Deterministic Log-CF Evaluation for Correlated Lognormal Sums

A Log-CF Quadrature/COS Scheme for Positive Sums of Correlated Lognormals

The distribution of a positive lognormal sum is determined by n(n+5)/2 parameters. A log-CF integral representation and a parameterized quadrature/COS evaluator provide deterministic CDF approximations from this finite latent.

Tamás Nagy, Ph.D.

tnagyphd@gmail.com

Working Paper • March 2026

Build-std v2.0+8309cd4685 | 2026-07-19 13:36 | abd244e7

Abstract

The moment-based Latent representation of correlated lognormal sums (Nagy, 2026, *The Exact Latent Distribution of Correlated Lognormal Sums*) relies on scaled moments $c_k = m_k/k!$ whose logarithms grow quadratically in k in non-degenerate positive-weight cases. This creates an extreme dynamic range in the Padé Toeplitz systems used by that implementation. The reported experiments become impractical for $\sigma_{\max} > 0.8$, but this paper does not prove a general condition-number asymptotic, an impossibility result, or an information-theoretic lower bound. The Smooth Latent Operator (Nagy, 2026, *The Smooth Latent Operator: Parameter-Free Distributional Representations via Kernel Moment Recovery*) mitigates the observed instability through regularization without eliminating it in the reported high-volatility cases.

We show that the moment-based route expands an unsuitable formal generating function. The moments are the formal Taylor coefficients associated with $M(z) = E[e^{zS}]$, which is infinite for positive real z in the non-degenerate lognormal setting and has zero Taylor radius at the origin. By contrast, the **Hermite-chaos expansion** — the Wiener polynomial chaos of S in the underlying Gaussian variables — uses coefficients that decay factorially. This occurs because the exponential function e^Y has a convergent Hermite series for $Y \sim N(\mu, \sigma^2)$, with coefficients $\sigma^k/k!$.

The resulting **Hermite Latent** $\Lambda^H = \{c^H_{\mathbf{k}}\}$ lives in ℓ^2 (no Gaussian weight is needed). The characteristic function is represented by a convergent inner product (no Padé resummation needed). The CDF can then be approximated by Fourier-cosine inversion (Fang and Oosterlee, 2008). The numerical conditioning and accuracy of the full evaluator are assessed here only in the reported test regimes.

We interpret this as a **grade-3 Latent**: a representation choice guided by the problem's structure. For lognormal sums, the Hermite-chaos basis is a natural candidate because it matches the underlying Gaussian variables. This is a problem-specific motivation, not a universal basis-selection theorem.

The exact CDF integral representation is determined by the finite generative latent $(w, \mu, \Sigma) \in \mathbb{R}^{n(n+5)/2}$ for positive weights and positive-definite covariance. Its numerical evaluator uses Gauss-

Hermite order, COS order, and a finite log-domain as explicit precision choices. The reported results motivate this route as a useful alternative to moment-based evaluation; rigorous parameter-selection and end-to-end convergence guarantees remain open.

Practitioner's Summary

This paper studies a deterministic CDF evaluator for low-dimensional sums of lognormal risk factors with strictly positive portfolio weights and positive-definite covariance. The implementation evaluates the characteristic function of $\log S$ by tensor-product Gauss–Hermite quadrature and reconstructs the CDF by a finite COS expansion. It is a numerical approximation, not a closed-form distribution and not a uniformly certified error method.

The reported positive-weight tests cover two-asset cases through volatility $\sigma = 2.0$ and one five-asset case. Their maximum discrepancies against seeded Monte Carlo range from 2.04×10^{-4} to 9.14×10^{-4}. Tensor cost grows as Q^n, so the method is not presented as practical for large portfolios without additional sparse-grid or factor-model work. Mixed-sign portfolios are outside the method's supported domain.

