Chilean Journal of Statistics
Volume 16, Number 1, June 2025, Pages 79–97
DOI: 10.32372/ChJS.16-01-05

Computational Statistics and Random Number Generation
Research Article

Tail-adaptive generation of random numbers from a gamma-order normal distribution using the Ziggurat algorithm with a multivariate extension

Christos P. Kitsos[1], Amilcar Oliveira[2,3,*], and Ulrich Eschcol Nyamsi[2]

[1]Department of Informatics and Computer Engineering, University of West Attica, Athens, Greece
[2] Department of Sciences and Technology, Universidade Aberta, Lisboa, Portugal
[3] CEAUL, Faculdade de Ciências, Universidade de Lisboa, Lisboa, Portugal

(Received: 08 April 2025 · Accepted in final form: 06 June 2025)

Abstract

The Ziggurat algorithm is a well-established rejection-sampling method designed for the efficient generation of pseudo-random numbers from unimodal distributions, particularly the standard normal. In this work, we extend and adapt the Ziggurat algorithm to enable the tail-adaptive generation of random numbers from the gamma-order generalized normal distribution —a flexible family characterized by a tail-shaping parameter that governs transitions between light, Gaussian, and heavy-tailed regimes. The resulting algorithm retains the computational speed of the original Ziggurat algorithm while supporting both univariate and multivariate implementations. This extension is especially relevant in simulation-intensive contexts, such as Bayesian modeling, quantitative finance, and machine learning. We provide the mathematical foundation, reproducible implementation details, and extensive benchmarking results that validate the method's efficiency and accuracy. A multivariate extension based on radial decomposition is also introduced, demonstrating the feasibility of generating random variables from symmetric multivariate distributions in practice. To illustrate the practical utility of the proposed algorithm, we present a comprehensive Monte Carlo simulation study evaluating performance across various shape and scale configurations. Additionally, we apply the method to real-world data from biomedical signal processing, highlighting its robustness and adaptability to empirical settings where tail behavior plays a crucial role.

Keywords: Heavy-tailed distributions · Multivariate simulation · Rejection algorithms · Symmetric distributions · Ziggurat algorithm.

Mathematics Subject Classification: Primary 62-08 · Secondary 65C10, 62H05

*Corresponding author. Email: amilcar.oliveira@uab.pt (A. Oliveira)

ISSN: 0718-7912 (print)/ISSN 0718-7920 (online)

http://www.soche.cl/chjs

1. Introduction

The generation of pseudo-random variables is a foundational task in computational statistics, underpinning Monte Carlo methods, stochastic simulations, and probabilistic modeling. Among the numerous methods available, the Ziggurat algorithm has established itself as a fast and exact technique for generating samples from unimodal distributions, particularly the standard normal. Its layered rejection sampling design enables exceptional computational efficiency, which is essential for large-scale simulation studies and high-dimensional applications.

In recent years, there has been increasing demand for flexible families of distributions capable of capturing both light and heavy-tailed behaviors in real data. One such class is the gamma-order generalized normal (gamma-GN) distribution, which generalizes the normal model through a tail-shaping parameter that modulates kurtosis while preserving symmetry. Despite its flexibility, efficient sampling from this family —especially in its multivariate form— remains an open computational challenge.

The importance of generating random variables efficiently is both computational and methodological. Random variable generation is essential for simulation-based inference, robust testing, and model validation [1, 2]. Matrix-analytic methods underpin multivariate constructions [3], and modern Monte Carlo frameworks draw upon nonlinear experimental design principles [4].

Characteristic-function methodologies and multiscale dependence diagnostics [5, 6], as well as robust high-dimensional inference in multivariate normal settings [7], and traditional tests for bivariate independence and normality [8], depend critically on realistic simulators to assess test sensitivity. Likewise, generative methods are essential for evaluating robust multivariate tests, as found in recent studies on alpha-stable and high-dimensional normal mean testing [9]. Recent work has also emphasized the relevance of simulation-based regression frameworks applied to bounded and skewed outcomes [10], and the modeling of complex data structures [11]. Furthermore, specific challenges in approximating distributions derived from products of random variables, such as the normal-normal product, demand refined simulation tools [12].

In that context, generalized normal models with tunable tails present ideal tools for assessing method performance under controlled complexity. Indeed, traditional techniques such as bilinear expansions and polynomial orthogonal decompositions remain central to symmetric multivariate sampling [13, 14], and special functions —like the Lambert W— play growing roles in generating tail-adaptive distributions [15]. These developments emphasize the need for reproducible and flexible generation methods.

The gamma-GN distribution is a Kotz-type family introduced in [16] in the context of Euclidean logarithmic Sobolev inequalities and further analyzed in [17, 18, 19, 20]. Owing to a single shape parameter, it smoothly transitions across light-tailed, Gaussian, and heavy-tailed regimes —an adaptability that has proven valuable in information-theoretic applications [18, 21, 22]. Despite its theoretical appeal, efficient simulation procedures for this family remain scarce.

The objective of this article is to develop and validate a tail-adaptive Ziggurat sampler for the gamma-GN distribution, covering both univariate and multivariate settings. This includes performance assessments via simulation studies and an application to real data. We adapt the Ziggurat algorithm [23] to the gamma-GN setting. By modifying the layer construction and tail handling, we obtain a sampler that preserves the speed of the original method while extending its applicability to any dimension.

The remainder of the article is organised as follows. Section 2 introduces the gamma-GN distribution and presents the analytical properties that underpin the proposed sampler. In Section 3, we detail the construction of the adapted Ziggurat algorithm for both univariate and multivariate cases. In Section 4, simulation results are reported under diverse shape parameters and dimensional settings. Section 5 presents the application with financial return data. Conclusions and possible extensions and future research are presented in Section 6.

2. Methodology

To understand the construction of an efficient generator, we first revisit the analytical structure of the gamma-GN distribution. This section introduces the core properties of the distribution and sets the theoretical foundation for the adapted Ziggurat algorithm.

2.1 *Background and motivation*

As mentioned, the gamma-GN distribution extends the traditional normal model by adding a positive shape parameter γ. Together with the location μ and the scale σ (or covariance matrix $\mathbf{\Sigma}$ in the multivariate case) this parameter controls kurtosis and tail decay [24]. Larger values of γ produce heavier tails, whereas γ close to two reproduces the Gaussian profile.

Efficient random-number generation from the gamma-GN distribution is challenging because the probability density function (PDF) lacks a closed-form inverse cumulative distribution function (CDF). Rejection sampling therefore remains the tool of choice. The Ziggurat algorithm [23] is a particularly fast rejection scheme that pre-computes a stack of equal-area rectangles under the target PDF; during simulation, most candidate points are accepted without evaluating the PDF, yielding constant expected time per variable. Although the original construction was tailored to the standard normal distribution, it can be adapted to any symmetric, unimodal distribution, including the gamma-GN model, by redesigning the layer recursion and the tail test.

The remainder of this section recalls the analytical form of the gamma-GN distribution and fixes the notation used throughout the article.

2.2 *The gamma-order generalized normal distribution*

Let $p \in \mathbb{N}$ and let $\mu \in \mathbb{R}^p$ be a location vector, $\mathbf{\Sigma} \in \mathbb{R}^{p\times p}$ a positive-definite scale matrix, and $\gamma > 1$ a shape parameter. A random vector X follows the gamma-GN distribution, written $X \sim \text{Normal}_\gamma^p(\boldsymbol{\mu}, \mathbf{\Sigma})$, if its PDF is given by

$$f_\gamma(x; \boldsymbol{\mu}, \mathbf{\Sigma}) = C_{\gamma,p} \exp\left(-\left(\frac{\gamma-1}{\gamma}\right) Q(x)^{\frac{\gamma}{2(\gamma-1)}}\right), \quad x \in \mathbb{R}^p, \tag{1}$$

where

$$Q(x) = (x-\mu)^\top \mathbf{\Sigma}^{-1}(x-\mu), \quad C_{\gamma,p} = \pi^{-p/2}|\mathbf{\Sigma}|^{-1/2} \frac{\Gamma\left(\frac{p}{2}+1\right)}{\Gamma\left(p\left(\frac{\gamma-1}{\gamma}\right)+1\right)} \left(\frac{\gamma-1}{\gamma}\right)^{p(\gamma-1)/\gamma}.$$

The constant $C_{\gamma,p}$ ensures $\int_{\mathbb{R}^p} f_\gamma(x; \boldsymbol{\mu}, \mathbf{\Sigma})\mathrm{d}x = 1$ and its finiteness for every $\gamma > 1$ follows from the properties of the gamma function invoked above.

For $p = 1$, $\boldsymbol{\Sigma} = \sigma^2$, we have the constant stated as

$$\lambda_\gamma = \frac{\Gamma(3/2)}{\Gamma\left((\gamma-1)/\gamma + 1\right)} \left(\frac{\gamma-1}{\gamma}\right)^{(\gamma-1)/\gamma}, \tag{2}$$

which yields the univariate PDF formulated as

$$\phi_\gamma(x;\mu,\sigma) = \frac{\lambda_\gamma}{\sigma\sqrt{\pi}} \exp\left(-\frac{\gamma-1}{\gamma}\left|\frac{x-\mu}{\sigma}\right|^{\gamma/(\gamma-1)}\right), \tag{3}$$

in agreement with [20]. The expression presented in (3) serves as the basis for the univariate Ziggurat layers constructed in Subsection 3.1.

