourier domain (integration step)
a	Imaginary shift of the integration path in the complex-plane

5. Error control

In this section, we provide general bounds for the numerical error introduced by the FGMC algorithm. We then derive specific error bounds for the main Lévy-driven OU processes, illustrating how our method translates into precise error control. Finally, we present practical guidelines for parameter selection.

5.1. Control of the numerical error

A key theoretical property of the FGMC algorithm is that its numerical error can be explicitly controlled for any CF that vanishes at infinity. This condition holds for all IA Lévy-driven OU, while for FA processes it can be achieved through a suitable transformation, as we have discussed in [Section 4.2](#).

For every fixed $t > 0$, we distinguish two main cases, depending on the decay of the CF modulus:

- (i) the decay is *exponential*, i.e. there exist constants $B, \ell, \omega > 0$ such that, for sufficiently large $|u|$, the following bound holds

$$|\phi_Z(u - ia, t)| < B e^{-\ell|u|^\omega} \quad \forall a \in (-p_-, p_+); \quad (20)$$

- (ii) the decay is *power-law*, i.e. there exist constants $B, \omega > 0$ such that, for sufficiently large $|u|$, the following bound holds

$$|\phi_Z(u - ia, t)| < B|u|^{-\omega} \quad \forall a \in (-p_-, p_+). \quad (21)$$

As we demonstrate in the following proposition, the numerical error of our algorithm is fully determined by the decay of the CF within the analyticity strip.

Proposition 4. For every shift $a \in (-p_-/2, p_+/2) \setminus \{0\}$ and for every $x \in \mathbb{R}$, the numerical error $\mathcal{E}(x) := |P(x) - \hat{P}(x)|$ of the CDF approximation satisfies the following bounds:

- (i) if the modulus of the CF has exponential decay,

$$\mathcal{E}(x) \leq \frac{B e^{-ax}}{\pi \omega} \frac{e^{-\ell(Nh)^\omega}}{\ell(Nh)^\omega} + \frac{1 + e^{-2ax} \phi_Z(-2ia, t)}{e^{2\pi|a|/h} - e^{-2\pi|a|/h}}; \quad (22)$$

- (ii) if the modulus of the CF has power-law decay,

$$\mathcal{E}(x) \leq \frac{B e^{-ax}}{\pi \omega} (Nh)^{-\omega} + \frac{1 + e^{-2ax} \phi_Z(-2ia, t)}{e^{2\pi|a|/h} - e^{-2\pi|a|/h}}. \quad (23)$$

Proof. See [Appendix A](#). $\square$

The approximation error resulting from FFT integration is controlled by two distinct terms: the first (the truncation error) depends on the tail behaviour of the CF; the second (the discretisation error) is related to the discretisation of the integral and is the same in both decay cases.

5.2. An example of error control: the OU-NTS process

To illustrate how the simulation error can be explicitly controlled and related to model parameters, we consider the OU-NTS process (see, e.g., [\[14\]](#)) as a representative example. Thanks to [Proposition 4](#), it suffices to analyse the decay of $|\phi_Z(u, t)| = \exp\{\operatorname{Re}(\Psi_Z(u, t))\}$.

The OU-NTS process, whose LCF is reported in [Appendix B](#), exhibits different jump regimes depending on the sign of the parameter α . For $\alpha > 0$, the process has IA (and its CF vanishes at infinity), allowing direct simulation via [Algorithm 1](#). For $\alpha < 0$, the process has FA, and the simulation requires [Algorithm 2](#). In this latter case, its LCF can be expressed as that of a compound Poisson r.v. [\(17\)](#):

$$\Psi_Z(u, t) = \underbrace{\frac{(1-\alpha)t}{\kappa|\alpha|}}_{\Lambda(t)} \left[\underbrace{\left(\frac{1-\alpha}{\kappa} \right)^{-\alpha} \frac{1}{bt} \int_{e^{-bt}}^1 \left(\frac{1}{2} \sigma^2 u^2 z^2 - i\theta uz + \frac{1-\alpha}{\kappa} \right) \frac{dz}{z} - 1}_{\phi_J(u, t)} \right].$$

This representation makes it straightforward to construct the conditional CF $\phi_Z^c(u, t)$ in [\(19\)](#).

To precisely determine the numerical error in both regimes, we establish asymptotic bounds on the CF decay. The following lemma demonstrates that for the OU-NTS the decay is exponential in the IA case, and power-law in the FA case.

Lemma 5. For an OU-NTS process with $\alpha \in (0, 1)$, the CF ϕ_Z has exponential decay as $|u| \rightarrow \infty$:

$$\log |\phi_Z(u - ia, t)| = -\left(\frac{\sigma^2}{2}\right)^\alpha \left(\frac{1-\alpha}{\kappa}\right)^{1-\alpha} \frac{1 - e^{-2\alpha b t}}{2\alpha^2 b} |u|^{2\alpha} + o(|u|^{2\alpha}) \quad \forall a \in (-p_-, p_+).$$

For an OU-NTS process with $\alpha < 0$, the conditional CF ϕ_Z^c has power-law decay as $|u| \rightarrow \infty$:

$$|\phi_Z^c(u - ia, t)| = \left(\frac{\sigma^2}{2}\right)^\alpha \left(\frac{1-\alpha}{\kappa}\right)^{1-\alpha} \frac{e^{-2\alpha b t} - 1}{2\alpha^2 b} |u|^{2\alpha} + o(|u|^{2\alpha}) \quad \forall a \in (-p_-, p_+).$$

In both cases, $(-p_+, p_-)$ denotes the NTS analyticity strip (see Lemma D.1).

Proof. See Appendix A. $\square$

Hence, for OU-NTS processes, both IA and FA regimes deliver effective error control: IA processes achieve a rapid exponential decrease of the numerical error, enabling high-precision simulations, while FA processes – when handled using the procedure in Section 4.2 – maintain a robust algebraic decrease suitable for all practical applications.

Furthermore, similar decay patterns hold for the other main Lévy-driven OU processes considered in the literature (namely OU-TS, TS-OU, NTS-OU), with detailed results provided in Appendices C and D. Collectively, the behaviour of the numerical error is summarised in the following proposition.

Proposition 6. For all Lévy-driven OU processes of TS and NTS type, the numerical error $\mathcal{E}(x)$ produced by the FGMC algorithm exhibits the following behaviour as a function of the FFT grid size:

- (i) for all IA specifications, $\mathcal{E}(x)$ decreases exponentially;
- (ii) for all FA specifications, $\mathcal{E}(x)$ decreases at a power-law rate.

Proof. See Appendix A. $\square$

5.3. A rule of thumb for numerical parameter selection

In Table 2, we have listed the three numerical parameters required by the FGMC algorithm; in what follows, we provide practical criteria for their selection based on the error analysis presented above. Regarding the grid size N , as we show in Section 6, even a relatively small grid ($M = 16$) achieves excellent accuracy. This indicates that the method can deliver high-precision results without requiring an excessively refined FFT grid.

Proposition 4 provides useful guidance for selecting the shift a and the integration step h for a given grid size N , thereby determining all parameters in the numerical scheme, as well as the convergence rate of the error bound. Specifically, the selection of h is informed by the decay of the CF:

- (i) In the exponential decay case, for fixed N , the integration step h can be chosen as

$$h(N) = \max \left(\left(\frac{2\pi|a|}{\ell N^\omega} \right)^{1/(\omega+1)}, 1\% \right).$$

The first term ensures that the exponential decay of the truncation error, i.e. $e^{-\ell(Nh)^\omega}$, is matched with that of the discretisation error, i.e. $e^{-2\pi|a|/h}$. With this choice of h , the overall error bound in (22) decreases exponentially with the FFT grid size N .⁹

- (ii) In the power-law decay case, the truncation error – the first term in (23) – dominates, because it decays more slowly than the discretisation error. Based on our numerical experiments, a choice of $h = 1\%$ proves sufficient to ensure accurate results across all tested cases. With this selection of the parameters, the error bound decreases with N as $O(N^{-\omega})$.

To minimise the total error on either tail of the CDF, the shift parameter a should be chosen as large as possible within the range permitted by Proposition 4, i.e. near either $-p_-/2$ or $p_+/2$. This places the integration path approximately midway between the singularity at zero and the corresponding boundary of the analyticity strip.

Finally, we suggest a practical rule of thumb to validate the implementation: verify that the condition $\max(\hat{P}(x_b), 1 - \hat{P}(x_e)) \leq 10^{-4}$ is satisfied. This ensures the computational domain (x_b, x_e) is sufficiently wide, as it confirms that negligible probability lies outside the boundaries. If this condition is not met, accuracy can be improved by performing two separate FFT reconstructions with different shifts a – one positive and one negative – for the respective positive and negative halves of the CDF grid. This trick can enhance the CDF reconstruction while keeping the total MC computational time substantially unchanged.

6. Applications and numerical results

In this section, we analyse the numerical performances of the proposed simulation algorithm and we present two applications in energy derivative pricing.¹⁰

⁹ The 1% lower bound is a conservative safeguard; as we have observed in extensive numerical experiments, even smaller values of h consistently yield excellent accuracy. Further details are available upon request.

¹⁰ All results are obtained using MATLAB R2025a on a machine with Apple M1 Chip, 8 GB RAM, macOS Sequoia 15.6.1.

First, we test the accuracy and speed of our simulation method, comparing the results with those obtained using the ED algorithms listed in Table 1. The simulation accuracy is measured, as standard in the literature, by considering the first four cumulants: we compare the MC estimates (obtained with 10^7 simulations) with the true cumulants – which can be computed analytically for all considered models (see Section 3). The speed is measured as the time in seconds required for the simulation.

Then, we illustrate how the proposed technique can be effectively used for pricing energy derivatives, considering two common instruments in the energy markets: European and Asian options. When dealing with options, it is not only important to assess numerical performance in pricing at-the-money (ATM) instruments, but also essential to verify that models and algorithms can correctly reproduce prices across different levels of moneyness (see, e.g., [36]). Indeed, the inherent seasonality of energy markets introduces fluctuations in the moneyness of the options over months and years (see, e.g., [14]). While the existing literature on simulation techniques for these processes focuses on ATM options (see, e.g., [13,37]), we perform an extended analysis by considering options in a range of moneyness around the ATM, covering the spectrum of strikes relevant for typical market volatility.