Reproduction uses the tracked script forge/fin_fenton_solved/demo_hermite_cos.py with NumPy and SciPy in float64, Monte Carlo seed 42, and the case-specific settings listed in §7.5–7.6. The study does not benchmark against Mahmoud (2010) or establish superiority over prior quadrature methods. VaR and ES are downstream computational proposals; neither risk measure is independently validated here.

---

1. Introduction

1.1 The Moment Growth Problem

The starting obstruction is visible in the growth of the moment coordinates.

Consider a sum $S = \sum_{i=1}^{n} w_i e^{Y_i}$ with positive weights and $Y \sim N(\mu, \Sigma)$, where $\Sigma \succ 0$ throughout the main construction. The moments are (Nagy, 2026, *The Exact Latent Distribution of Correlated Lognormal Sums*, Theorem 1):

$$m_k = E[S^k] = \sum_{|\mathbf{j}|=k} \binom{k}{\mathbf{j}} \prod_{i=1}^{n} w_i^{j_i} \cdot \exp\left(\mathbf{j}^T \mu + \tfrac{1}{2}\mathbf{j}^T \Sigma \mathbf{j}\right)$$

For any component i with $w_i > 0$, the multinomial term with $j_i = k$ gives the exact lower bound

$$m_k \geq w_i^k \exp\left(k\mu_i + \frac{k^2\sigma_i^2}{2}\right).$$

Choose an index i^* with $\sigma_{i^*} = \sigma_{\max}$. Then

$$c_k = \frac{m_k}{k!} \geq \frac{w_{i^*}^k \exp(k\mu_{i^*} + k^2\sigma_{\max}^2/2)}{k!},$$

and hence

$$\log c_k \geq \frac{\sigma_{\max}^2}{2} k^2 + k(\log w_{i^*} + \mu_{i^*}) - \log(k!).$$

The factorial contributes only a lower-order term to $\log |c_k|$ as $k \to \infty$, but it is material at finite order. For the one-asset normalization $\mu = 0$, $\sigma = 1$, and $k = 20$, the exact scaled moment is

$$c_{20} = \frac{e^{200}}{20!} \approx 2.9701 \times 10^{68}.$$

1.2 Why This Growth Is Fatal

The Padé $[N_P - 1/N_P]$ approximant requires solving a Toeplitz system with entries $H_{ij} = c_{N_P+i-j}$. Its entries can span an enormous range because $\log c_k$ grows quadratically in k up to lower-order terms. Entrywise dynamic range alone is not a lower bound for the spectral condition number, so no general formula for $\kappa(H)$ is asserted here. The referenced implementation nevertheless reports singular or numerically unusable solves at moderate Padé orders in its high-volatility cases (Nagy, 2026, *The Exact Latent Distribution of Correlated Lognormal Sums*, Section 7.3).

In the reported Padé experiments, neither regularization (such as the Smooth Latent Operator) nor increased precision resolved the observed singular/ill-conditioned Toeplitz solves in the high-volatility regimes. This suggests that higher-order moments can become increasingly hard to use numerically. One or more diagonal multinomial terms with maximal variance control the leading quadratic rate in $\log c_k$; ties are possible, and no single-term dominance claim is needed.

1.3 Exact Integral Representation and Numerical Evaluator

For strictly positive weights and $\Sigma \succ 0$, the CDF of $S = \sum_{i=1}^n w_i e^{Y_i}$, $Y \sim N(\mu, \Sigma)$, is determined by the finite generative latent $(w, \mu, \Sigma) \in \mathbb{R}^{n(n+5)/2}$:

$$\boxed{F_S(x) = \frac{1}{2} - \frac{1}{\pi} \int_0^\infty \mathrm{Im}\left[\frac{\tilde{\phi}(t)}{t}\, e^{-it\log x}\right] dt, \qquad \tilde{\phi}(t) = \int_{\mathbb{R}^n} S(z)^{it}\, \gamma_n(z)\, dz, \quad x > 0}$$

where $S(z) = \sum_{i=1}^n w_i \exp(\mu_i + \ell_i^T z)$ is evaluated exactly at each state point from the Cholesky factor $L = \mathrm{chol}(\Sigma)$, and γ_n is the standard Gaussian density. The formula composes a bounded Gaussian expectation with Gil–Pelaez inversion under the usual continuity conditions. The implementation approximates these stages by Gauss-Hermite quadrature and Fourier-cosine synthesis respectively. VaR and ES then follow numerically from the resulting CDF and density approximations (§8).