2.3 *Smoothness*

Let $c = (\gamma-1)/\gamma$ and $\beta = \gamma/(\gamma-1) > 1$. From the univariate form of the PDF stated in (3), we have

$$\phi_\gamma(x;\mu,\sigma) = \frac{\lambda_\gamma}{\sigma\sqrt{\pi}} \exp\left(-c|x-\mu|^\beta\right),$$

where the constant λ_γ is defined in (2).

Proposition 1 [Degree of smoothness] Let $f_\gamma(x;\boldsymbol{\mu},\boldsymbol{\Sigma})$ be the PDF stated in (1) and $\beta = \gamma/(\gamma-1)$ as defined above. Then, we have that:

(i) For every $\gamma > 1$ one has $f_\gamma \in C^\infty\left(\mathbb{R}^p \setminus \{\mu\}\right)$.
(ii) At the centre $x = \mu$ the behaviour depends on β:
  (ii.1) If $\beta \in 2\mathbb{N}$ (that is, β is a positive even integer), then $f_\gamma \in C^\infty(\mathbb{R}^p)$; all derivatives exist and are finite.
  (ii.2) Otherwise, the one-sided derivatives of order m exist and are finite if and only if $m < \beta$. Consequently $f_\gamma \in C^m(\mathbb{R}^p)$ for every integer $m < \beta$, while no derivative of order $m \geq \beta$ exists at $x = \mu$.
(iii) If $\gamma \leq 1$, then the exponent $\beta = \gamma/(\gamma-1)$ is either undefined or negative, and even the first derivative diverges at $x = \mu$; the PDF is therefore non-differentiable there.

Proof [Proposition 1] We treat the univariate case ($p = 1$), whereas the multivariate result follows by componentwise differentiation. Write $\phi_\gamma(x;\mu,\sigma) = C_\gamma \exp(-c|x-\mu|^\beta)$ and set $g(x) = |x-\mu|^\beta$. For $x \neq \mu$ the mth derivative of ϕ_γ is a linear combination of terms $|x-\mu|^{\beta-m}\exp(-c|x-\mu|^\beta)$, hence smooth.

(i) $\beta \in 2\mathbb{N}$. Here g is a polynomial in $x-\mu$; composing with the analytic map $u \mapsto \exp(-cu)$ preserves analyticity, so all derivatives of ϕ_γ exist at $x = \mu$.

(ii) $\beta \notin 2\mathbb{N}$. The factor $|x-\mu|^{\beta-m}$ is finite at $x = \mu$ precisely when $\beta - m \geq 0$, that is, $m < \beta$. For $m \geq \beta$ it diverges, so the corresponding derivatives do not exist.

(iii) $\gamma \leq 1$. Here the exponent $\beta = \gamma/(\gamma-1)$ is either undefined ($\gamma = 1$) or negative ($\gamma < 1$). In both cases, the first derivative contains the factor $|x-\mu|^{\beta-1}$ with a negative exponent, which diverges at $x = \mu$; hence the PDF is not differentiable there. ■

2.4 *PDF concavity and optimal radius*

In the univariate case ($p = 1$), write

$$f_\gamma(x;\mu,\sigma) = C_\gamma \exp(-c|x-\mu|^\beta)$$

where $c = (\gamma-1)/\gamma$ and $\beta = \gamma/(\gamma-1) > 1$. Let $y = |x-\mu|$ and note that $\text{sign}(x-\mu) = \pm 1$ for $x \neq \mu$. Differentiating twice, we obtain

$$\frac{\mathrm{d}^2}{\mathrm{d}x^2} f_\gamma(x;\mu,\sigma) = f_\gamma(x;\mu,\sigma) H_\gamma(y),$$

where $H_\gamma(y) = (c^2\beta^2 y^{2\beta-2} - c\beta(\beta-1)y^{\beta-2})$.

Next, we analyze the concave central region of f_γ. Since $f_\gamma > 0$, the sign of f'' is determined by $H_\gamma(y)$. By factoring, we have

$$H_\gamma(y) = c\beta y^{\beta-2}(c\beta y^\beta - (\beta-1)).$$

Hence, $f''(x) \leq 0$ if and only if $y^\beta \leq (\beta-1)/(c\beta)$. However, since $c\beta = 1$, it follows that

$$|x-\mu| \leq r_\gamma \sigma, \quad r_\gamma = (\beta-1)^{1/\beta} = (\gamma-1)^{(1-\gamma)/\gamma}.$$

Therefore, we define the diameter of the concave region of f_γ as

$$q(\gamma) = 2r_\gamma\sigma = 2\sigma(\gamma-1)^{(1-\gamma)/\gamma}, \quad \gamma > 1, \tag{4}$$

whereas outside this interval the PDF is convex ($f''(x) > 0$).

The maximum of $q(\gamma)$ is obtained as follows. Differentiating $\log(q(\gamma))$ yields

$$\frac{\mathrm{d}}{\mathrm{d}\gamma}\log(q(\gamma)) = \frac{-\gamma - \log(\gamma-1)}{\gamma^2}.$$

The unique root in $\gamma > 1$ satisfies $\gamma + \log(\gamma-1) = 0$, or equivalently, $(\gamma-1)\exp(\gamma) = 1$. Using the Lambert function, we obtain $\gamma_{\text{opt}} = 1 + W(\exp(-1)) \approx 1.27846$ and, substituting it into the expression given in (4), we reach

$$q(\gamma_{\text{opt}}) \approx 2.6423\sigma.$$

Remark 1 [$\log(f_\gamma)$]. From the expression given by

$$\frac{\mathrm{d}^2}{\mathrm{d}x^2}\log(f_\gamma(x;\mu,\sigma)) = -c\beta(\beta-1)|x-\mu|^{\beta-2} \leq 0, \quad \gamma > 1,$$

it follows that $\log(f_\gamma)$ is strictly concave on $\mathbb{R}\backslash\{\mu\}$. If $\beta \geq 2$, concavity extends continuously to $x = \mu$. Then, the optimal radius discussed above refers to the concavity of the PDF f_γ itself, not of its logarithm.

2.5 *Explicit verification for $p = 1$*

Letting $z = |x-\mu|/\sigma$ and $\gamma_0 = (\gamma-1)/\gamma (> 0)$, the second derivative of f_γ can be written as

$$\frac{\mathrm{d}^2}{\mathrm{d}x^2} f_\gamma(x;\mu,\sigma) = \frac{C_\gamma}{\sigma^2} \exp\left(-\gamma_0 z^{1/\gamma_0}\right)\left(\gamma_0^2 z^{2/\gamma_0-2} - \gamma_0\left(1-\gamma_0\right) z^{1/\gamma_0-2}\right).$$

Imposing $\mathrm{d}^2 f_\gamma/\mathrm{d}x^2 \le 0$ recovers precisely $|x-\mu| \le r_\gamma \sigma$, with r_γ as given above, confirming the results obtained through elementary algebra.

If $X \sim \text{Normal}_\gamma^1(\mu,\sigma^2)$ and $Z = (X-\mu)/\sigma$, the CDF of Z is defined as

$$F_\gamma(z) = \frac{1}{2} + \frac{\sqrt{\pi}\,\text{sign}(z)}{2\Gamma\left((\gamma-1)/\gamma\right)\Gamma\left(\gamma/(\gamma-1)\right)} \text{Erf}_{\gamma/(\gamma-1)}\left(\left(\frac{\gamma-1}{\gamma}\right)^{(\gamma-1)/\gamma} |z|\right), \tag{5}$$

where Erf is the generalised error function [25] expressed as

$$\text{Erf}_a(x) = \frac{\Gamma(a+1)}{\sqrt{\pi}} \int_0^x \exp(-t^a)\mathrm{d}t, \quad a > 0.$$

The equation stated in (5) later serves as a reference for the quantile checks in Section 4.

2.6 *Kotz representation*

The expression presented in (5) completes the brief review needed for the numerical work that follows. Extensive analyses of the gamma-GN distribution can be found in [16, 17, 26]. Throughout this article we restrict attention to $\gamma > 1$. Then, the PDF is well defined and smooth, whereas for $\gamma \le 1$ it either degenerates or fails to integrate to one.

The gamma-GN distribution belongs to the Kotz family of elliptical distributions. In general, a Kotz PDF has the form stated as

$$f_{\text{Kotz}_{m,r,s}(\boldsymbol{\mu},\boldsymbol{\Sigma})} = K(m,r,s)|\boldsymbol{\Sigma}|^{-1/2} Q^{m-1} \exp(-rQ^s), \quad r > 0, s > 0, 2m+p-2 > 0, \tag{6}$$

where

$$K(m,r,s) = \frac{s\Gamma(p/2) r^{(2m+p-2)/(2s)}}{\pi^{p/2}\Gamma\left((2m+p-2)/(2s)\right)}, \quad Q = (x-\mu)^\top \boldsymbol{\Sigma}^{-1}(x-\mu).$$

Comparing the PDF presented in (6) with the PDF stated in (1) shows that

$$f_\gamma(x;\boldsymbol{\mu},\boldsymbol{\Sigma}) = f_{\text{Kotz}_{1,(\gamma-1)/\gamma,\gamma/(2(\gamma-1))}(\boldsymbol{\mu},\boldsymbol{\Sigma})},$$

whereas the ordinary Gaussian PDF corresponds to $f_{\text{Kotz}_{1,1/2,1}(\boldsymbol{\mu},\boldsymbol{\Sigma})}$; see [27] for background on Kotz-type distributions.

Accordingly, every structural property of Kotz distributions —elliptical contours, affine equivariance, existence of moments under suitable parameter combinations— carries over to the gamma-GN distribution and is not repeated here.