6.1. Accuracy and computational time

We consider for the comparison all processes listed in Table 1. When more than one ED scheme exists, we implement the fastest (and most recently published) algorithm for each class: *Algorithm 1* and *Algorithm 2* in [16] for finite-variation OU-TS and TS-OU, *Algorithm 1* in [12] for finite-activity OU-TS, *Algorithm 3.1* in [14] for finite- and infinite-variation OU-NTS, and *Algorithm 1* in [13] for finite- and infinite-variation NTS-OU.¹¹

All processes are simulated on a yearly time-horizon ($t = 1$), with the initial condition $X_0 = 0$. For OU-TS and TS-OU processes, as in [16], we select the parameters $(b, \beta_p, \beta_n, c_p, c_n, \gamma_c) = (0.1, 2.5, 3.5, 0.5, 1, 0)$ and let α_p and α_n vary over the range of admissible values, setting $\alpha := \alpha_p = \alpha_n$. For OU-NTS and NTS-OU processes, following [41], we consider the parameters $(b, \sigma, \kappa, \theta) = (0.2162, 0.201, 0.256, 0)$ and let α vary.¹² Finally, for the FFT numerical parameters, we choose $M = 16$, while h and a are selected so as to minimise the numerical error, as suggested in Section 5.3.

Table 3 shows the accuracy results, obtained by simulating 10^7 realisations for each process and for different choices of α . The true cumulants are deduced using the relationships (7) and (10), together with the formulas in Lemmas C.2 and D.2. In all cases, our FGMC simulation algorithm appears to be extremely accurate, reproducing the true cumulants up to the fourth decimal digit.

Table 4 highlights the other important result: the improvement in terms of computational time is extremely relevant. Across all processes, our simulation technique consistently outperforms the ED schemes, being at least one order of magnitude faster, and the computational times remain relatively constant across the values of the parameter α . This is not true for the ED methods, where the results are highly sensitive to parameter choice, as a consequence of the acceptance/rejection algorithms.

For reference, the time required to simulate 10^7 samples of the Gaussian OU – the simplest OU process – is 0.17 seconds. Hence, FGMC achieves results that are comparable in magnitude to this benchmark model, which is simulated by sampling only a Gaussian r.v.

Finally, the FGMC allows the extension of the MC simulation also to the asymmetric Lévy-driven OU processes of NTS type. The existing literature on ED simulation has focused exclusively on the symmetric case (cf. [13]) with the only exception of $\alpha = 0$ (cf. [11]). In Table 5, we report the results for the simulation of asymmetric OU-NTS and NTS-OU. We consider a yearly horizon and model parameters equal to $(b, \sigma, \kappa, \theta) = (0.2162, 0.201, 0.256, 0.1)$, maintaining consistency with those used previously, with the only difference being the parameter θ . Also in this case, the empirical estimates of the cumulants demonstrate exceptional accuracy, with precision extended up to the fourth decimal digit consistently across all levels of the stability parameter α . The availability of a simulation scheme for such processes opens up a broader variety of modelling possibilities, which have not yet been explored in the current state of play.

6.2. Energy derivative pricing

In this subsection, we discuss the application of FGMC in energy derivative pricing. We assume that the market model is specified in the risk-neutral measure and that the spot price is driven by the following one-factor process:

$$S_t = F_{0,t} e^{X_t + f_t}, \quad (24)$$

where X_t is a Lévy-driven OU starting in $X_0 = 0$, f_t is a deterministic function of time, and $F_{0,t}$ is the forward price with maturity t derived from quoted products at value date $t_0 = 0$. As carried out by Hambly et al. [2], the risk-neutral conditions are met by imposing that $f_t = -\Psi_Z(-i, t)$ for all $t \geq 0$.

We consider in the following the pricing of two main derivatives, a European call option and an Asian path-dependent option, demonstrating that in both cases FGMC is an extremely well-suited methodology for pricing in terms of both accuracy and computational time. Although our analysis focuses on one-factor models, the simulation methodology extends straightforwardly to the

¹¹ To sample the required TS r.v.s, besides the standard acceptance/rejection algorithm in [38], we implemented the fast algorithms proposed in [39] and [40], following the joint implementations suggested in the latter paper. These algorithms for simulating TS r.v.s can be up to 50 times faster than the one in [38] for the parameters considered in the numerical analysis, but in any case they are at least one order of magnitude slower than the proposed FGMC algorithm. Details are available upon request.

¹² Similar conclusions have been derived with different parameter sets and are available upon request.

Table 3

Comparison of the first four cumulants (multiplied by 1000) obtained via the analytical expression (True), via the proposed method (FGMC) and via the Exact Decomposition (ED). We consider the TS and the NTS cases for both OU-Lévy and Lévy-OU processes, generating 10^7 trials with a time horizon of $t = 1$, and selecting different values for the stability parameter α . The symbol $\times$ indicates that no ED algorithm is available for that class of processes. We notice that FGMC is extremely accurate in all cases; the obtained accuracy is at least at the fourth decimal digit for all cumulants.

OU-TS												
α	$c_1(X_t) \cdot 10^3$			$c_2(X_t) \cdot 10^3$			$c_3(X_t) \cdot 10^3$			$c_4(X_t) \cdot 10^3$		
	True	FGMC	ED	True	FGMC	ED	True	FGMC	ED	True	FGMC	ED
1.6	0.000	-0.627	$\times$	1914.8	1914.3	$\times$	-26.426	-27.566	$\times$	107.41	107.58	$\times$
1.2	0.000	-0.365	$\times$	640.81	640.65	$\times$	-7.070	-7.309	$\times$	94.514	94.547	$\times$
0.8	0.000	-0.264	-0.027	323.64	323.54	323.51	2.916	2.799	2.660	89.495	89.532	89.836
0.4	0.000	-0.213	0.033	202.58	202.51	202.70	9.473	9.381	9.202	90.270	90.327	90.784
-1.0	0.000	-0.204	-0.028	100.28	100.32	100.18	31.807	31.743	31.642	138.94	139.50	138.45
-2.0	0.000	-0.061	-0.100	105.85	105.76	105.58	66.683	66.533	66.179	256.36	256.98	254.28
TS-OU												
α	$c_1(X_t) \cdot 10^3$			$c_2(X_t) \cdot 10^3$			$c_3(X_t) \cdot 10^3$			$c_4(X_t) \cdot 10^3$		
	True	FGMC	ED	True	FGMC	ED	True	FGMC	ED	True	FGMC	ED
1.6	0.000	-0.282	$\times$	382.96	382.85	$\times$	-7.927	-8.016	$\times$	42.962	43.089	$\times$
1.2	0.000	-0.168	$\times$	128.16	128.12	$\times$	-2.121	-2.141	$\times$	37.805	37.945	$\times$
0.8	0.000	-0.120	0.005	64.727	64.695	64.644	0.874	0.865	0.734	35.798	35.971	35.965
0.4	0.000	-0.091	-0.021	40.517	40.480	40.469	2.841	2.798	2.988	36.108	36.060	36.448
OU-NTS												
α	$c_1(X_t) \cdot 10^3$			$c_2(X_t) \cdot 10^3$			$c_3(X_t) \cdot 10^3$			$c_4(X_t) \cdot 10^3$		
	True	FGMC	ED	True	FGMC	ED	True	FGMC	ED	True	FGMC	ED
0.8	0.000	-0.083	-0.065	32.800	32.790	32.799	0.000	-0.000	-0.007	0.839	0.839	0.838
0.6	0.000	-0.084	-0.086	32.800	32.790	32.809	0.000	-0.003	-0.007	0.839	0.838	0.842
0.4	0.000	-0.084	-0.072	32.800	32.791	32.809	0.000	-0.003	-0.008	0.839	0.837	0.836
0.2	0.000	-0.084	-0.075	32.800	32.791	32.823	0.000	-0.003	0.000	0.839	0.837	0.844
-1.0	0.000	-0.120	$\times$	32.800	32.790	$\times$	0.000	-0.007	$\times$	0.839	0.841	$\times$
-2.0	0.000	-0.121	$\times$	32.800	32.790	$\times$	0.000	-0.006	$\times$	0.839	0.842	$\times$
NTS-OU												
α	$c_1(X_t) \cdot 10^3$			$c_2(X_t) \cdot 10^3$			$c_3(X_t) \cdot 10^3$			$c_4(X_t) \cdot 10^3$		
	True	FGMC	ED	True	FGMC	ED	True	FGMC	ED	True	FGMC	ED
0.8	0.000	-0.057	-0.049	14.183	14.176	14.175	0.000	0.001	-0.001	0.725	0.726	0.719
0.6	0.000	-0.057	-0.059	14.183	14.177	14.171	0.000	-0.001	-0.007	0.725	0.724	0.715
0.4	0.000	-0.057	-0.058	14.183	14.178	14.196	0.000	-0.001	-0.006	0.725	0.724	0.729
0.2	0.000	-0.057	-0.039	14.183	14.178	14.182	0.000	-0.002	-0.002	0.725	0.724	0.726

Table 4

Computational times in seconds for simulating 10^7 trials with our simulation technique (FGMC) and Exact Decomposition (ED). We simulate the processes on the yearly time horizon ($t = 1$), maintaining consistent parameter values and ranges as in Table 3. The symbol $\times$ denotes that, for that choice of α , no ED algorithm is available in the literature, while the writing *Not Defined* indicates that the process is not defined. The comparison reveals a significant disparity between the two approaches, with FGMC outpacing exact decomposition by at least one order of magnitude.

α	OU-TS		TS-OU		α	OU-NTS		NTS-OU	
	FGMC	ED	FGMC	ED		FGMC	ED	FGMC	ED
1.6	0.87	$\times$	0.25	$\times$	0.8	0.74	5.44	0.39	3.46
1.2	0.88	$\times$	0.27	$\times$	0.6	0.72	13.75	0.39	4.24
0.8	0.89	26.88	0.29	10.11	0.4	0.73	15.15	0.35	5.82
0.4	0.89	16.84	0.31	5.02	0.2	0.75	15.85	0.33	7.09
-1.0	0.59	4.81	<i>Not Defined</i>		-1.0	0.69	$\times$	<i>Not Defined</i>	
-2.0	0.58	3.01	<i>Not Defined</i>		-2.0	0.70	$\times$	<i>Not Defined</i>	

Table 5

True and empirical cumulants (multiplied by 1000) obtained by generating 10^7 samples with our technique (FGMC), together with the corresponding simulation times. We simulate the processes considering a yearly horizon, i.e. $t = 1$, using the same parameters as in Table 3 – with the exception of the parameter θ , which is chosen equal to 0.1, to ensure the asymmetry of the Lévy driver. The symbol $\times$ indicates that no ED is available for that choice of process and parameters. The results show that the FGMC obtains excellent results also in the asymmetric case, highlighting the great level of generality, speed and accuracy of the methodology.