The symbol β denotes the upper endpoint, and α the lower endpoint, of the finite log-domain. The integral representation has no calibration parameters, while its numerical evaluator requires choices of Q quadrature nodes, N COS terms, and the domain $[\alpha, \beta]$; §7.6 gives practical heuristics rather than a proved selection rule. Boundedness of $|S(z)^{it}|$ supports the log-CF integral for positive S. In the reported high-volatility cases ($\sigma > 0.8$), max discrepancies against Monte Carlo are on the 10^{-4} scale in standard float64 arithmetic (§7.5).

1.4 The Key Observation

The moments are the Taylor coefficients of the MGF $M(z)$. Because the MGF has zero convergence radius, moment-based methods attempt to resum a divergent series.

However, this divergence is not a statement about the distribution's existence or smoothness; it is a consequence of using formal MGF/moment coordinates as the intermediate object in a numerical evaluation chain. The lognormal law is moment-indeterminate, so its complete moment sequence does not uniquely determine the distribution among all probability laws on the positive half-line. Within the specified finite-dimensional lognormal family the parameters (w, μ, Σ) determine the law, but the monomial sequence remains poorly suited to the Padé computation studied here.

The question thus becomes: is there a representation whose coefficients are well-behaved?

1.5 Gaussian-Adapted Coordinates

For lognormal sums, the relevant structural fact is direct: the underlying variables Y_i are Gaussian. Hermite polynomials form an orthogonal basis for functions under Gaussian measure, making Hermite chaos a natural representation to analyze. This observation motivates the coefficient calculation in §2 without implying that the basis is universally optimal or that it controls the conditioning of the separate GH/COS evaluator.

1.6 Organization

Sections 2–3 develop the Hermite-chaos expansion and define the Hermite Latent. Section 4 builds the evaluation chain: log-space characteristic function via Gauss–Hermite quadrature, CDF via Fourier-cosine inversion, and the boxed integral representation. Section 5 discusses grade-3 Latents and a representation-structure matching heuristic. Section 6 compares moment and Hermite Latents. Section 7 examines three routes from Latent to CDF (Edgeworth, Hermite-COS, quadratic special case) and provides numerical validation. Section 8 derives numerical VaR and ES evaluation from the CDF approximation. Section 9 addresses the curse of dimensionality and factor-model mitigation. Sections 10–11 discuss implications for the Latent framework and conclude.

---

2. The Hermite-Chaos Expansion

2.1 Setup

To replace moments with Gaussian-adapted coordinates, first standardize the latent Gaussian vector. Unless a rank-deficient extension is explicitly discussed, the model domain in this paper is

$$w_i > 0, \qquad \mu \in \mathbb{R}^n, \qquad \Sigma \in \mathbb{S}^n_{++}.$$

Let $L = \mathrm{chol}(\Sigma)$ denote the lower-triangular Cholesky factor, and let $Z = L^{-1}(Y - \mu) \sim N(0, I_n)$ denote the standardized variables. We then have:

$$Y_i = \mu_i + \sum_{j=1}^{n} L_{ij} Z_j = \mu_i + \ell_i^T Z$$

where $\ell_i = (L_{i1}, \dots, L_{in})^T$ is the i-th row of L.

From this, each lognormal component is:

$$e^{Y_i} = e^{\mu_i + \ell_i^T Z} = e^{\mu_i + \|\ell_i\|^2/2} \cdot e^{\ell_i^T Z - \|\ell_i\|^2/2}$$

The second factor is precisely the stochastic exponential (Wick exponential) of $\ell_i^T Z$.