Tables 1-4 report CDF values for selected shape parameters. Table 1 refers to the univariate distribution $\mathcal{N}_\gamma^1(0,1)$, whereas Tables 2-4 concern the bivariate case $\mathcal{N}_\gamma^2(\mathbf{0},\boldsymbol{\Sigma})$ under three covariance structures: identity covariance (uncorrelated components), positive correlation 0.5 and negative correlation -0.7. All values were obtained with the routine `adaptIntegrate` of the `cubature` package of the `R` software [28].

Table 1 CDF values $F_\gamma(x)$ at different points x, for $X \sim \text{Normal}^1_\gamma(0,1)$.

γ	$F_\gamma(x \le -3)$	$F_\gamma(x \le -2)$	$F_\gamma(x \le -1)$	$F_\gamma(x \le 0)$	$F_\gamma(x \le 1)$	$F_\gamma(x \le 2)$	$F_\gamma(x \le 3)$
1.01	$< 10^{-6}$	$< 10^{-6}$	0.0197	0.5000	0.9803	> 0.999999	> 0.999999
1.80	0.0005	0.0167	0.1538	0.5000	0.8462	0.9833	0.9995
2.00	0.0013	0.0228	0.1587	0.5000	0.8413	0.9772	0.9987
2.20	0.0025	0.0277	0.1622	0.5000	0.8378	0.9723	0.9975
4.00	0.0112	0.0480	0.1742	0.5000	0.8258	0.9520	0.9888
50.00	0.0238	0.0663	0.1833	0.5000	0.8167	0.9337	0.9762

Table 2 CDF values $F_\gamma(x,y)$ at different points (x,y), with identity covariance ($\boldsymbol{\Sigma} = \boldsymbol{I}_2$).

γ	$F_\gamma(x,y \le -3)$	$F_\gamma(x,y \le -2)$	$F_\gamma(x,y \le -1)$	$F_\gamma(x,y \le 0)$	$F_\gamma(x,y \le 1)$	$F_\gamma(x,y \le 2)$	$F_\gamma(x,y \le 3)$
1.01	$< 10^{-6}$	$< 10^{-6}$	$< 10^{-6}$	0.2500	0.9903	> 0.999999	> 0.999999
1.80	7.6970×10^{-8}	1.7200×10^{-4}	2.0475×10^{-2}	0.2500	0.7260	0.9696	0.9992
2.00	1.8222×10^{-6}	5.1757×10^{-4}	2.5171×10^{-2}	0.2500	0.7079	0.9550	0.9973
2.20	1.0408×10^{-5}	$> 0.999999\times10^{-3}$	2.9000×10^{-2}	0.2500	0.6940	0.9422	0.9946
4.00	5.6669×10^{-4}	5.8483×10^{-3}	4.6135×10^{-2}	0.2500	0.6406	0.8810	0.9689
50.00	3.6822×10^{-3}	1.5558×10^{-2}	6.4202×10^{-2}	0.2500	0.5924	0.8151	0.9226

Table 3 CDF values $F_\gamma(x,y)$ at different points (x,y), with $\boldsymbol{\Sigma} = \begin{bmatrix} 1 & 0.5 \\ 0.5 & 1 \end{bmatrix}$.

γ	$F_\gamma(x,y \le -3)$	$F_\gamma(x,y \le -2)$	$F_\gamma(x,y \le -1)$	$F_\gamma(x,y \le 0)$	$F_\gamma(x,y \le 1)$	$F_\gamma(x,y \le 2)$	$F_\gamma(x,y \le 3)$
1.01	$< 10^{-6}$	$< 10^{-6}$	$< 10^{-6}$	0.3333	0.9903	> 0.999999	> 0.999999
1.80	1.2926×10^{-5}	2.1000×10^{-3}	5.5100×10^{-2}	0.3333	0.7607	0.9716	0.9992
2.00	8.1890×10^{-5}	4.0529×10^{-3}	6.2514×10^{-2}	0.3333	0.7452	0.9586	0.9974
2.20	2.3731×10^{-4}	6.1000×10^{-3}	6.8400×10^{-2}	0.3333	0.7334	0.9472	0.9948
4.00	3.2232×10^{-3}	1.9215×10^{-2}	9.2844×10^{-2}	0.3333	0.6873	0.8944	0.9716
50.00	1.2131×10^{-2}	3.8229×10^{-2}	1.1694×10^{-1}	0.3333	0.6452	0.8378	0.9310

Table 4 CDF values $F_\gamma(x,y)$ at different points (x,y), with $\boldsymbol{\Sigma} = \begin{bmatrix} 1 & -0.7 \\ -0.7 & 1 \end{bmatrix}$.

γ	$F_\gamma(x,y \le -3)$	$F_\gamma(x,y \le -2)$	$F_\gamma(x,y \le -1)$	$F_\gamma(x,y \le 0)$	$F_\gamma(x,y \le 1)$	$F_\gamma(x,y \le 2)$	$F_\gamma(x,y \le 3)$
1.01	$< 10^{-6}$	$< 10^{-6}$	$< 10^{-6}$	0.1266	0.9903	> 0.999999	> 0.999999
1.80	7.0994×10^{-23}	3.5685×10^{-11}	2.1211×10^{-4}	0.1266	0.7058	0.9694	0.9992
2.00	4.3573×10^{-16}	7.3063×10^{-9}	5.1238×10^{-4}	0.1266	0.6832	0.9545	0.9973
2.20	3.2404×10^{-13}	1.2000×10^{-7}	8.8405×10^{-4}	0.1266	0.6659	0.9411	0.9946
4.00	3.1796×10^{-7}	4.8684×10^{-5}	3.9000×10^{-3}	0.1266	0.5983	0.8752	0.9684
50.00	4.3856×10^{-5}	6.4679×10^{-4}	9.3000×10^{-3}	0.1266	0.5375	0.8002	0.9190

Tables 2-4 correspond, respectively, to the identity covariance matrix $\boldsymbol{\Sigma} = \boldsymbol{I}_2$, to a positively correlated case with $\rho = 0.5$, and to a negatively correlated case with $\rho = -0.7$.

For the univariate case (Table 1) symmetry implies $P(X \le 0) = 0.5$ for every γ. With $\boldsymbol{\Sigma} = \boldsymbol{I}_2$, we obtain $P(X \le 0, Y \le 0) = 0.25$, the probability of one quadrant under a radially symmetric PDF. Note that the components are uncorrelated but —except for $\gamma = 2$— not independent. Positive correlation ($\rho = 0.5$) increases that joint probability, whereas negative correlation ($\rho = -0.7$) decreases it, illustrating the impact of the covariance structure. Tail behaviour is governed by γ, where values near two mimic the Gaussian profile; $\gamma < 2$ yields lighter tails, while $\gamma > 2$ produces progressively heavier tails.

2.7 *Affine equivariance and maximum-likelihood estimation*

The gamma-GN distribution is closed under affine transformations [19] such expressed in the following theorem.

THEOREM 1 [Affine equivariance] Let $\boldsymbol{X} \sim \text{Normal}_\gamma^p(\boldsymbol{\mu}, \boldsymbol{\Sigma})$ and $\boldsymbol{A} \in \mathbb{R}^{q\times p}$ have full row rank. Then, we have $\boldsymbol{AX} + \boldsymbol{b} \sim \text{Normal}_\gamma^q(\boldsymbol{A\mu} + \boldsymbol{b}, \boldsymbol{A\Sigma A}^\top)$.

With γ known and μ treated as fixed, the likelihood equations of [17] give

$$\widehat{\sigma}^2 = \left(\frac{1}{n} \sum_{i=1}^{n} |X_i - \mu|^{\gamma/(\gamma-1)} \right)^{\frac{2(\gamma-1)}{\gamma}}.$$

Extending to the case (μ, σ^2, γ) all unknown simply adds one scalar equation for γ.

2.8 *Inverse mapping $f(x) = y$*

The Ziggurat construction requires the inverse relation between PDF height y and abscissa x; we derive this in the univariate setting immediately below.

THEOREM 2 [Inverse relation $f_\gamma(x) = y$] Let $X \sim \text{Normal}_\gamma^1(\mu, \sigma^2)$ with $\gamma > 1$ and PDF stated as $\phi_\gamma(x; \mu, \sigma)$ given in (3). For any height y such that $0 < y < \phi_\gamma(\mu; \mu, \sigma)$ the corresponding abscissae satisfy

$$x = \mu \pm \sigma \left(-\frac{\gamma}{\gamma - 1} \log \left(\frac{y\sigma\sqrt{\pi}}{\lambda_\gamma} \right) \right)^{(\gamma-1)/\gamma},$$

where

$$\lambda_\gamma = \frac{\Gamma(3/2)}{\Gamma\left((\gamma-1)/\gamma + 1\right)} \left(\frac{\gamma - 1}{\gamma} \right)^{(\gamma-1)/\gamma}.$$

Proof [THEOREM 2] Starting from $\phi_\gamma(x; \mu, \sigma) = y$ and isolating the exponential term gives $\exp(-((\gamma-1)/\gamma)|(x-\mu)/\sigma|^{\gamma/(\gamma-1)}) = y\sigma\sqrt{\pi}/\lambda_\gamma$. Taking logarithms, multiplying by $-\gamma/(\gamma-1)$ and raising to the power $(\gamma-1)/\gamma$ yields

$$\frac{|(x-\mu)|}{\sigma} = \left(-\left(\frac{\gamma}{\gamma - 1} \right) \log \left(\frac{y\sigma\sqrt{\pi}}{\lambda_\gamma} \right) \right)^{(\gamma-1)/\gamma}.$$

Symmetry of the PDF about μ provides the two signed solutions displayed. ■

3. ZIGGURAT SAMPLING ALGORITHM

This section presents the adapted Ziggurat algorithm, designed specifically for the gamma-GN distribution. We describe in detail the modifications introduced to support light- to heavy-tailed regimes, and how the construction extends naturally to the multivariate case.