Asymmetric OU-NTS										
α	$c_1(X_t) \cdot 10^3$		$c_2(X_t) \cdot 10^3$		$c_3(X_t) \cdot 10^3$		$c_4(X_t) \cdot 10^3$		Time [s]	
	True	FGMC	True	FGMC	True	FGMC	True	FGMC	FGMC	ED
0.8	89.929	89.840	34.879	34.862	2.572	2.568	1.551	1.570	0.68	$\times$
0.6	89.929	89.840	34.879	34.863	2.451	2.446	1.234	1.237	0.65	$\times$
0.4	89.929	89.840	34.879	34.863	2.411	1.135	1.135	1.135	0.72	$\times$
0.2	89.929	89.840	34.879	34.863	2.391	2.385	1.087	1.086	0.71	$\times$
-1.0	89.929	89.801	34.879	34.859	2.355	2.345	1.002	1.001	0.58	$\times$
-2.0	89.929	89.800	34.879	34.861	2.347	2.340	0.983	0.986	0.56	$\times$

Asymmetric NTS-OU										
α	$c_1(X_t) \cdot 10^3$		$c_2(X_t) \cdot 10^3$		$c_3(X_t) \cdot 10^3$		$c_4(X_t) \cdot 10^3$		Time [s]	
	True	FGMC	True	FGMC	True	FGMC	True	FGMC	FGMC	ED
0.8	19.443	19.382	15.081	15.070	1.668	1.668	1.341	1.369	0.27	$\times$
0.6	19.443	19.382	15.081	15.071	1.590	1.586	1.067	1.072	0.28	$\times$
0.4	19.443	19.381	15.081	15.071	1.564	1.560	0.982	0.983	0.29	$\times$
0.2	19.443	19.381	15.081	15.072	1.551	1.547	0.940	0.940	0.32	$\times$

multi-factor case, as well as to the modelling of the forward price (see, e.g., [42,43]), as already discussed in the existing literature (see, e.g., [14]).

6.2.1. European option

In the European option case, we evaluate calls with different moneyness levels for the four processes described in the previous sections. To ensure our simulation technique correctly prices both in-the-money (ITM) and out-of-the-money (OTM) options, we select 30 call options (with one-year maturity, i.e. $T = 1$), having moneyness χ in the range $\sqrt{T}(-0.2, 0.2)$. We recall that the moneyness of an option having strike K and maturity T is defined as

$$\chi := \log \frac{K}{F_{0,T}}.$$

Moreover, as in [13], we consider a flat forward curve, assuming $F_{0,t} = 1 \forall t \geq 0$.

To ensure consistency with the previous experiments, we opt for the same set of model parameters used in Section 6.1. In particular, we consider symmetric OU-NTS and NTS-OU processes (i.e. $\theta = 0$) to allow a direct comparison with ED algorithms. For the FGMC, the selected numerical parameters are $M = 16$, and both h and a as suggested in Section 5.3. The number of simulations is set to 10^7 .

For European call options an exact formula exists (cf. [44]), allowing us to verify whether the retrieved MC prices align with the true prices of the contracts. To properly assess the accuracy of FGMC, we compute different metrics – the Maximum absolute error (MAX), the Root Mean Squared Error (RMSE), the Mean Absolute Percentage Error (MAPE) and the average of the MC Standard Deviation ($\overline{SD}$) – over the set of 30 moneyness levels.

Table 6 shows the comparison in terms of accuracy and computational time. The results are noteworthy: we observe that FGMC achieves an accuracy on the order of one basis point (bp), not only for a few specific moneyness levels, but for all considered levels. Moreover, this high accuracy is consistent across all four processes and for all choices of the parameter α .

This analysis highlights once again the two main strengths of our methodology. First of all, it is fast: the computational time is always remarkably faster than the corresponding ED algorithm. Second, it is general, and we are indeed able to compute the prices for all the admissible choices of the parameters, whereas – in many cases – the ED algorithms do not allow the computation of the price.

6.2.2. Asian option

As a second application in derivative pricing, we consider an Asian call option, as in [12]. In gas and power markets, these contracts typically have expiry dates spanning several years, and play an important role as effective hedging devices (see, e.g., [3]). The payoff of the derivative is

$$\max \left(\frac{1}{Q} \sum_{j=1}^Q S_{t_j} - K, 0 \right) \quad (25)$$

Table 6

Comparison between European call pricing error for FGMC and Exact Decomposition (ED) methods. After computing the true prices with Lewis formula, we compute MAX, RMSE, MAPE and SD for 30 call options (one-year maturity), with moneyness in the range $\sqrt{T}(-0.2, 0.2)$; in addition, we report the computational time in seconds. All results are obtained using 10^7 trials, and the symbol $\times$ indicates that no ED algorithm is available for that class of processes. We notice that in all cases the error in pricing is of the order of one bp and that FGMC is always remarkably faster than the ED.

OU-TS										
α	MAX [bp]		RMSE [bp]		MAPE [%]		$\overline{SD}$ [bp]		Time [s]	
	FGMC	ED	FGMC	ED	FGMC	ED	FGMC	ED	FGMC	ED
1.6	2.21	$\times$	1.65	$\times$	0.03	$\times$	7.47	$\times$	0.99	$\times$
1.2	2.72	$\times$	2.09	$\times$	0.06	$\times$	2.76	$\times$	0.98	$\times$
0.8	2.16	1.69	1.52	1.34	0.06	0.05	1.88	1.85	0.97	26.28
0.4	1.64	1.45	1.08	1.35	0.05	0.08	1.63	1.59	0.97	17.35
-1.0	1.47	1.96	1.32	1.59	0.12	0.15	2.01	2.20	0.71	3.05
-2.0	1.65	3.99	1.48	3.87	0.12	0.32	3.31	3.35	0.61	2.13

TS-OU										
α	MAX [bp]		RMSE [bp]		MAPE [%]		$\overline{SD}$ [bp]		Time [s]	
	FGMC	ED	FGMC	ED	FGMC	ED	FGMC	ED	FGMC	ED
1.6	2.53	$\times$	1.87	$\times$	0.07	$\times$	1.81	$\times$	0.49	$\times$
1.2	1.46	$\times$	0.97	$\times$	0.05	$\times$	1.06	$\times$	0.47	$\times$
0.8	0.75	0.36	0.45	0.23	0.03	0.02	0.92	0.90	0.42	13.39
0.4	1.53	1.53	0.35	1.48	0.04	0.29	0.90	0.91	0.45	5.51

OU-NTS										
α	MAX [bp]		RMSE [bp]		MAPE [%]		$\overline{SD}$ [bp]		Time [s]	
	FGMC	ED	FGMC	ED	FGMC	ED	FGMC	ED	FGMC	ED
0.8	1.01	1.17	0.66	0.77	0.06	0.08	0.37	0.37	0.65	6.35
0.6	1.01	1.02	0.66	0.66	0.07	0.06	0.37	0.37	0.68	13.39
0.4	1.00	0.89	0.66	0.55	0.07	0.05	0.37	0.37	0.74	15.28
0.2	1.00	1.41	0.66	1.03	0.08	0.14	0.37	0.37	0.72	16.34
-1.0	0.49	$\times$	0.42	$\times$	0.08	$\times$	0.37	$\times$	0.60	$\times$
-2.0	0.53	$\times$	0.46	$\times$	0.08	$\times$	0.37	$\times$	0.63	$\times$

NTS-OU										
α	MAX [bp]		RMSE [bp]		MAPE [%]		$\overline{SD}$ [bp]		Time [s]	
	FGMC	ED	FGMC	ED	FGMC	ED	FGMC	ED	FGMC	ED
0.8	0.71	0.58	0.42	0.39	0.04	0.07	0.05	0.25	0.34	3.73
0.6	0.69	0.75	0.41	0.46	0.05	0.05	0.25	0.25	0.37	4.58
0.4	0.67	0.71	0.40	0.46	0.05	0.06	0.25	0.25	0.41	5.75
0.2	0.65	0.84	0.38	0.52	0.05	0.06	0.25	0.25	0.39	7.23

where S_t is the spot price, specified as in (24), K the strike price, $\{t_j\}_{j=1,\dots,Q}$ the Q monitoring dates, and $T = t_Q$ the expiry of the option. We consider in our analysis an Asian option having two-year expiry and monthly-monitoring, i.e. $Q = 24$.

Unlike the European call case, for these derivatives no exact pricing formula exists, and MC simulation is required. To evaluate such a contract, one needs to simulate discrete trajectories of the process X_{t_j} , with $0 = t_0 < t_1 < \dots < t_{24} = T$, where $T = 2$ represents the contract expiry.

We conduct a comprehensive comparison of pricing accuracy using FGMC and ED methods. To evaluate performance across different moneyness levels, we focus on three key contracts: an OTM option ($\chi = 0.2\sqrt{T}$), an ATM option ($\chi = 0$), and an ITM option ($\chi = -0.2\sqrt{T}$). We set a flat forward curve $F_{0,t} = 1 \forall t \geq 0$, and we select model parameters and FGMC numerical parameters as in the previous experiments.

Table 7 presents the results of our analysis. For both simulation methods, we provide the prices of the three contracts along with their corresponding MC standard deviations, and the simulation time. We observe that, for all considered cases, the obtained option prices align between the two simulation approaches. Furthermore, the SD consistently remains around 1 bp, the typical accuracy required in most financial markets, confirming the suitability of the technique for real-world applications.

In terms of computational time, the results are remarkable and mirror those observed previously. As expected, when simulating trajectories with multiple monitoring dates, FGMC showcases a substantial gain in computational speed compared to ED methods, while maintaining analogous pricing accuracy.

Table 7

Pricing Asian options using FGMC and Exact Decomposition (ED) methods. We evaluate an arithmetic average Asian option with two-year ($T = 2$) expiry and monthly monitoring. We consider 3 relevant levels of moneyness (OTM, ATM, ITM), corresponding to $\chi = 0.2\sqrt{T}$, $\chi = 0$, $\chi = -0.2\sqrt{T}$ respectively, and we compute the price and the MC standard deviation for different values of the stability parameters α . As in the previous experiments, estimates are obtained simulating 10^7 paths, using the same parameters as in Table 3, along with the simulation times in seconds for both methods. The symbol $\times$ denotes that, for that choice of α , no ED is available in the literature. In all cases prices are consistent across the two methods, but FGMC is always remarkably faster than the ED.