Positive-semidefinite extension. If $\Sigma \succeq 0$ has rank $r < n$, choose a full-column factor $B \in \mathbb{R}^{n \times r}$ with $\Sigma = BB^T$ and write $Y = \mu + BZ$ for $Z \sim N(0, I_r)$. The same Hermite and log-CF constructions then use r Gaussian variables and an r-dimensional Gaussian integral. This boundary case does not use L^{-1} or an ordinary $n \times n$ Cholesky factor. It is mathematically valid but is not the implementation exercised in §7.

2.2 Hermite Expansion of a Single Lognormal

For a single standard normal Z_j, the Hermite polynomials $\{H_k(z)\}_{k=0}^{\infty}$ form an orthogonal basis of $L^2(\mathbb{R}, \gamma)$ where γ is the standard Gaussian measure. The exponential has the well-known expansion:

$$e^{aZ_j} = e^{a^2/2} \sum_{k=0}^{\infty} \frac{a^k}{k!} H_k(Z_j)$$

Fact: The coefficients $a^k/k!$ decay factorially for any fixed a.

For the multivariate case, define multi-index Hermite polynomials $H_{\mathbf{k}}(Z) = \prod_{j=1}^{n} H_{k_j}(Z_j)$ for $\mathbf{k} = (k_1, \dots, k_n) \in \mathbb{N}_0^n$. The stochastic exponential expands as:

$$e^{\ell_i^T Z - \|\ell_i\|^2/2} = \sum_{\mathbf{k} \in \mathbb{N}_0^n} \prod_{j=1}^{n} \frac{L_{ij}^{k_j}}{k_j!} H_{\mathbf{k}}(Z)$$

Therefore:

$$e^{Y_i} = e^{\mu_i + \|\ell_i\|^2/2} \sum_{\mathbf{k}} \left(\prod_{j=1}^{n} \frac{L_{ij}^{k_j}}{k_j!} \right) H_{\mathbf{k}}(Z)$$

2.3 Hermite Expansion of the Sum

The weighted sum is:

$$S = \sum_{i=1}^{n} w_i e^{Y_i} = \sum_{\mathbf{k} \in \mathbb{N}_0^n} c_{\mathbf{k}}^H H_{\mathbf{k}}(Z)$$

where the **Hermite-chaos coefficients** are:

$$\boxed{c_{\mathbf{k}}^H = \sum_{i=1}^{n} w_i\, e^{\mu_i + \|\ell_i\|^2/2} \prod_{j=1}^{n} \frac{L_{ij}^{k_j}}{k_j!}}$$

This is a **closed-form** expression: a finite sum of products of known quantities.

2.4 Growth Rate

The coefficient formula now permits a direct quantitative growth estimate.

Proof provenance. Theorem 1, Proposition 1, Theorem 2, and Proposition 2 are conventional analytical statements proved in this manuscript. No machine-verification claim is made for them, and no machine-checkable companion statement currently exists. Their publication route is therefore the ordinary one: explicit hypotheses and peer-checkable paper proofs. The absence of a formal companion is a limitation of proof provenance, not evidence against the mathematics.

Theorem 1 (Hermite Coefficient Bound). *The Hermite-chaos coefficients satisfy:*

$$|c^H_{\mathbf{k}}| \le n \cdot \frac{(n\sigma_{\max})^{|\mathbf{k}|}}{|\mathbf{k}|!} \cdot \max_i |w_i| e^{\mu_i + \sigma_i^2/2}$$

where $|\mathbf{k}| = k_1 + \cdots + k_n$*.*

In particular, for fixed total degree $K = |\mathbf{k}|$*:*

$$\sum_{|\mathbf{k}|=K} |c^H_{\mathbf{k}}|^2 \le D \cdot \frac{(n\sigma_{\max}^2)^K}{K!}$$

where D *depends on* w*,* μ*, and the component variances (equivalently, it may depend on* Σ*), but not on* K*.*