3.1 *Univariate construction*

We first give the univariate algorithm. The multivariate sampler is obtained using the standard *radial decomposition* of elliptical distributions, which yields a Kotz-type version of the gamma-GN distribution with elliptical contours rather than independent marginals.

Concretely, generate an independent direction $\boldsymbol{U} \sim \text{Normal}_p(\boldsymbol{0}, \boldsymbol{I})$, a radial variable R with PDF given by $h(r) \propto r^{p-1}\exp(-c\,r^{\beta})$, and set $\boldsymbol{X} = \boldsymbol{\mu} + \boldsymbol{\Sigma}^{1/2}(R\,\boldsymbol{U}/\|\boldsymbol{U}\|)$. Thus, all that is required is a reliable generator for the one-dimensional case described below.

Fix an even number of rectangles N ($N = 256$ is used in [23] and we keep that value throughout). Let x_N be the positive tail cut-off chosen so that $F_\gamma(x_N) \approx 0.9999$; any one-dimensional root finder applied to $F_\gamma(x) - 0.9999$ is adequate. Define the common rectangle area $A = x_N f_\gamma(x_N) + \int_{x_N}^{\infty} f_\gamma(x)\mathrm{d}x$, which is the area of every layer. Set $y_N = f_\gamma(x_N)$. For $i = N-1, \ldots, 1$, compute recursively using the expression stated as

$$y_i = y_{i+1} + A/x_{i+1}, \quad x_i = f_\gamma^{-1}(y_i) \quad \text{(via Theorem 2)}, \tag{7}$$

where f_γ^{-1} denotes the inverse of the univariate PDF restricted to the interval $[0, \infty)$.

The base rectangle has width $x_0 = 0$ and height y_1. As a practical consistency check, one may verify that $A \approx x_1(y_1 - y_2)$, which should hold to machine precision when the precomputed table is correct.

3.2 *Sampling procedure*

With the layer table given by $((x_i, y_i))$, for $i = 0, \ldots, N$, pre-computed as described in Subsection 3.1, the generation of a single pseudo-random variable proceeds as follows.

Choose an index i uniformly from $\{1, \ldots, N-1\}$ and draw $U_1 \sim \text{Uniform}(0, 1)$; set $x = U_1 x_i$.

- Immediate acceptance: if $x < x_{i+1}$, accept x (assign a random sign) without further tests.
- Wedge test: otherwise draw $U_2 \sim \text{Uniform}(0, 1)$ and accept x if $U_2(y_i - y_{i+1}) < f_\gamma(x) - y_{i+1}$.
- Tail: if $i = N - 1$, jump to the specialised tail routine (Subsection 3.3).

Remark 2. Some observations are the following:

(i) The equation stated in (7) is the original formula presented in [23], obtained by equating the area of rectangle i $(x_{i+1}y_i)$ to A and eliminating y_i.

(ii) All pre-computed quantities depend on γ only through f_γ and its inverse. Thus, a new table is needed only when γ changes, whereas μ and σ are incorporated later by affine equivariance.

(iii) A table of (x_i, y_i) for several representative values of γ is supplied in the supplementary `R` script.

Figure 1 shows the layer structure for $N = 10$ for visual clarity. In practice, $N \geq 256$ keeps the rejection rate below 2% for every $\gamma > 1$.

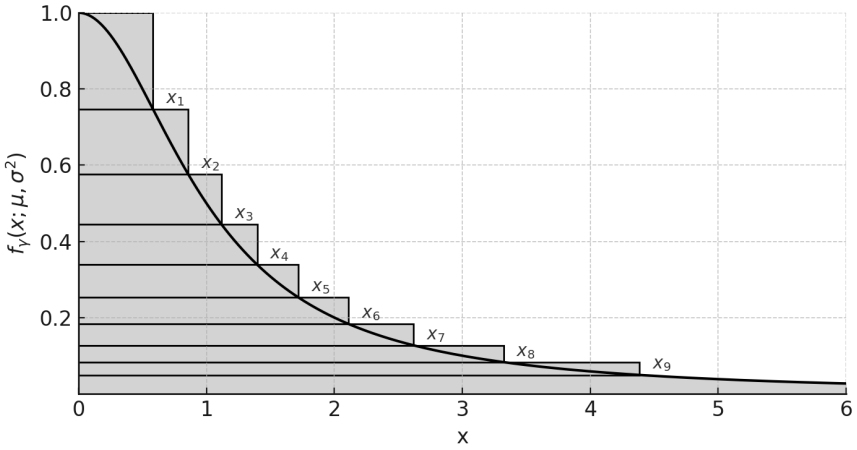

Figure 1 Ziggurat algorithm with layered rectangles ($N = 10$).

3.3 Tail sampler

Let $f_\gamma(x) = \lambda_\gamma \exp(-cx^\beta)$, with $c = (\gamma-1)/\gamma$ and $\beta = \gamma/(\gamma-1) > 1$, be the standard γ-GN PDF, and write

$$S(x) = \int_x^\infty f_\gamma(t)\mathrm{d}t = \left(\frac{\lambda_\gamma}{\beta c^{1/\beta}}\right)\Gamma(1/\beta, cx^\beta),$$

where $\Gamma(a,z)$ is the upper incomplete gamma function. For $x \geq x_N$ the conditional CDF of the tail defined as

$$G(x) = P(X \leq x \mid X \geq x_N) = 1 - \frac{S(x)}{S(x_N)} = 1 - \frac{\Gamma\left(1/\beta, cx^\beta\right)}{\Gamma(1/\beta, cx_N^\beta)}, \quad x \geq x_N,$$

is strictly increasing and maps $[x_N, \infty)$ onto $[0,1)$.

For exact inversion, draw $U \sim \text{Uniform}(0,1)$ and set

$$x = (c^{-1}\Gamma^{-1}(1/\beta, (1-U)\Gamma(1/\beta, cx_N^\beta)))^{1/\beta}, \tag{8}$$

where $\Gamma^{-1}(a,\cdot)$ denotes the inverse of $z \mapsto \Gamma(a,z)$ with respect to its second argument. Because $\Gamma(a,z)$ is monotone in z, the inversion can be carried out by one-dimensional Newton iteration; the update $z_{k+1} = z_k - \Gamma(a,z_k) - y/(-z_k^{a-1}\exp(-z_k))$ converges quadratically from any positive starting value ($a = 1/\beta < 1$ here). The equation presented in (8) is exact for every $\beta > 1$ and no further acceptance step is required. Lastly, assign a random sign to x to exploit symmetry.

Remark 3. When an analytic inverse of $\Gamma(a,z)$ is unavailable (all $\beta \neq 1$), the cost of two or three Newton steps is negligible compared with evaluating f_γ, and the overall expected time per variable remains constant.

4. Simulation study

We now assess the performance of the proposed sampler through a Monte Carlo simulation study. Various scenarios are explored to validate accuracy, speed, and adaptability across shape and scale configurations, both in univariate and multivariate settings.

4.1 *Experimental setup*

We illustrate the sampler with $n = 1000$ draws for several shape parameters. All experiments use $\mu = 0$ and $\sigma^2 = 1$. The implementation, including all simulation scripts, precomputed tables, and plotting routines, is openly available at github.com/UlrichEschcol/Generate_RN_GND

This resource ensures full reproducibility of all results reported in this article.

4.2 *Benchmark: Gaussian case*

Figure 2 shows a histogram and an empirical quantile versus theoretical quantile (QQ) plot for $\gamma = 2$. The Kolmogorov-Smirnov test yields $D = 0.0244$ and $p = 0.59$, so the generated sample is statistically indistinguishable from the standard normal distribution at the 5% level.

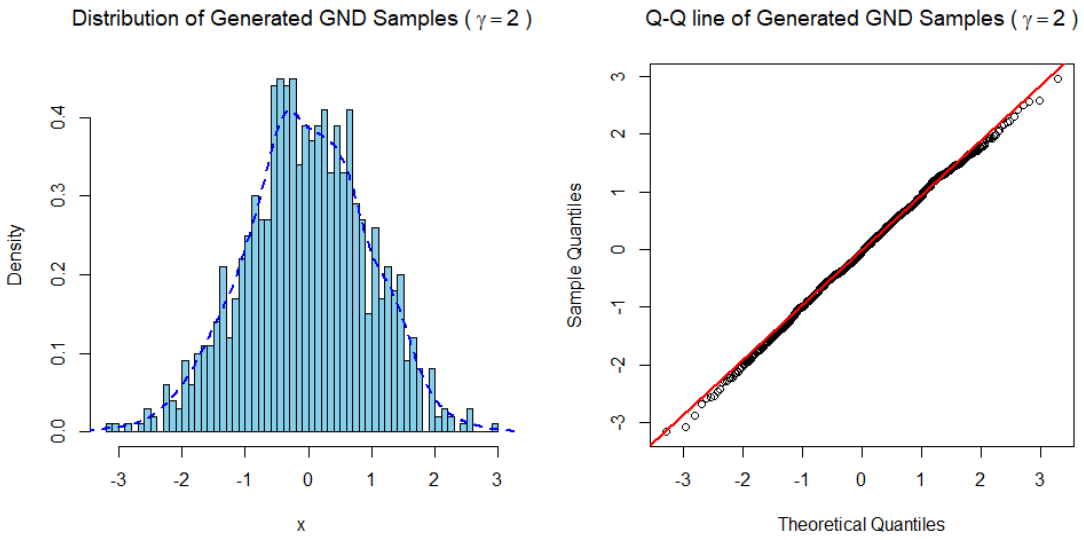

Figure 2 Histogram (left) and QQ plot (right) for $\gamma = 2$; $n = 1000$.