OU-TS														
α	Price FGMC [%]			SD FGMC [bp]			Price ED [%]			SD ED [bp]			Time [s]	
	OTM	ATM	ITM	OTM	ATM	ITM	OTM	ATM	ITM	OTM	ATM	ITM	FGMC	ED
1.6	168.5	184.5	199.0	22.6	22.8	22.8	$\times$	$\times$	$\times$	$\times$	$\times$	$\times$	5.88	$\times$
1.2	36.33	50.42	65.47	3.18	3.42	3.58	$\times$	$\times$	$\times$	$\times$	$\times$	$\times$	6.15	$\times$
0.8	15.99	28.44	44.24	1.62	1.87	2.03	15.96	28.40	44.19	1.63	1.87	2.03	6.69	98.58
0.4	9.28	20.00	36.70	1.23	1.45	1.59	9.29	20.00	36.70	1.25	1.47	1.60	7.12	68.26
-1.0	6.11	10.84	31.16	1.50	1.64	1.69	6.14	10.87	31.19	1.67	1.80	1.84	2.89	13.74
-2.0	8.72	12.02	32.66	3.12	3.21	3.23	8.78	12.08	32.73	2.90	3.00	3.03	2.53	12.36
TS-OU														
α	Price FGMC [%]			SD FGMC [bp]			Price ED [%]			SD ED [bp]			Time [s]	
	OTM	ATM	ITM	OTM	ATM	ITM	OTM	ATM	ITM	OTM	ATM	ITM	FGMC	ED
1.6	14.97	26.17	40.18	1.45	1.72	1.91	$\times$	$\times$	$\times$	$\times$	$\times$	$\times$	4.81	$\times$
1.2	3.23	10.46	25.73	0.61	0.81	0.98	$\times$	$\times$	$\times$	$\times$	$\times$	$\times$	4.96	$\times$
0.8	2.45	10.44	31.03	0.61	0.77	0.85	2.45	10.44	31.03	0.61	0.77	0.85	5.38	89.79
0.4	1.63	5.46	27.04	0.56	0.68	0.73	1.68	5.75	27.43	0.63	0.74	0.79	6.07	77.05
OU-NTS														
α	Price FGMC [%]			SD FGMC [bp]			Price ED [%]			SD ED [bp]			Time [s]	
	OTM	ATM	ITM	OTM	ATM	ITM	OTM	ATM	ITM	OTM	ATM	ITM	FGMC	ED
0.8	0.27	6.17	26.11	0.08	0.31	0.45	0.26	6.19	26.10	0.08	0.31	0.45	5.25	47.92
0.6	0.27	6.18	26.11	0.08	0.31	0.45	0.27	6.18	26.10	0.08	0.31	0.45	5.99	46.85
0.4	0.27	6.17	26.11	0.08	0.31	0.45	0.28	6.16	26.10	0.08	0.31	0.45	6.32	45.76
0.2	0.28	6.17	26.11	0.08	0.31	0.45	0.28	6.16	26.10	0.08	0.31	0.45	6.28	73.62
-1.0	0.28	6.16	26.11	0.08	0.31	0.45	$\times$	$\times$	$\times$	$\times$	$\times$	$\times$	5.16	$\times$
-2.0	0.28	6.15	26.11	0.08	0.31	0.45	$\times$	$\times$	$\times$	$\times$	$\times$	$\times$	5.69	$\times$
NTS-OU														
α	Price FGMC [%]			SD FGMC [bp]			Price ED [%]			SD ED [bp]			Time [s]	
	OTM	ATM	ITM	OTM	ATM	ITM	OTM	ATM	ITM	OTM	ATM	ITM	FGMC	ED
0.8	0.07	3.74	25.26	0.04	0.20	0.29	0.07	3.74	25.25	0.05	0.20	0.29	5.40	42.47
0.6	0.07	3.66	25.26	0.04	0.20	0.29	0.07	3.66	25.26	0.04	0.20	0.29	5.70	46.87
0.4	0.07	3.61	25.26	0.04	0.20	0.29	0.07	3.60	25.25	0.04	0.20	0.29	6.24	49.97
0.2	0.07	3.61	25.26	0.04	0.20	0.29	0.07	3.56	25.26	0.04	0.21	0.29	7.22	51.77

7. Conclusion

In this paper, we have introduced a new unified framework for simulating Lévy-driven OU, a class of processes which is extremely relevant in the energy sector due to its ability to capture key features of market dynamics. The proposed approach, named FGMC, relies on the reconstruction of the CDF from the CF of the stochastic process, leveraging complex-integration, FFT, and spline interpolation. Following a detailed discussion of the method, we have performed an empirical analysis of its accuracy and computational efficiency, and we have illustrated two applications in option pricing.

There are three main contributions of this paper. First, we have developed a *fast* and *general* simulation method for Lévy-driven OU processes. Algorithmic stability is ensured through complex-plane shifts, leveraging the novel results on CF regularity for these processes established in Propositions 1 and 2. Notably, we have also demonstrated FGMC applicability to novel subclasses of processes, where no alternative simulation algorithms currently exist.

Second, we have designed a specialised algorithm for FA Lévy-driven OU processes that incorporates a Bernoulli-based correction to optimally handle discontinuous CDFs. For both simulation algorithms, we have demonstrated that the CDF numerical error $\mathcal{E}(x)$ can be explicitly computed and controlled. Moreover, we have shown that for all main processes considered in the literature, the error $\mathcal{E}(x)$ decays exponentially with the FFT grid size for infinite-activity processes, and exhibits power-law decay for finite-activity processes, as established in Proposition 6.

Third, we have conducted a comprehensive experimental comparison of our method against the existing ED algorithms, which represent the state-of-the-art techniques for simulating Lévy-driven OU processes. The results show that, in all cases, FGMC achieves equivalent accuracy while delivering runtimes at least one order of magnitude faster. Simulation times are consistently comparable to those of the Gaussian OU – the least computationally demanding process within the entire OU family – rendering FGMC suitable for real-time applications. Furthermore, the proposed algorithms can be easily parallelised across multiple cores, offering additional opportunities for performance gains.

A natural extension of this work concerns high-frequency sampling with very small time increments. In such settings, the FFT inversion may become numerically sensitive, as the distribution of the process increments increasingly concentrates around zero. The development of robust FFT-based simulation schemes suitable for this regime constitutes a promising direction for future research.

Data availability

Data will be made available on request.

Disclosure of competing interest

The authors have no relevant financial or non-financial interests to disclose.

Acknowledgment

We are grateful to M. Azzone, S. Castellano, A. Pallavicini, L. Torricelli, and F. Ziel for their valuable suggestions. We also wish to thank all participants of the XXV Quantitative Finance Workshop (QFW) in Bologna, of the Energy Finance Italia Conference (EF19) in Bari, and of the International Ruhr Energy Conference (INREC13) in Essen for their insightful feedback. Finally, a special thank goes to the R33 team at EDF R&D Paris-Saclay, where part of this research has been developed.

Appendix A. Proofs

In this Appendix, we collect the proofs for the results presented in the main text. We first prove a useful proposition, which is instrumental to the proofs of [Propositions 1](#) and [2](#).

Proposition A.1. *Let X be an infinitely divisible r.v. with CF ϕ_X and LCF Ψ_X . Then, the analyticity strips of ϕ_X and Ψ_X coincide.*

Proof. Let $S := \{u \in \mathbb{C} : -p_+ < \text{Im}(u) < p_-\} \neq \emptyset$ be the analyticity strip of the CF ϕ_X . Since X is infinitely divisible, ϕ_X has no zeros in S (cf. [\[28, Theorem 5.4a, p.21\]](#)).

The strip S is simply connected, and ϕ_X is holomorphic and non-vanishing there. By a standard result of complex analysis (cf., e.g., [\[45, Theorem 16, p.143\]](#)), the complex logarithm of ϕ_X admits a single-valued holomorphic branch on S . Hence, the LCF $\Psi_X(u) := \log \phi_X(u)$ – originally defined only on the real axis – can be analytically continued at least to all of S (see also [\[25, Theorem 25.17, p.165\]](#)); therefore, the analyticity strip of Ψ_X , denoted as S' , satisfies $S' \supseteq S$.

Conversely, since the complex exponential e^z is an entire function, $\phi_X(u) = e^{\Psi_X(u)}$ is analytic at every point where Ψ_X is analytic, implying $S \supseteq S'$. Together, the two inclusions yield $S = S'$. $\square$

Proof of Proposition 1. Fix $t > 0$, and let S_L denote the analyticity strip of the characteristic exponent ψ_L . We assume the strip S_L to be bounded; the unbounded case follows immediately. Define

$$\zeta(u, s) := \psi_L(ue^{-bs}), \quad \text{with } (u, s) \in S_L \times [0, t].$$

The function $\zeta(u, s)$ satisfies two properties:

- (i) For each fixed $s \in [0, t]$, we have $0 < e^{-bs} \leq 1$, so the function $u \mapsto ue^{-bs}$ maps S_L into itself. Since ψ_L is holomorphic on S_L , it follows that $\zeta(\cdot, s)$ is holomorphic on S_L .
- (ii) The function $\zeta(u, s)$ is jointly continuous in (u, s) on $S_L \times [0, t]$, because it is a composition of continuous functions.

By Theorem 5.4 in Stein and Shakarchi [\[46, p. 56\]](#), these two conditions guarantee that the integral

$$\Psi_Z(u, t) = \int_0^t \zeta(u, s) ds = \int_0^t \psi_L(ue^{-bs}) ds$$

is holomorphic for every $u \in S_L$.

Conversely, for any singularity of ψ_L on the boundary of S_L there exist points $\tilde{u} \in \mathbb{C}$ arbitrarily close to the boundary, but outside S_L , such that the set $\{\tilde{u}e^{-bs} : s \in [0, t]\}$ intersects this singularity. For such $\tilde{u}$, the integral is ill-defined, entailing that $\Psi_Z(\cdot, t)$ cannot be analytically continued outside S_L .

Hence, the LCF Ψ_Z is holomorphic in exactly the same strip as the characteristic exponent ψ_L . Since both Z_t and L_t are infinitely divisible, [Proposition A.1](#) allows us to translate this analyticity result into the CFs $\phi_Z(u, t)$ and $\phi_L(u)$, concluding the proof. $\square$

Proof of Proposition 2. Fix $t > 0$, and let S_X denote the analyticity strip of the characteristic exponent ψ_X . We assume the strip S_X to be bounded; the unbounded case follows immediately. Since $0 < e^{-bt} \leq 1$, for every $u \in S_X$ both terms $\psi_X(u)$ and $\psi_X(ue^{-bt})$ in (8) are holomorphic. Hence, $\Psi_Z(u, t)$ is holomorphic for all $u \in S_X$.

Conversely, let $\tilde{u} \in \mathbb{C}$ be a singularity of ψ_X on the boundary of S_X . At $\tilde{u}$, the first term $\psi_X(\tilde{u})$ is non-holomorphic, while the second term $\psi_X(\tilde{u}e^{-bt})$ remains holomorphic since $\tilde{u}e^{-bt} \in S_X$. Thus, $\Psi_Z(\tilde{u}, t)$ is non-holomorphic on the boundary of S_X .