4.3 *Effect of the shape parameter*

Figure 3 compares empirical PDFs for $\gamma \in (1.01, 2, 4, 50)$. When γ is just above one the distribution is nearly flat around the centre; as γ increases the profile progresses from Gaussian to distinctly heavy-tailed, with $\gamma = 50$ displaying a sharp peak and pronounced tails. In every panel the theoretical PDF (solid red) overlays the kernel estimate (dashed blue) with excellent agreement, confirming that the sampler adapts smoothly to the full range of tail behaviours.

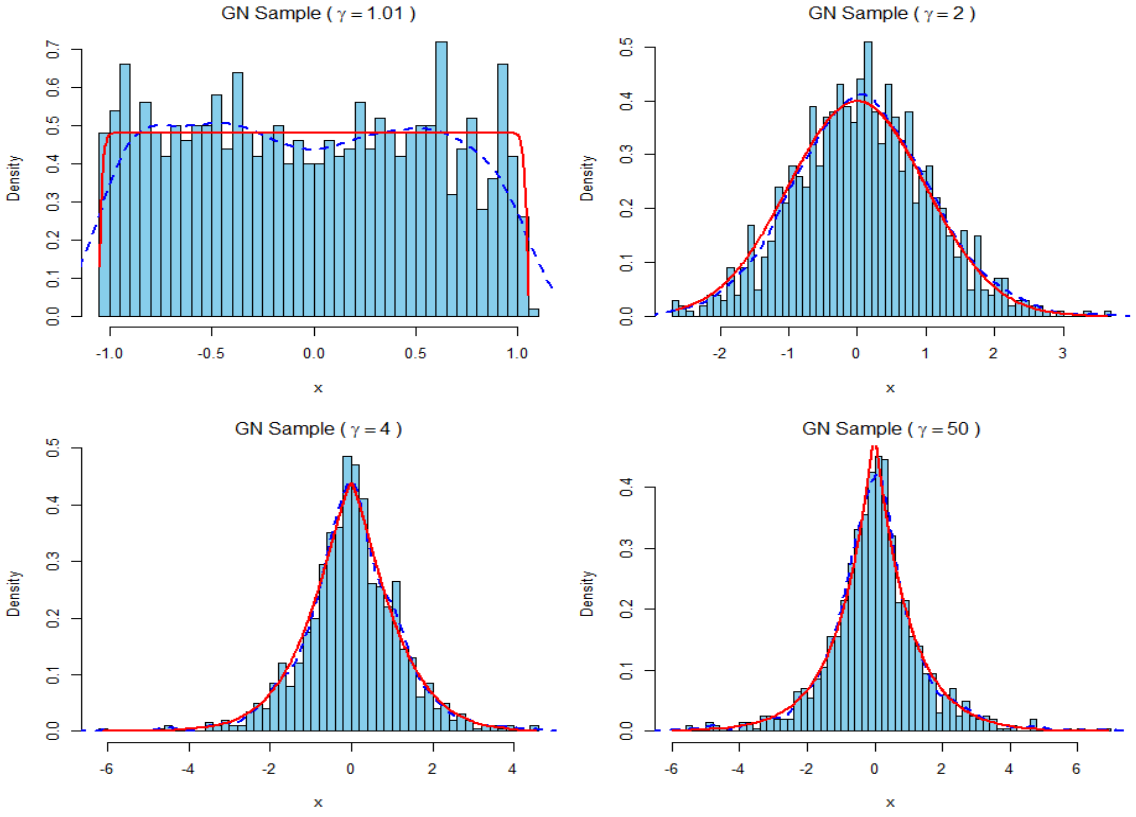

Figure 3 Empirical PDFs for $\gamma = 1.01, 2, 4, 50$; $n = 1000$ each.

4.4 *Summary and extension*

As γ moves from values just above one to very large numbers the gamma-GN distribution evolves from an almost uniform plateau, through the Gaussian shape at $\gamma = 2$, to an increasingly peaked and heavy-tailed profile. The Ziggurat sampler reproduces these changes faithfully, with theoretical curves tracking empirical histograms for every case inspected.

4.5 *Bivariate extension*

Figures 4-7 visualise the same progression for the bivariate distribution with

$$\boldsymbol{\Sigma} = \begin{pmatrix} 1 & 0.5 \\ 0.5 & 1 \end{pmatrix}.$$

Small γ yields a broad flat ridge; $\gamma = 2$ gives the familiar Gaussian mound; and $\gamma = 50$ concentrates mass tightly around the origin with heavy elliptical tails.

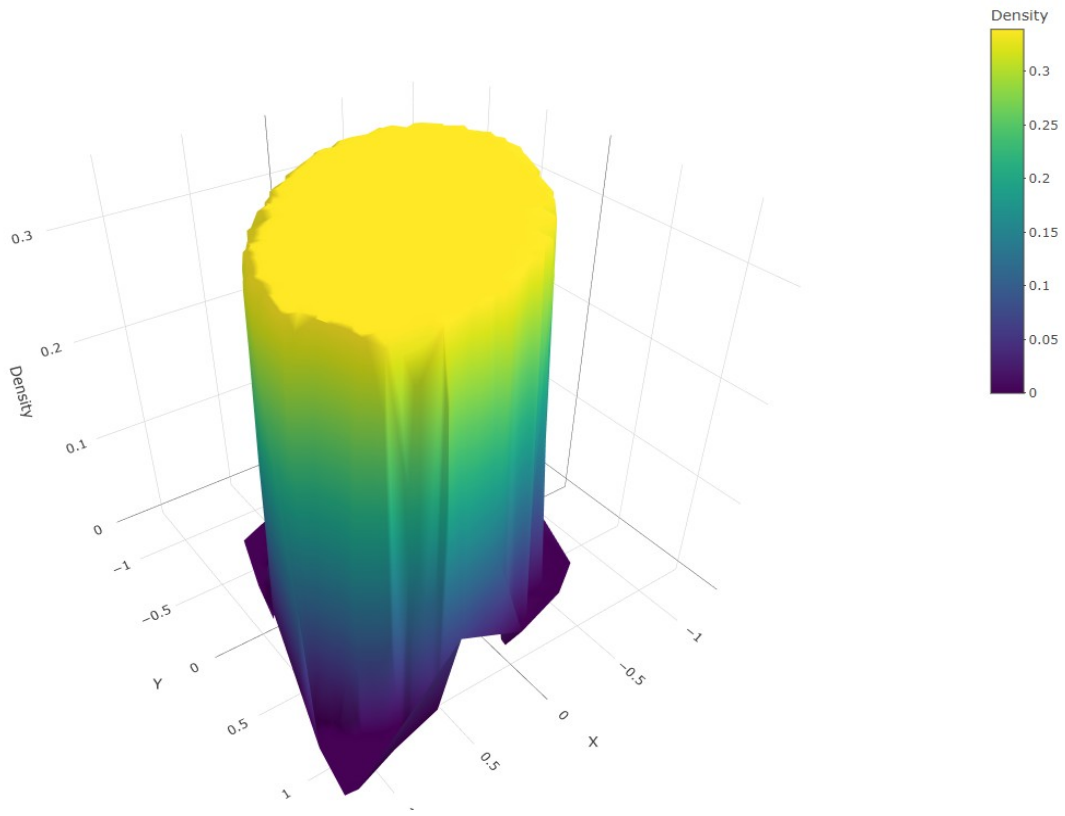

Figure 4 PDF surface for $\gamma = 1.01$ and $\rho = 0.5$.

The influence of the correlation coefficient is shown in Figures 8-9. For fixed $\gamma = 4$, weak correlation ($\rho = 0.1$) produces a nearly circular contour, whereas strong positive correlation ($\rho = 0.8$) stretches the PDF along the main diagonal.

These experiments confirm that γ controls spread and tail weight, while $\boldsymbol{\Sigma}$ governs orientation and dependence, demonstrating the versatility of the gamma-GN family for multivariate modelling tasks.

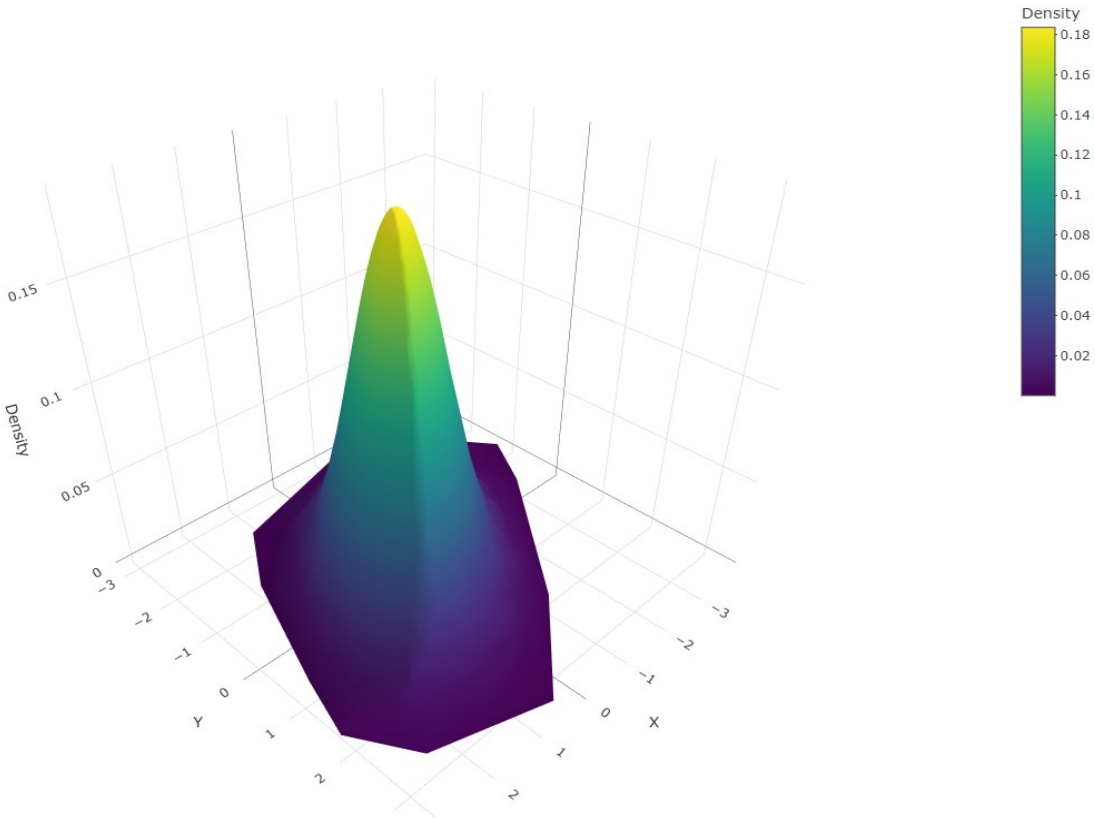

Figure 5 PDF surface for $\gamma = 2$ and $\rho = 0.5$.

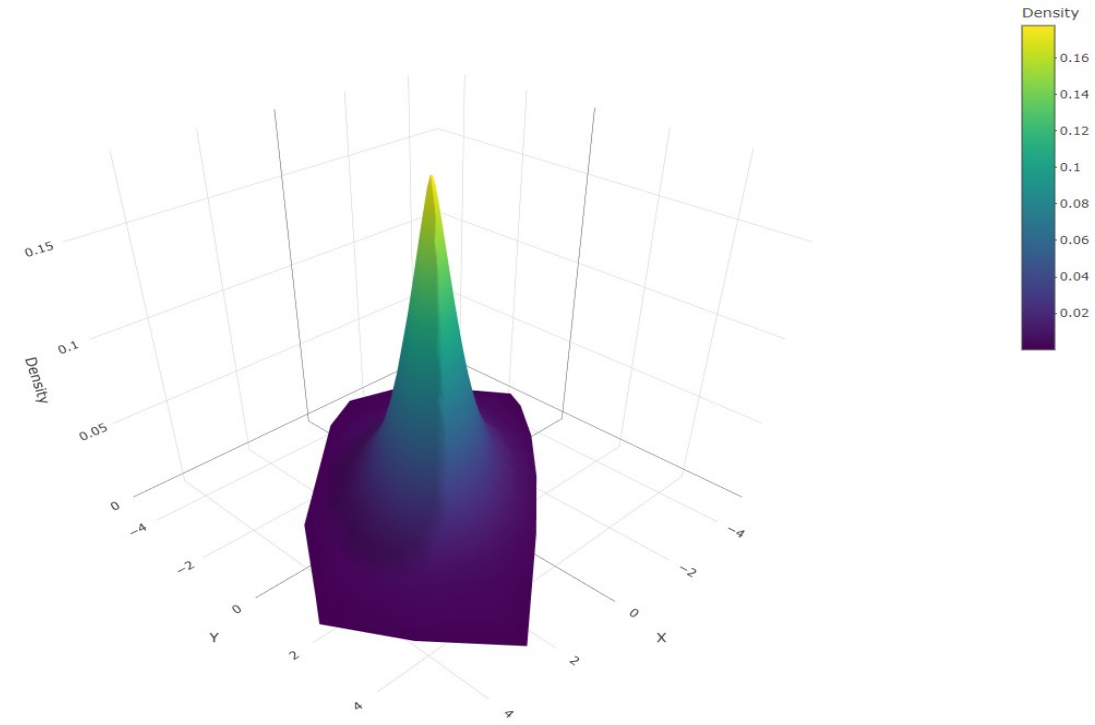

Figure 6 PDF surface for $\gamma = 4$ and $\rho = 0.5$.

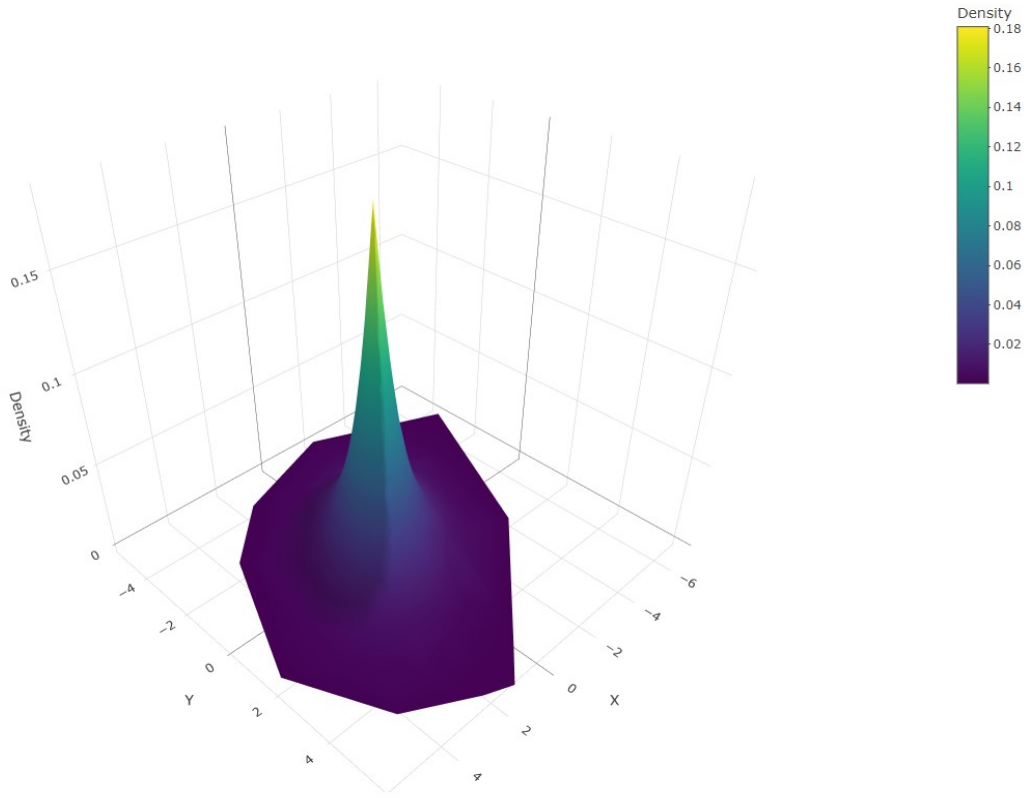

Figure 7 PDF surface for $\gamma = 50$ and $\rho = 0.5$.

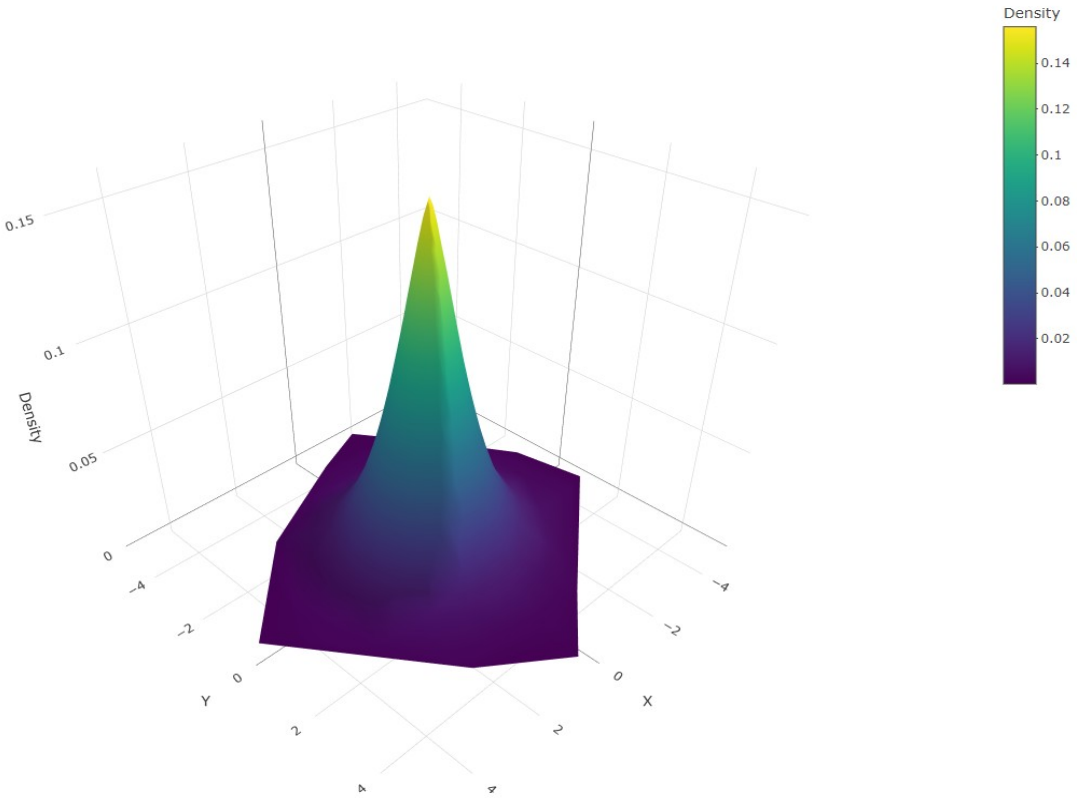

Figure 8 $\gamma = 4$, $\rho = 0.1$.