Therefore, the LCF $\Psi_Z(u, t)$ is holomorphic in exactly the same strip as ψ_X . Since both Z_t and X are infinitely divisible, Proposition A.1 allows us to translate this result into the respective CFs. $\square$

Proof of Proposition 3. Since the Lévy driver L_t is a pure-jump process with Lévy density $\nu_L(\cdot)$ (i.e. its Lévy measure is absolutely continuous with respect to the Lebesgue measure), the Lévy measure of Z_t is also absolutely continuous for every $t > 0$, with Lévy density $\nu_Z(\cdot, t)$ given by (6). For any such t , two cases are possible:

- (i) Z_t has infinite-activity, i.e. $\int_{\mathbb{R}} \nu_Z(x, t) dx = \infty$. In this case, a result of Sato [25, Theorem 27.7, p.177] ensures the absolute continuity of the law of the process at any finite time.
- (ii) Z_t has finite-activity, i.e. $\int_{\mathbb{R}} \nu_Z(x, t) dx < \infty$. In this case, Z_t is equivalent in law to a compound Poisson r.v. whose i.i.d. jumps are absolutely continuous, as discussed in detail in Section 4.2. A mixture of absolutely continuous laws remains absolutely continuous; therefore, conditionally on at least one jump occurring, Z_t is absolutely continuous. Unconditionally, Z_t has a mixed distribution with an atom at zero (if no jumps occur) and an absolutely continuous part elsewhere. $\square$

Proof of Proposition 4. The approximation error $\mathcal{E}(x)$ is given by the absolute difference between the exact CDF $P(x)$ and its approximation $\hat{P}(x)$ in (15). Letting $\hat{P}_\infty(x)$ denote the limit of $\hat{P}(x)$ as the number of terms N in the summation goes to infinity, we write

$$\mathcal{E}(x) := |P(x) - \hat{P}(x)| \leq |P(x) - \hat{P}_\infty(x)| + |\hat{P}_\infty(x) - \hat{P}(x)|,$$

decomposing the approximation error into a discretisation term and a truncation term, respectively.

For the truncation error, we distinguish two cases:

- (i) if the decay of the CF modulus is exponential, one has

$$\begin{aligned} |\hat{P}_\infty(x) - \hat{P}(x)| &= \frac{he^{-ax}}{\pi} \left| \sum_{n=N}^{\infty} \frac{\phi\left(\left(n + \frac{1}{2}\right)h - ia\right)}{i\left(n + \frac{1}{2}\right)h + a} e^{-ix\left(n + \frac{1}{2}\right)h} \right| \leq \frac{he^{-ax}}{\pi} \sum_{n=N}^{\infty} \left| \frac{\phi\left(\left(n + \frac{1}{2}\right)h - ia\right)}{i\left(n + \frac{1}{2}\right)h + a} \right| \leq \\ &\stackrel{(A)}{\leq} \frac{Bhe^{-ax}}{\pi} \sum_{n=N}^{\infty} \frac{e^{-\ell\left(n + \frac{1}{2}\right)h}}{\left(n + \frac{1}{2}\right)h} \stackrel{(B)}{\leq} \frac{Be^{-ax}}{\pi\omega} \int_{\ell(Nh)^\omega}^{\infty} \frac{e^{-y}}{y} dy \stackrel{(C)}{\leq} \frac{Be^{-ax}}{\pi\omega} \frac{e^{-\ell(Nh)^\omega}}{\ell(Nh)^\omega}. \end{aligned}$$

Inequality (A) is due to the fact that $|i(n + 1/2)h + a| \geq (n + 1/2)h$; inequality (B) exploits the eventual convexity of the function $e^{-\ell x^\omega}/x$ to bound the summation with the integral; inequality (C) is given by (5.1.20) in Abramowitz et al. [47], p.229.

- (ii) if the decay of the CF modulus is power-law, one has

$$\begin{aligned} |\hat{P}_\infty(x) - \hat{P}(x)| &= \frac{he^{-ax}}{\pi} \left| \sum_{n=N}^{\infty} \frac{\phi\left(\left(n + \frac{1}{2}\right)h - ia\right)}{i\left(n + \frac{1}{2}\right)h + a} e^{-ix\left(n + \frac{1}{2}\right)h} \right| \leq \frac{he^{-ax}}{\pi} \sum_{n=N}^{\infty} \left| \frac{\phi\left(\left(n + \frac{1}{2}\right)h - ia\right)}{i\left(n + \frac{1}{2}\right)h + a} \right| \leq \\ &\stackrel{(A)}{\leq} \frac{Bhe^{-ax}}{\pi} \sum_{n=N}^{\infty} \left(\left(n + \frac{1}{2}\right)h\right)^{-\omega-1} \stackrel{(B)}{\leq} \frac{Be^{-ax}}{\pi} \int_{Nh}^{\infty} x^{-\omega-1} dx = \frac{Be^{-ax}}{\pi\omega} (Nh)^{-\omega}. \end{aligned}$$

Inequality (A) is due to the fact that $|i(n + 1/2)h + a| \geq (n + 1/2)h$, as in the previous case, and inequality (B) exploits the convexity of the function $x^{-\omega-1}$.

For the discretisation error, we rely on Theorems 6.2 and 6.5 in Lee [29]. For our purposes, the relevant case is Lee's G_3 payoff with $b_0 = 1$ and $b_1 = 0$, and we set $p = 2a$ to obtain the appropriate form of the discretisation bound. $\square$

Proof of Lemma 5. In the IA case, we consider the complex term appearing inside the integral in (B.3), evaluated at $u - ia$:

$$w := \frac{1}{2}\sigma^2(u - ia)^2 z^2 - i\theta(u - ia)z + \frac{1 - \alpha}{\kappa}.$$

Let us first observe that, for every a in the NTS analyticity strip, we have $\text{Re}(w) > 0$: indeed,

$$\text{Re}(w) = \frac{1}{2}\sigma^2(u^2 - a^2)z^2 - \theta az + \frac{1 - \alpha}{\kappa} \geq -\frac{1}{2}\sigma^2 a^2 z^2 - \theta az + \frac{1 - \alpha}{\kappa} > 0,$$

where the last inequality coincides with the analyticity condition in Lemma D.1. Then, as $|u| \rightarrow \infty$,

$$\text{Re}(w^\alpha) = |w|^\alpha \cos\left[\alpha \arctan\left(\frac{\text{Im}(w)}{\text{Re}(w)}\right)\right] = \left(\frac{1}{2}\sigma^2 z^2\right)^\alpha |u|^{2\alpha} + o(|u|^{2\alpha}).$$

By computing the integral in (B.3), we deduce that the behaviour for large $|u|$ is

$$\int_{e^{-bt}}^1 \left(\frac{1}{2} \sigma^2 z^2 \right)^\alpha |u|^{2\alpha} \frac{dz}{z} + o(|u|^{2\alpha}) = \left(\frac{\sigma^2}{2} \right)^\alpha \frac{1 - e^{-2\alpha bt}}{2\alpha} |u|^{2\alpha} + o(|u|^{2\alpha}).$$

In the FA case, we study the modulus of the conditional CF ϕ_Z^c in (19), namely

$$\left| \phi_Z^c(u - ia, t) \right| = \frac{\left| e^{f(u-ia, t)} - 1 \right|}{e^{\Lambda(t)} - 1}, \quad (\text{A.1})$$

where

$$f(u, t) := \Lambda(t) \phi_J(u, t) = \left(\frac{1-\alpha}{\kappa} \right)^{1-\alpha} \frac{1}{|\alpha|b} \int_{e^{-bt}}^1 \left(\frac{1}{2} \sigma^2 u^2 z^2 - i\theta uz + \frac{1-\alpha}{\kappa} \right)^\alpha \frac{dz}{z}.$$

By continuity of the complex function $z \mapsto |e^z - 1|$ at $z = 0$, it follows that, as $|u| \rightarrow \infty$,

$$\left| e^{f(u-ia, t)} - 1 \right| = |f(u - ia, t)| + o(|f(u - ia, t)|). \quad (\text{A.2})$$

As in the IA case, it is straightforward to deduce that – for every a in the analyticity strip – the asymptotic behaviours of the real and imaginary parts of $f(u - ia, t)$ for large $|u|$ are:

$$\begin{cases} \operatorname{Re}(f(u - ia, t)) = \left(\frac{\sigma^2}{2} \right)^\alpha \left(\frac{1-\alpha}{\kappa} \right)^{1-\alpha} \frac{e^{-2\alpha bt} - 1}{2\alpha^2 b} |u|^{2\alpha} + o(|u|^{2\alpha}), \\ \operatorname{Im}(f(u - ia, t)) = o(|u|^{2\alpha}). \end{cases}$$

The thesis follows by substituting these expressions into (A.1) and (A.2). $\square$

Proof of Proposition 6. Using Algorithm 1 for IA processes and Algorithm 2 for FA processes, we observe distinct asymptotic behaviours of the CF moduli: exponential decay for the former and power-law decay for the latter, as proved in Lemmas 5, C.5, C.6, and D.5. The numerical error $\mathcal{E}(x)$ exhibits a decay rate consistent with that of the CFs, as proved in Proposition 4. $\square$

Appendix B. Lévy-driven OU processes of TS and NTS type

In this Appendix, we characterise the four main Lévy-driven OU processes considered in the literature, which are based on two well-known Lévy processes – TS and NTS – and the associated self-decomposable distributions. In the following, we report the LCFs of these processes, derived using (5) and (8).¹³

(i) An **OU-TS process** is an OU-Lévy whose driver is a TS process. It is defined by 7 parameters: $\alpha_p, \alpha_n < 2$, $\beta_p, \beta_n > 0$, $c_p, c_n \geq 0$, and $\gamma_c \in \mathbb{R}$. The LCF of the process Z_t in (4) reads

$$\begin{aligned} \Psi_Z(u, t) = iu \frac{1 - e^{-bt}}{b} \gamma_c + \frac{c_p \beta_p^{\alpha_p} \Gamma(-\alpha_p)}{b} \left[\int_{\beta_p}^{\beta_p e^{bt}} \frac{(z - iu)^{\alpha_p}}{z^{\alpha_p+1}} dz - bt + \frac{\alpha_p}{\beta_p} iu(1 - e^{-bt}) \right] + \\ + \frac{c_n \beta_n^{\alpha_n} \Gamma(-\alpha_n)}{b} \left[\int_{\beta_n}^{\beta_n e^{bt}} \frac{(z + iu)^{\alpha_n}}{z^{\alpha_n+1}} dz - bt - \frac{\alpha_n}{\beta_n} iu(1 - e^{-bt}) \right], \end{aligned} \quad (\text{B.1})$$

valid when both $\alpha_p, \alpha_n \notin \{0, 1\}$. As noted by [12], the integrals in (B.1) can also be expressed via the Gaussian hypergeometric function ${}_2F_1(a, b, c; x)$, for which fast computational routines exist.