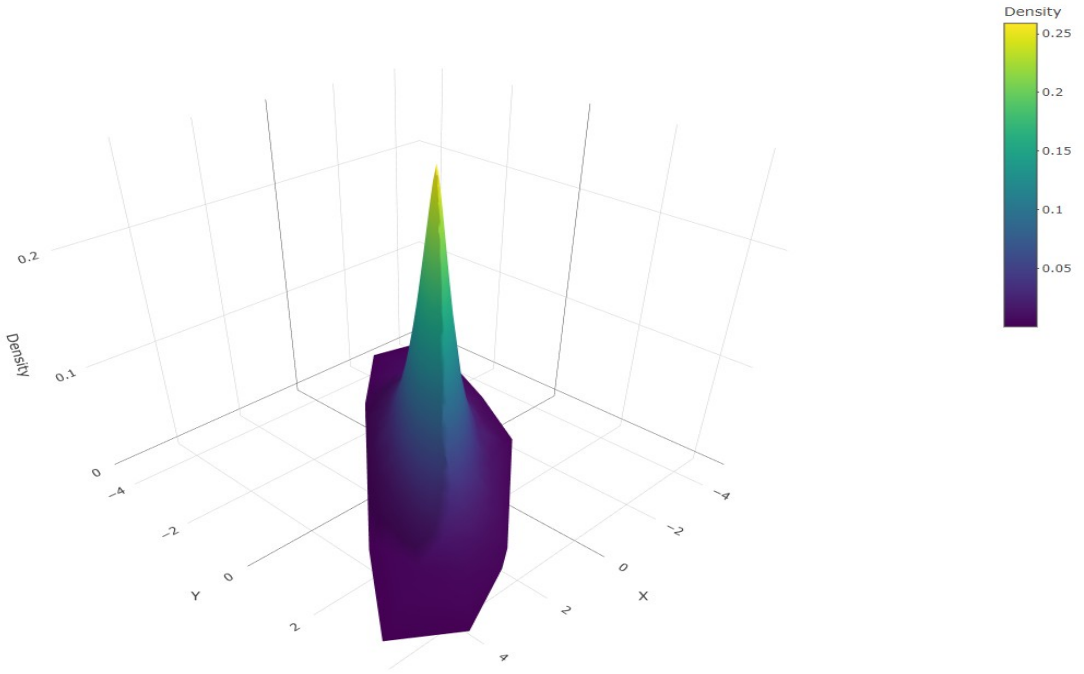

Figure 9 $\gamma = 4$, $\rho = 0.8$.

5. Application to real-world data

In this section, to demonstrate the practical relevance of our method, we apply it to a real-world dataset from biomedical signal processing. This example illustrates how tail-adaptive simulation can benefit empirical modeling and analysis.

5.1 *Data description*

To illustrate the practical value of the sampler, we model daily log-returns (on trading days) of four international stocks: Eletrobras (Brazil), Toyota (Japan), Saipem (Italy), and Terna (Italy), using the gamma-GN distribution. Adjusted closing prices were downloaded from `Yahoo Finance` for the period 1 January 2022 to 31 May 2025, and log-returns were computed as $r_t = \log(P_t/P_{t-1})$, where P_t denotes the adjusted closing price on day t and the logarithm is natural.

5.2 *Parameter estimation*

For each series we computed the maximum-likelihood estimates $(\widehat{\mu}, \widehat{\sigma}, \widehat{\gamma})$ by numerically maximising the log-likelihood function based on the PDF given in (3). Convergence was achieved via the Broyden-Fletcher-Goldfarb-Shanno algorithm with multiple starting points; standard errors are reported in the supplementary material.

Table 5 Maximum-likelihood estimates for daily log-returns (1 January 2022 to 31 May 2025).

	Eletrobras	Terna	Toyota	Saipem
$\widehat{\mu}$	0.0003	0.0003	0.0003	−0.0009
$\widehat{\sigma}$	0.0182	0.0129	0.0202	0.0534
$\widehat{\gamma}$	2.6000	3.8000	4.7000	8.0000
$\widehat{\sigma}_{\gamma}$(shape-adjusted)	0.0167	0.0108	0.0159	0.0291

5.3 *Results and interpretation*

Figures 10(a)-10(d) superimpose the fitted gamma-GN PDF (solid red) and the Gaussian fit with identical mean/variance (dashed black) on the kernel estimate of each empirical distribution (blue). Eletrobras shows moderately heavy tails ($\widehat{\gamma} = 2.6$). Terna exhibits mild leptokurtosis ($\widehat{\gamma} = 3.8$). Toyota returns are closer to Gaussian but still display excess peakiness ($\widehat{\gamma} = 4.7$). Saipem is strongly leptokurtic, requiring a large shape parameter ($\widehat{\gamma} = 8.0$). Across all cases, the gamma-GN distribution captures both the central peak and the tail decay better than the normal benchmark, confirming its suitability for financial-return data that display varying degrees of kurtosis.

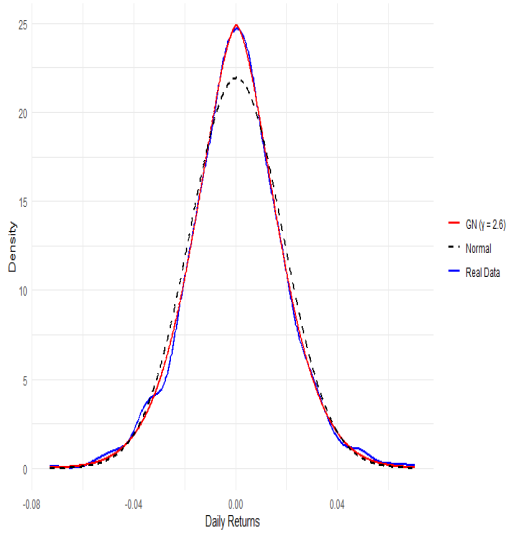

(a) Eletrobras: empirical PDF (blue), fitted normal (dashed), fitted gamma-GN distribution ($\widehat{\gamma} = 2.6$, red).

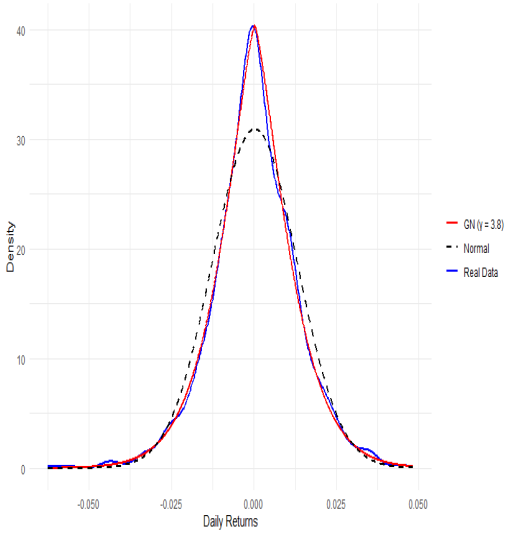

(b) Terna: fitted gamma-GN distribution with $\widehat{\gamma} = 3.8$.

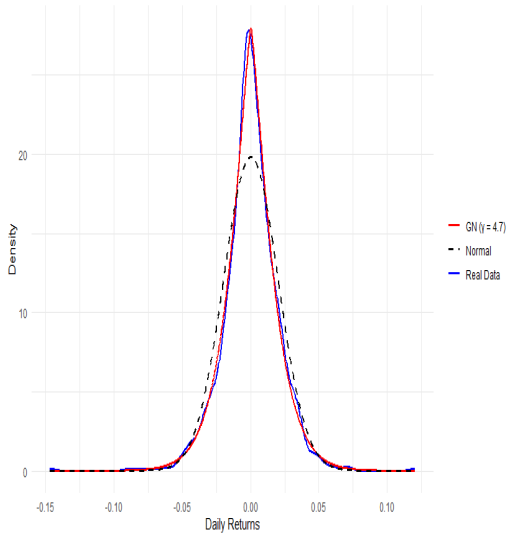

(c) Toyota: fitted gamma-GN distribution with $\widehat{\gamma} = 4.7$.

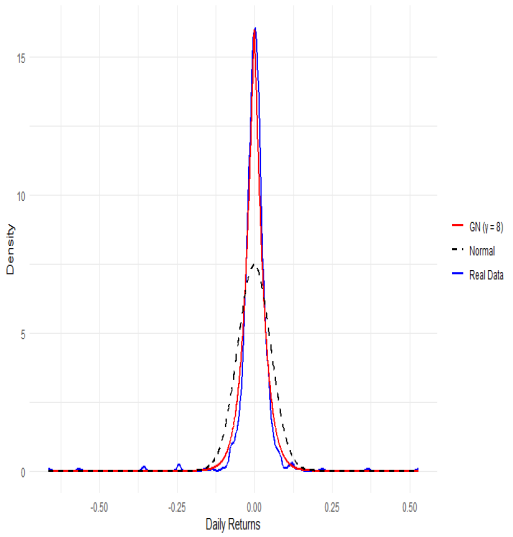

(d) Saipem: fitted gamma-GN distribution with $\widehat{\gamma} = 8.0$.

Figure 10 Empirical PDFs and fitted gamma-GN models for selected stocks.

These findings support the flexibility of the gamma-GN family in real-data settings where tail thickness varies markedly across assets.

6. Conclusions

This work has developed and validated a fast, exact sampling scheme for the gamma-order generalized normal distribution by extending the traditional Ziggurat algorithm. The extended Ziggurat algorithm enables efficient generation of random variates from light-tailed, Gaussian, and heavy-tailed distributions through a single shape parameter, that is, gamma. Importantly, the method preserves affine equivariance, allowing users to impose any valid covariance structure in the multivariate setting. The proposed adaptation generalizes the standard Ziggurat algorithm while retaining its constant expected cost per sample.

Comprehensive numerical simulations demonstrated that the algorithm is both accurate and computationally efficient across a wide range of dimensionalities and shape parameters. In addition, an application to financial return data confirmed the method's practical relevance and adaptability to real-world heavy-tailed datasets. The gamma-order generalized normal family proved capable of capturing tail behavior more effectively than the Gaussian counterpart, especially in the modeling of asset returns.

Despite these advantages, the current implementation relies on symmetric distributions and does not accommodate skewed gamma-order generalized normal extensions. Furthermore, although the expected cost is constant, the rejection rate in extreme tails is sensitive to the chosen cut-off threshold, which remains manually tuned in the present version. These factors may limit performance when sampling from highly asymmetric or extremely heavy-tailed data.

Future work will focus on three main directions. First, the development of asymmetric versions of the gamma-order generalized normal distribution and corresponding extensions of the Ziggurat algorithm will allow broader modeling capabilities. Second, automatic tuning procedures for the tail cut-off parameter —perhaps using adaptive rejection schemes or optimization heuristics— could further enhance stability and reduce the rejection rate in extreme cases. Third, extending the implementation to exploit parallel hardware and modern multithreading architectures would enable massive-scale simulations for high-dimensional applications.

The proposed algorithm, along with its precomputed tables and source code, is freely available under an open-source license at github.com/UlrichEschcol/Generate_RN_GND, ensuring reproducibilityand facilitating future extensions by the scientific community.

Statements

Acknowledgements

The authors would like to thank anonymous reviewers for their insightful and constructive comments. Their sharp observations and valuable suggestions are greatly appreciated and have contributed to extending and improving the final version of the article.

Author contributions

Conceptualization: C.P. Kitsos; data curation: A. Oliveira, U.E. Nyamsi; formal analysis: C.P. Kitsos, A. Oliveira, U.E. Nyamsi; investigation: C.P. Kitsos, A. Oliveira, U.E. Nyamsi; methodology: C.P. Kitsos, A. Oliveira, U.E. Nyamsi; writing original draft: A. Oliveira, C.P. Kitsos, U.E. Nyamsi writing review and editing: C.P. Kitsos, A. Oliveira, U.E. Nyamsi. All authors have read and agreed to the published version of the article.

Conflicts of interest

The authors declare no conflict of interest.

Data and code availability

The data and computational code are publicly available at github.com/UlrichEschcol/Generate_RN_GND.

Declaration on the use of artificial intelligence (AI) technologies

The authors declare that no generative AI was used in the preparation of this article.

Funding

This work is partially financed to Amilcar Oliveira by national funds through FCT —Fundação para a Ciência e Tecnologia— under the project UIDB/00006/2020 (DOI: 10.54499/UIDB/00006/2020).

References

[1] von Neumann, J., 1951. The future of high-speed computing. Proceedings of the Computer Seminar, IBM, p. 13.

[2] Tukey, W.J., 1962. The future of data analysis. Annals of Mathematical Statistics, 33, 1–67. DOI: 10.1007/978-1-4612-4380-9_31

[3] Horn, R.A. and Johnson, C.R., 2012. Matrix analysis. Cambridge University Press, Cambridge, UK.

[4] Ford, I., Titterington, D.M., and Kitsos, C.P., 1989. Recent advances in nonlinear experimental design. Technometrics, 31, 49–60. DOI: 10.1080/00401706.1989.10488475

[5] Ejsmont, W., Milosevic, B., and Obradovic, M., 2023. A test of normality and independence based on characteristic function. Statistical Papers, 64, 1864–1889. DOI: 10.1007/s00362-022-01365-1

[6] Gorsky, S. and Ma, L., 2022. Multiscale Fisher's independence test for multivariate dependence. Biometrika, 109, 569–587. DOI: 10.1093/biomet/asac013

[7] de Paula Alves, H.J. and Furtado Ferreira, D., 2020. On new robust tests for the multivariate normal mean vector with high-dimensional data and applications. Chilean Journal of Statistics, 11, 117–135. URL: soche.cl/chjs/volumes/11/ChJS-11-02-03.pdf

[8] Kallenberg, W.C.M., Ledwina, T., and Rafałowicz, E., 1977. Testing bivariate independence and normality. Sankhyā A, 39, 42–59. URL: www.jstor.org/stable/25051133

[9] Coulibaly, B.D., Chaibi, G., and El Khomssi, M., 2024. Parameter estimation of the alpha-stable distribution and applications to financial data. Chilean Journal of Statistics, 15, 60–80. DOI: 10.32372/ChJS.15-01-04

[10] Barros, M., Galea, M., Gonzalez, M., and Leiva, V., 2010. Influence diagnostics in the tobit censored response model. Statistical Methods and Applications, 19, 379–397. DOI: 10.1007/s10260-010-0135-y

[11] Vila, R., Saulo, H., Marchant, C., Leiva, V., and Castro, C., 2025. Bivariate extended skew-elliptical Heckman models: Mathematical characterization and an application in economic sciences. Computational and Applied Mathematics, 44, 236. DOI: 10.1007/s40314-025-03169-z

[12] Seijas-Macias, A., Oliveira, A., Oliveira, T., and Leiva, V., 2020. Approximating the distribution of the product of two normally distributed random variables. Symmetry, 12, 1201. DOI: 10.3390/sym12081201

[13] Sarmanov, O.V. and Bratova, Z.N., 1967. Probabilistic properties of bilinear expansions of Hermite polynomials. Theory of Probability and its Applications, 12, 470–481. DOI: 10.1137/1112056

[14] Suzuki, T. and Senba, T., 2011. Applied Analysis: Mathematical Methods in Natural Science. World Scientific Publishing, Singapore.

[15] Lehtonen, J., 2016. The Lambert W function in ecological and evolutionary models. Methods in Ecology and Evolution, 7, 1110–1118. DOI: 10.1111/2041-210X.12568

[16] Kitsos, C.P. and Tavoularis, K.N., 2009. Logarithmic Sobolev inequalities for information measures. IEEE Transactions on Information Theory, 55, 2554–2561. DOI: 10.1109/TIT.2009.2018179

[17] Kitsos, C.P., Vassiliadis, V.G., and Toulias, T.L., 2014. MLE for the gamma-order generalized normal distribution. Discussiones Mathematicae Probability and Statistics, 34, 143–158. DOI: 10.7151/dmps.1168

[18] Kitsos, C.P. and Toulias, T., 2014. Inequalities for the Fisher's information measures. In Rassias, T.M. (editor). Handbook of Functional Equations: Functional Inequalities, pp. 281–313. Springer, New York, NY, US.

[19] Toulias, L. and Kitsos, C., 2014. On the properties of the generalized normal distribution. Discussiones Mathematicae Probability and Statistics, 34, 35–49. DOI: 10.7151/dmps.1167

[20] Kitsos, C.P. and Stamatiou, I.S., 2024. The gamma-order generalized chi-square distribution. Research in Statistics, 2, 2377684. DOI: 10.1080/27684520.2024.2377684

[21] Kitsos, C.P. and Toulias, T.L., 2012. Bounds for the generalized entropy type information measure. Journal of Communication and Computer, 9, 56–64. URL: www.davidpublisher.com/Public/uploads/Contribute/5539f7c1969ba.pdf

[22] Toulias, T.L. and Kitsos, C.P., 2012. Kullback-Leibler divergence of the gamma-ordered normal over t-distribution. Journal of Advances in Mathematics and Computer Science, 2, 198–212. DOI: 10.9734/BJMCS/2012/1358

[23] Marsaglia, G. and Tsang, W.W., 2000. The Ziggurat method for generating random variables. Journal of Statistical Software, 5(8), 1–7. DOI: 10.18637/jss.v005.i08

[24] Kitsos, C.P., Toulias, T.L., and Trandafir, P.C., 2011. On the multivariate gamma-ordered normal distribution. Far East Journal of Theoretical Statistics, 38, 49–73. URL: www.pphmj.com/abstract/6616

[25] Abramowitz, M. and Stegun, I.A., 1972. Handbook of Mathematical Functions with Formulas, Graphs, and Mathematical Tables, Volume 55. National Bureau of Standards, Washington, DC, US.

[26] Kitsos, C.P. and Toulias, L., 2010. New information measures for the generalized normal distribution. Information, 1, 13–27. DOI: 10.3390/info1010013

[27] Kotz, S. and Nadarajah, S., 2006. Kotz-type distribution. Encyclopedia of Statistical Sciences. DOI: 10.1002/0471667196.ess4087.pub2

[28] R Core Team, 2023. R: A language and environment for statistical computing. R Foundation for Statistical Computing, Vienna, Austria.