(ii) A **TS-OU process** is a Lévy-OU whose stationary distribution is a TS. It is defined by 7 parameters: $\alpha_p, \alpha_n \in [0, 2)$, $\beta_p, \beta_n > 0$, $c_p, c_n \geq 0$, and $\gamma_c \in \mathbb{R}$. The LCF of the process Z_t is

$$\begin{aligned} \Psi_Z(u, t) = iu(1 - e^{-bt}) \left[\gamma_c - c_p \Gamma(1 - \alpha_p) \beta_p^{\alpha_p-1} + c_n \Gamma(1 - \alpha_n) \beta_n^{\alpha_n-1} \right] + \\ + c_p \Gamma(-\alpha_p) \left[(\beta_p - iu)^{\alpha_p} - (\beta_p - iue^{-bt})^{\alpha_p} \right] + \\ + c_n \Gamma(-\alpha_n) \left[(\beta_n + iu)^{\alpha_n} - (\beta_n + iue^{-bt})^{\alpha_n} \right], \end{aligned} \quad (\text{B.2})$$

valid when both $\alpha_p, \alpha_n \notin \{0, 1\}$. The TS-OU is defined only for non-negative α_p and α_n , as the TS distribution is self-decomposable only in this parameter range (see, e.g., Lemma C.3).

¹³ For brevity, we omit the discussion of special parameter cases, namely $\alpha_p, \alpha_n \in \{0, 1\}$ for OU-TS and TS-OU processes, and $\alpha = 0$ for OU-NTS and NTS-OU processes.

(iii) An **OU-NTS process** is a OU-Lévy whose driver is an NTS process. It is defined by 4 parameters: $\alpha < 1$, $\kappa > 0$, $\sigma > 0$, and $\theta \in \mathbb{R}$. The LCF of the process Z_t is

$$\Psi_Z(u, t) = \frac{1-\alpha}{\kappa\alpha} t - \left(\frac{1-\alpha}{\kappa}\right)^{1-\alpha} \frac{1}{\alpha b} \int_{e^{-bt}}^1 \left(\frac{1}{2}\sigma^2 u^2 z^2 - i\theta uz + \frac{1-\alpha}{\kappa}\right)^{\frac{\alpha}{1-\alpha}} \frac{dz}{z}, \quad (\text{B.3})$$

valid when $\alpha \neq 0$. Importantly, the case $\theta \neq 0$ is not considered in the literature on ED algorithms. Sabino [14] treats only the case $\theta = 0$, calling it the *symmetric* case.

(iv) Finally, an **NTS-OU process** is a Lévy-OU whose stationary distribution is an NTS. It is defined by 4 parameters: $\alpha \in [0, 1)$, $\kappa > 0$, $\sigma > 0$, and $\theta \in \mathbb{R}$. The LCF of the process Z_t is

$$\Psi_Z(u, t) = \frac{1-\alpha}{\kappa\alpha} \left[\left(1 - \frac{i\kappa}{1-\alpha} \left(\theta u e^{-bt} + i \frac{\sigma^2 u^2 e^{-2bt}}{2} \right) \right)^{\alpha} - \left(1 - \frac{i\kappa}{1-\alpha} \left(\theta u + i \frac{\sigma^2 u^2}{2} \right) \right)^{\alpha} \right], \quad (\text{B.4})$$

valid when $\alpha \neq 0$. The NTS-OU is defined only for non-negative α , since the NTS distribution is self-decomposable only in this parameter range (see, e.g., Lemma D.3). Existing simulation methods for this process are also limited to the case $\theta = 0$ (cf. [13]).

Appendix C. The Tempered Stable process

The TS process is a widely used Lévy process that encompasses several well-known subcases such as the CGMY process [48]. It is obtained by exponential tempering of stable Lévy densities, which ensures finite moments of all orders while preserving high flexibility for financial modelling (cf. [26, Section 4.5]).

Definition. A TS process is a Lévy process with the following characteristic triplet $(0, \nu, \gamma)$:

$$\begin{cases} \nu(x) = c_p \frac{e^{-\beta_p x}}{x^{1+\alpha_p}} \mathbb{1}_{x>0} + c_n \frac{e^{-\beta_n |x|}}{|x|^{1+\alpha_n}} \mathbb{1}_{x<0}, \\ \gamma = \gamma_c - c_p \beta_p^{\alpha_p-1} \Gamma_U(1-\alpha_p, \beta_p) + c_n \beta_n^{\alpha_n-1} \Gamma_U(1-\alpha_n, \beta_n), \end{cases} \quad (\text{C.1})$$

where $\alpha_p, \alpha_n < 2$, $\beta_p, \beta_n > 0$, $c_p, c_n \geq 0$, $\gamma_c \in \mathbb{R}$ are the model parameters, and $\Gamma_U(\alpha, x)$ denotes the upper incomplete gamma function. For $\alpha_p, \alpha_n \notin \{0, 1\}$, its characteristic exponent reads

$$\psi(u) = iu\gamma_c + c_p \Gamma(-\alpha_p) \left[(\beta_p - iu)^{\alpha_p} - \beta_p^{\alpha_p} + iu\alpha_p \beta_p^{\alpha_p-1} \right] + c_n \Gamma(-\alpha_n) \left[(\beta_n + iu)^{\alpha_n} - \beta_n^{\alpha_n} - iu\alpha_n \beta_n^{\alpha_n-1} \right]. \quad (\text{C.2})$$

The following two lemmas establish important properties of the TS process that are used throughout the paper. The first defines the analyticity strip of the process, which determines the valid range of complex shifts for our simulation method (see Section 4).

Lemma C.1. For a TS process, the analyticity strip of the CF is given by

$$(-p_+, p_-) = (-\beta_p, \beta_n). \quad (\text{C.3})$$

In the unilateral positive case (i.e. when $c_n = 0$), the strip reduces to $(-\beta_p, +\infty)$.

Proof. See, e.g., [49, Lemma 2.8, p.7]. $\square$

The second lemma provides exact expressions for the cumulants of the TS process, which are used to test the accuracy of our simulation method in Section 6.

Lemma C.2. Let L_t be a TS process. Then, the cumulants of L_1 are given by:

$$c_k(L_1) = \begin{cases} \gamma_c & \text{if } k = 1, \\ c_p \beta_p^{\alpha_p-k} \Gamma(k-\alpha_p) + (-1)^k c_n \beta_n^{\alpha_n-k} \Gamma(k-\alpha_n) & \text{if } k \geq 2. \end{cases} \quad (\text{C.4})$$

Proof. See, e.g., [13, p.21]. $\square$

C.1. Additional properties of OU-TS and TS-OU

While TS processes describe continuous-time trajectories, the distribution at a fixed time is often the primary object of interest in applications. The TS distribution is defined as the law of a TS process at time $t = 1$.

Definition. A r.v. X follows a TS distribution if its LCF is $\psi(u)$ in (C.2). Equivalently, the r.v. X has the same law as L_1 , where L_t is a TS process.

The following lemma discusses the self-decomposability of the TS distribution, which is key for the construction of OU-Lévy processes (see Section 3).

Lemma C.3. The TS distribution is self-decomposable if and only if both α_p and α_n are non-negative.

Proof. The result follows directly from Theorem 15.11 in [25, p.95]. $\square$

The Lévy-driven OU processes of TS type defined in [Appendix B](#) exhibit distinct behaviours depending on the stability parameters α_p and α_n . Following the literature (see, e.g., [\[26, Section 3.5\]](#)), we classify these processes by their activity and variation properties. This classification is crucial for our methodology, as it determines the selection of the simulation algorithm (see [Section 4](#)).

Lemma C.4. *The TS, OU-TS and TS-OU processes can be classified, based on the values of α_p and α_n , as in the following table:*

	Finite-Activity	Infinite-Activity	
		Finite-Variation	Infinite-Variation
TS	$\max\{\alpha_p, \alpha_n\} < 0$	$\max\{\alpha_p, \alpha_n\} \in [0, 1)$	$\max\{\alpha_p, \alpha_n\} \in [1, 2)$
OU-TS	$\max\{\alpha_p, \alpha_n\} < 0$	$\max\{\alpha_p, \alpha_n\} \in [0, 1)$	$\max\{\alpha_p, \alpha_n\} \in [1, 2)$
TS-OU	$\alpha_p = \alpha_n = 0$	$\max\{\alpha_p, \alpha_n\} \in (0, 1)$	$\max\{\alpha_p, \alpha_n\} \in [1, 2)$

Proof. The TS density is given in [\(C.1\)](#); those of OU-TS and TS-OU are derived from [\(6\)](#) and [\(9\)](#). The classification follows from the standard definitions of activity and variation (see, e.g., [\[25\]](#)). $\square$

Lastly, we establish in the following lemmas the asymptotic behaviour of the CF for the OU-TS and TS-OU processes, which complements the results presented in [Section 5.2](#). For simplicity, we focus on the unilateral positive case ($c_n = 0$); the generalisation to the bilateral case is straightforward.

Lemma C.5. *For an OU-TS process with $\alpha_p \in (0, 1) \cup (1, 2)$, the CF ϕ_Z has exponential decay as $|u| \rightarrow \infty$:*

$$\log |\phi_Z(u - ia, t)| = c_p \Gamma(-\alpha_p) \cos\left(\alpha_p \frac{\pi}{2}\right) \frac{1 - e^{-\alpha_p b t}}{\alpha_p b} |u|^{\alpha_p} + o(|u|^{\alpha_p}) \quad \forall a \in (-\infty, \beta_p).$$

For an OU-TS process with $\alpha_p < 0$, the conditional CF ϕ_Z^c has power-law decay as $|u| \rightarrow \infty$:

$$|\phi_Z^c(u - ia, t)| = \frac{c_p \Gamma(-\alpha_p)}{e^{\Lambda(t)} - 1} \frac{1 - e^{-\alpha_p b t}}{\alpha_p b} |u|^{\alpha_p} + o(|u|^{\alpha_p}) \quad \forall a \in (-\infty, \beta_p),$$

where $\Lambda(t) = c_p \beta_p^{\alpha_p} \Gamma(-\alpha_p) t$.

Lemma C.6. *For a TS-OU process with $\alpha_p \in (0, 1) \cup (1, 2)$, the CF ϕ_Z has exponential decay as $|u| \rightarrow \infty$:*

$$\log |\phi_Z(u - ia, t)| = c_p \Gamma(-\alpha_p) \cos\left(\alpha_p \frac{\pi}{2}\right) (1 - e^{-\alpha_p b t}) |u|^{\alpha_p} + o(|u|^{\alpha_p}) \quad \forall a \in (-\infty, \beta_p).$$

Proof. The proofs of both lemmas are analogous to [Lemma 5](#). $\square$

Appendix D. The Normal Tempered Stable process

The NTS process L_t is a time-changed Lévy process of the form $L_t = \theta S_t + \sigma W_{S_t}$, where $\theta \in \mathbb{R}$, $\sigma > 0$, W_t is a standard BM, and S_t is a TS subordinator¹⁴ (cf. [\[26, Section 4.4\]](#)). Throughout this work, we consider the most general specification of the NTS process, in which the stability parameter α is allowed to take any value in $(-\infty, 1)$, ensuring consistency with the TS process in [Appendix C](#).

Definition. An NTS process is a Lévy process with the following characteristic triplet $(0, \nu, \gamma)$:

$$\begin{cases} \nu(x) = \frac{C(\alpha, \kappa, \sigma, \theta)}{|x|^{\alpha+1/2}} e^{\theta x / \sigma^2} K_{\alpha+1/2}\left(\frac{|x| A(\alpha, \kappa, \sigma, \theta)}{\sigma^2}\right), \\ \gamma = \frac{1}{\Gamma(1-\alpha)} \left(\frac{1-\alpha}{\kappa}\right)^{1-\alpha} \int_0^{+\infty} \frac{dt}{t^{\alpha+1}} e^{-\frac{(1-\alpha)t}{\kappa}} \int_{-1}^1 dx \frac{x}{\sqrt{2\pi t \sigma^2}} e^{-\frac{(x-\theta)^2}{2t\sigma^2}}, \end{cases} \quad (\text{D.1})$$

where $\alpha < 1$, $\kappa > 0$, $\theta \in \mathbb{R}$, $\sigma > 0$ are the model parameters, $K_m(\cdot)$ is the modified Bessel function of second kind with order m , and

$$C(\alpha, \kappa, \sigma, \theta) := \frac{2}{\Gamma(1-\alpha)\sigma\sqrt{2\pi}} \left(\frac{1-\alpha}{\kappa}\right)^{1-\alpha} A(\alpha, \kappa, \sigma, \theta)^{\alpha+1/2}, \quad (\text{D.2})$$

$$A(\alpha, \kappa, \sigma, \theta) := \sqrt{\theta^2 + 2\sigma^2(1-\alpha)/\kappa}. \quad (\text{D.3})$$

For $\alpha \neq 0$, its characteristic exponent reads

$$\psi(u) = \frac{1-\alpha}{\kappa\alpha} \left[1 - \left(1 - \frac{i\kappa}{1-\alpha} \left(\theta u + i \frac{u^2 \sigma^2}{2} \right) \right)^\alpha \right]. \quad (\text{D.4})$$

The next two lemmas state fundamental properties of the NTS process. The first defines its analyticity strip, establishing the range of valid complex shifts for our simulation method in [Section 4](#).

¹⁴ The TS subordinator S_t (i.e. a positive unilateral TS process) is parametrised using only two parameters, $\alpha < 1$ and $\kappa > 0$, by setting $\alpha_p = \alpha$, $\beta_p = (1-\alpha)/\kappa$, $c_p = ((1-\alpha)/\kappa)^{1-\alpha}/\Gamma(1-\alpha)$ and $\gamma_c = 1$ (cf. [\[26, p. 116\]](#)).

Lemma D.1. The analyticity strip of the NTS process is given by

$$(-p_+, p_-) = \left(-\frac{A(\alpha, \kappa, \sigma, \theta) - \theta}{\sigma^2}, \frac{A(\alpha, \kappa, \sigma, \theta) + \theta}{\sigma^2} \right). \quad (\text{D.5})$$

Proof. See Appendix D.2. $\square$

The second lemma derives a closed-form expression for the NTS cumulants for general order k , which, to our knowledge, has not appeared explicitly in the literature before.

Lemma D.2. Let L_t be an NTS process. Then, the cumulants of L_1 are given by:

$$c_k(L_1) = \sum_{n=0}^{\lfloor k/2 \rfloor} \frac{k!}{n! (k-2n)!} \theta^{k-2n} \left(\frac{\sigma^2}{2} \right)^n \frac{\Gamma(k-n-\alpha)}{\Gamma(1-\alpha)} \left(\frac{\kappa}{1-\alpha} \right)^{k-n-1}, \quad (\text{D.6})$$

where the usual convention $0^0 = 1$ applies.

Proof. See Appendix D.2. $\square$

D.1. Additional properties of OU-NTS and NTS-OU

Following the standard convention for Lévy processes, the NTS distribution is defined as the law of an NTS process at time $t = 1$.

Definition. A r.v. X follows an NTS distribution if its CF is $\phi_X(u) = e^{\psi(u)}$, with $\psi(u)$ in (D.4). Equivalently, the r.v. X has the same law as L_1 , where L_t is an NTS process.

The next lemma addresses the self-decomposability of the NTS distribution, which is essential for the construction of OU-Lévy processes (see Section 3).

Lemma D.3. The NTS distribution is self-decomposable if and only if α is non-negative.

Proof. See, e.g., [50, Section 4, p.8]. $\square$

The Lévy-driven OU processes of NTS type described in Appendix B exhibit distinct behaviours depending on the stability parameter α . The following lemma classifies these processes by activity and variation, directly informing algorithm selection and error control (see Section 4).

Lemma D.4. The NTS, OU-NTS and NTS-OU processes can be classified based on the values of α as in the following table:

	Finite-Activity	Infinite-Activity	
		Finite-Variation	Infinite-Variation
NTS	$\alpha < 0$	$\alpha \in \left[0, \frac{1}{2} \right)$	$\alpha \in \left[\frac{1}{2}, 1 \right)$
OU-NTS	$\alpha < 0$	$\alpha \in \left[0, \frac{1}{2} \right)$	$\alpha \in \left[\frac{1}{2}, 1 \right)$
NTS-OU	$\alpha = 0$	$\alpha \in \left(0, \frac{1}{2} \right)$	$\alpha \in \left[\frac{1}{2}, 1 \right)$

Proof. The NTS density is given in (D.1); those of OU-NTS and NTS-OU are derived from (6) and (9). The classification follows from the standard definitions of activity and variation (see, e.g., [25]). $\square$

Lastly, the following lemma characterises the asymptotic behaviour of the CF of the NTS-OU process, complementing the OU-NTS results presented in Section 5.2.

Lemma D.5. For an NTS-OU process with $\alpha \in (0, 1)$, the CF ϕ_Z has exponential decay as $|u| \rightarrow \infty$:

$$\log |\phi_Z(u - ia, t)| = -\frac{(\sigma^2/2)^\alpha}{\alpha} \left(\frac{1-\alpha}{\kappa} \right)^{1-\alpha} (1 - e^{-2abt}) |u|^{2\alpha} + o(|u|^{2\alpha}) \quad \forall a \in (-p_-, p_+)$$

where $(-p_-, p_-)$ denotes the NTS analyticity strip (see Lemma D.1).

Proof. The proof is analogous to Lemma 5. $\square$

D.2. Proofs for NTS properties

Proof of Lemma D.1. By Theorem 3.1 in [28], the analyticity strip is determined by the singularities of the CF on the imaginary axis. Setting $u = i\xi$ in (D.4), with $\xi \in \mathbb{R}$, the characteristic exponent is well-defined if

$$1 + \frac{\kappa}{1-\alpha} \left(\theta \xi - \frac{1}{2} \sigma^2 \xi^2 \right) > 0.$$

The boundaries of the strip are obtained by solving this quadratic inequality in ξ . $\square$

Proof of Lemma D.2. Let L_t be an NTS process, and L_1 an NTS random variable. By definition, the cumulants are given by

$$c_k(L_1) = (-i)^k \frac{d^k \psi_L(u)}{du^k} \Big|_{u=0} = (-i)^k \frac{d^k}{du^k} \psi_S \left(\theta u + i \frac{\sigma^2}{2} u^2 \right) \Big|_{u=0}, \quad (\text{D.7})$$

where $\psi_S(\cdot)$ denotes the characteristic exponent of the TS subordinator S_t (cf., e.g., [26, Theorem 4.2, p. 115]).

For $k = 1$, direct differentiation of (D.4) gives $c_1(L_1) = \theta$.

For $k \geq 2$, we set $g(u) := \theta u + i\sigma^2 u^2/2$ and apply the Faà di Bruno formula:

$$\frac{d^k}{du^k} \psi_S(g(u)) \Big|_{u=0} = \sum_{n=1}^k \psi_S^{(n)}(g(0)) B_{k,n}(g'(0), g''(0), \dots, g^{(k-n+1)}(0)),$$

where $B_{k,n}$ are the partial Bell polynomials (see, e.g., [51]). Since $g(u)$ is quadratic, all derivatives of order higher than two are equal to zero. Hence, for every $k \geq 2$ and $n = 1, \dots, k$, the Bell polynomials reduce to

$$B_{k,n}(x_1, x_2, 0, 0, \dots) = \frac{k!}{(2n-k)!(k-n)! 2^{k-n}} x_1^{2n-k} x_2^{k-n} \mathbb{1}_{\{[k/2] \leq n \leq k\}}.$$

Substituting these expressions into (D.7) yields

$$c_k(L_1) = (-i)^k \sum_{n=[k/2]}^k i^n c_n(S_1) \frac{k!}{(2n-k)!(k-n)! 2^{k-n}} \theta^{2n-k} (i\sigma^2)^{k-n}, \quad (\text{D.8})$$

where $c_n(S_1) = (-i)^n \psi_S^{(n)}(0)$ are the cumulants of the TS subordinator, given in (C.4). Under the TS subordinator parametrisation, these cumulants read

$$c_n(S_1) = \left(\frac{\kappa}{1-\alpha} \right)^{n-1} \frac{\Gamma(n-\alpha)}{\Gamma(1-\alpha)}.$$

Plugging this expression into (D.8), simplifying the powers of i , and reindexing the summation via $n \mapsto k-n$, we obtain the stated closed-form result. $\square$

Symbols and Acronyms

Symbol	Description
L_t	Background driving Lévy process
X_t	Lévy-driven OU process defined in (4)
Z_t	Integral Process defined in (4)
$Z_t^{\mathbb{C}}$	For FA processes, <i>conditional</i> law of Z_t , cf. (18)
W_t	Brownian Motion
γ	Drift in the characteristic triplet
$\nu_L(x)$	Lévy density of a Lévy process L_t
$\phi_L(x)$	Characteristic Function of a Lévy process L_t
$\psi_L(x)$	Characteristic Exponent of a Lévy process L_t
$\nu_Z(x, t)$	Lévy density of an additive process Z_t
$\phi_Z(x, t)$	Characteristic Function of an additive process Z_t
$\Psi_Z(x, t)$	Log-Characteristic Function of an additive process Z_t
p_-, p_+	Boundaries of the analyticity strip
$c_k(Y)$	Cumulant of order k of a r.v. Y
$P(x)$	True CDF of the process innovation
$\hat{P}(x)$	Numerical approximation of the CDF of the process innovation
$\mathcal{E}(x)$	Numerical error in the CDF reconstruction
a	Imaginary shift of the integration path in the complex plane
h	Grid step in the Fourier domain
M	Integer number equal to $\log_2(N)$, with N number of FFT points
N	Number of points in the FFT grid, equal to 2^M
α	Stability parameter for TS and NTS process/distribution
b	Mean-reversion parameter in the OU equation
β, c, γ_c	Other parameters of TS process/distribution
κ, θ, σ	Other parameters of NTS process/distribution
B, ℓ, ω	Coefficients for the CF bounds, cf. (20) and (21)
$\Lambda(t)$	Intensity of Z_t in the compound Poisson representation
$\Phi(x)$	Standard Normal CDF
$\Gamma(x)$	Gamma function
$\Gamma_U(x)$	Incomplete upper gamma function
$E_1(x)$	Exponential integral function
$K_m(x)$	Modified Bessel function of second kind with order m
${}_2F_1(a, b; c; x)$	Gaussian hypergeometric function
χ	Moneyness level of an option
$\mathcal{U}(a, b)$	Uniform distribution on $[a, b]$
$O(\cdot)$	Landau's <i>Big-O</i> symbol
$o(\cdot)$	Landau's <i>Little-o</i> symbol

Acronym	Description
ATM	At-the-money
BM	Brownian Motion
bp	Basis Point
CDF	Cumulative Distribution Function
CF	Characteristic Function
ED	Exact Decomposition
FA	Finite-Activity
FFT	Fast Fourier Transform
FGMC	Fast and General Monte Carlo
IA	Infinite-Activity
i.i.d.	Independent and Identically Distributed
ITM	In-the-money
LCF	Log-Characteristic Function
MAPE	Mean Absolute Percentage Error
MC	Monte Carlo
NTS	Normal Tempered Stable
OTM	Out-of-the-money
OU	Ornstein-Uhlenbeck
PDF	Probability Density Function
RMSE	Root Mean Squared Error
r.v.	Random Variable
SD	Standard Deviation
SDE	Stochastic Differential Equation
TS	Tempered Stable
VG	Variance Gamma

References

- [1] F.E. Benth, J.S. Benth, S. Koekebakker, *Stochastic modelling of electricity and related markets*, World Scientific, 2008.
- [2] B. Hambly, S. Howison, T. Kluge, Modeling spikes and pricing swing options in electricity markets, *Quant. Finance* 9 12 (2009) 937–949.
- [3] G. Fusai, M.arena, A. Roncoroni, Analytical pricing of discretely monitored asian-style options: theory and application to commodity markets, *J. Bank. Finance* 32 (10) (2008) 2033–2045.
- [4] S. Zhang, X. Zhang, Exact simulation of IG-OU processes, *Methodol. Comput. Appl. Probab.* 10 (3) (2008) 337–355.
- [5] S. Zhang, X. Zhang, On the transition law of tempered stable Ornstein–Uhlenbeck processes, *J. Appl. Probab.* 46 (3) (2009) 721–731.
- [6] S. Zhang, Transition law-based simulation of generalized inverse Gaussian Ornstein–Uhlenbeck processes, *Methodol. Comput. Appl. Probab.* 13 (3) (2011) 619–656.
- [7] R. Kawai, H. Masuda, Exact discrete sampling of finite variation tempered stable Ornstein–Uhlenbeck processes, *Monte Carlo Methods Appl.* (2011) 279–300.
- [8] R. Kawai, H. Masuda, Infinite variation tempered stable Ornstein–Uhlenbeck processes with discrete observations, *Commun. Stat. Simul. Comput.* 41 (1) (2012) 125–139.
- [9] M.L. Bianchi, S.T. Rachev, F.J. Fabozzi, Tempered stable Ornstein–Uhlenbeck processes: a practical view, *Commun. Stat.-Simul. Comput.* 46 (1) (2017) 423–445.
- [10] Y. Qu, A. Dassios, H. Zhao, Exact simulation of Ornstein–Uhlenbeck tempered stable processes, *J. Appl. Probab.* 58 (2) (2021) 347–371.
- [11] P. Sabino, Exact simulation of variance gamma-related OU processes: application to the pricing of energy derivatives, *Appl. Math. Finance* 27 (3) (2020) 207–227.
- [12] P. Sabino, Pricing energy derivatives in markets driven by tempered stable and CGMY processes of Ornstein–Uhlenbeck type, *Risks* 10 (8) (2022) 148.
- [13] P. Sabino, Exact simulation of normal tempered stable processes of OU type with applications, *Stat. Comput.* 32 (5) (2022) 81.
- [14] P. Sabino, Normal tempered stable processes and the pricing of energy derivatives, *SIAM J. Financ. Math.* 14 (1) (2023) 99–126.
- [15] P. Sabino, N. Cufaro Petroni, Gamma-related Ornstein–Uhlenbeck processes and their simulation, *J. Stat. Comput. Simul.* 91 (6) (2021) 1108–1133.
- [16] P. Sabino, N. Cufaro Petroni, Fast simulation of tempered stable Ornstein–Uhlenbeck processes, *Comput. Stat.* 37 (5) (2022) 2517–2551.
- [17] Z. Chen, L. Feng, X. Lin, Simulating Lévy processes from their characteristic functions and financial applications, *ACM Trans. Model. Comput. Simul.* 22 (3) (2012) 1–26.
- [18] L. Ballotta, I. Kyriakou, Monte carlo simulation of the CGMY process and option pricing, *J. Futures Mark.* 34 (12) (2014) 1095–1121.
- [19] O.E. Barndorff-Nielsen, N. Shephard, Non-Gaussian Ornstein–Uhlenbeck-based models and some of their uses in financial economics, *J. R. Stat. Soc. B* 63 (2) (2001) 167–241.
- [20] B.D. Flury, Acceptance–rejection sampling made easy, *Siam Rev.* 32 (3) (1990) 474–476.
- [21] J. Gil-Pelaez, Note on the inversion theorem, *Biometrika* 38 (3-4) (1951) 481–482.
- [22] J.G. Wendel, The non-absolute convergence of Gil-Pelaez’ inversion integral, *Ann. Math. Stat.* 32 (1) (1961) 338–339.
- [23] J.W. Cooley, J.W. Tukey, An algorithm for the machine calculation of complex fourier series, *Math. Comput.* 19 (90) (1965) 297–301.
- [24] A. Quarteroni, R. Sacco, F. Saleri, *Numerical mathematics*, 37, Springer Science & Business Media, 2006.
- [25] K.I. Sato, *Lévy processes and infinitely divisible distributions*, 68, Cambridge university press, 1999.
- [26] R. Cont, P. Tankov, *Financial Modelling with Jump Processes*, Chapman & Hall/CRC, 2003.
- [27] E. Lukacs, A characterization of stable processes, *J. Appl. Probab.* 6 (2) (1969) 409–418.
- [28] E. Lukacs, A survey of the theory of characteristic functions, *Adv. Appl. Probab.* 4 (1) (1972) 1–37.
- [29] R.W. Lee, Option pricing by transform methods: extensions, unification and error control, *J. Comput. Finance* 7 (3) (2004) 51–86.
- [30] R. Baviera, M.D. Massaria, The additive Bachelier model with an application to the oil option market in the Covid period, *J. Comput. Appl. Math.*, 2026, 117741.
- [31] S.J. Wolfe, On a continuous analogue of the stochastic difference equation $X_n = \rho X_{n-1} + B_n$, *Stoch. Process. Their Appl.* 12 (3) (1982) 301–312.
- [32] E. Taufer, N. Leonenko, Simulation of Lévy-driven Ornstein-Uhlenbeck processes with given marginal distribution, *Comput. Stat. Data Anal.* 53 (6) (2009) 2427–2437.
- [33] M. Reed, B. Simon, *Methods of Modern Mathematical Physics, II: Fourier analysis, self-adjointness*, 2, Elsevier, 1975.
- [34] E. Lukacs, *Characteristic Functions*, Griffin, London, 2nd edition, 1970.

- [35] W.H. Press, S.A. Teukolsky, W.T. Vetterling, B.P. Flannery, Numerical recipes: The art of scientific computing, Cambridge University Press, third edition, 2007.
- [36] M. Azzone, R. Baviera, A fast Monte Carlo scheme for additive processes and option pricing, *Comput. Manag. Sci.* 20 (1) (2023) 31.
- [37] F.E. Benth, L. Di Persio, S. Lavagnini, Stochastic modeling of wind derivatives in energy markets, *Risks* 6 (2) (2018) 56.
- [38] A. Brix, Generalized gamma measures and shot-noise cox processes, *Adv. Appl. Probab.* 31 (4) (1999) 929–953.
- [39] L. Devroye, Random variate generation for exponentially and polynomially tilted stable distributions, *ACM Trans. Model. Comput. Simul.* 19 (4) (2009) 1–20.
- [40] M. Hofert, Sampling exponentially tilted stable distributions, *ACM Trans. Model. Comput. Simul.* 22 (1) (2011) 1–11.
- [41] M. Cummins, G. Kiely, B. Murphy, Gas storage valuation under Lévy processes using the fast fourier transform, *J. Energy Mark.* 10 (4) (2017) 43–86.
- [42] F.E. Benth, M. Piccirilli, T. Vargiolu, Mean-reverting additive energy forward curves in a Heath-Jarrow-Morton framework, *Math. Financ. Econ.* 13 (4) (2019) 543–577.
- [43] L. Latini, M. Piccirilli, T. Vargiolu, Mean-reverting no-arbitrage additive models for forward curves in energy markets, *Energy Econ.* 79 (2019) 157–170.
- [44] A.L. Lewis, A simple option formula for general jump-diffusion and other exponential Lévy processes, Available SSRN 282110 (2001).
- [45] L.V. Ahlfors, Complex analysis, 3, McGraw-Hill New York, 1979.
- [46] E.M. Stein, R. Shakarchi, Complex analysis, 2, Princeton University Press, 2010.
- [47] M. Abramowitz, I.A. Stegun, R.H. Romer, Handbook of mathematical functions with formulas, graphs, and mathematical tables, American Association of Physics Teachers, 1988.
- [48] P. Carr, H. Geman, D.B. Madan, M. Yor, The fine structure of asset returns: an empirical investigation, *J. Bus.* 75 (2) (2002) 305–332.
- [49] U. Küchler, S. Tappe, Tempered stable distributions and processes, *Stoch. Process. Their Appl.* 123 (12) (2013) 4256–4293.
- [50] O.E. Barndorff-Nielsen, N. Shephard, Normal modified stable processes, *Theory Probab. Math. Stat.* 65 (2001) 1–19.
- [51] L. Comtet, Advanced Combinatorics: The art of finite and infinite expansions, Springer Science & Business Media, 2012.