Proof. Fix a multi-index $\mathbf{k}$. Each term in $c^H_{\mathbf{k}}$ is bounded by

$$|w_i| e^{\mu_i + \sigma_i^2/2} \prod_j \frac{|L_{ij}|^{k_j}}{k_j!}.$$

Since $|L_{ij}| \le \sigma_{\max}$, the product is at most

$$\frac{\sigma_{\max}^{|\mathbf{k}|}}{\prod_j k_j!}.$$

The identity $\prod_j k_j! = |\mathbf{k}|!/\binom{|\mathbf{k}|}{\mathbf{k}}$ and the bound $\binom{|\mathbf{k}|}{\mathbf{k}} \le n^{|\mathbf{k}|}$ give the first estimate after summing over i.

For the degree-wise bound, use $|\sum_i a_i|^2 \le n \sum_i |a_i|^2$ together with $1/(k_j!)^2 \le 1/k_j!$. The multinomial theorem then gives

$$\sum_{|\mathbf{k}|=K} |c^H_{\mathbf{k}}|^2 \le \frac{n}{K!} \sum_i w_i^2 e^{2\mu_i + \sigma_i^2} \|\ell_i\|^{2K},$$

which implies the displayed estimate after enlarging D. $\square$

Comparison with moments. For one asset with $w = 1$, $\mu = 0$, and $\sigma = 1$, the scaled raw moment is $e^{K^2/2}/K!$, whereas the degree-K Hermite coefficient is $e^{1/2}/K!$. Their ratio is therefore exactly $e^{(K^2-1)/2}$. The shared factorial must not be omitted when quoting either coefficient; the distinction is the quadratic exponential factor in the moment coordinate.

---

3. The Hermite Latent

3.1 Definition

The coefficient sequence from Section 2 is the representation used below.

Definition 1 (Hermite Latent). The **Hermite Latent** of the sum $S = \sum w_i e^{Y_i}$ is the sequence:

$$\Lambda^H = \{c_{\mathbf{k}}^H\}_{\mathbf{k}\in\mathbb{N}_0^n} \in \ell^2(\mathbb{N}_0^n)$$

By Theorem 1, we have $\|\Lambda^H\|_2^2 = \sum_{\mathbf{k}} |c_{\mathbf{k}}^H|^2 < \infty$ (a convergent factorial series after grouping by total degree). Consequently, no Gaussian weight is needed. The Hermite Latent naturally lives in the standard ℓ^2 space.

3.2 The Hilbert Space Structure

The Parseval identity for Hermite-chaos gives:

$$\mathrm{Var}(S) = \sum_{\mathbf{k}\neq\mathbf{0}} |c_{\mathbf{k}}^H|^2 \cdot \mathbf{k}! = \|\Lambda^H\|_w^2$$

where $\mathbf{k}! = \prod k_j!$ and $\|\cdot\|_w$ is the weighted ℓ^2 norm. Because lognormal sums have finite variance, this directly confirms ℓ^2 membership.

Comparison with the moment Hilbert space: The moment Latent $\Lambda^M = \{c_k\}$ required a Gaussian-weighted space $\mathcal{H}_{\beta_{\mathrm{wt}}}$ with weight exponent $\beta_{\mathrm{wt}} > \sigma_{\max}^2/2$ to ensure convergence (Nagy, 2026, *The Smooth Latent Operator: Parameter-Free Distributional Representations via Kernel Moment Recovery*, Section 2.1). By contrast, the Hermite Latent needs no such weight. It lives naturally in ℓ^2.

3.3 Smoothness

Beyond square summability, the representation varies regularly with the model parameters.

Proposition 1. *On $\mathbb{R}^n \times \mathbb{R}^n \times \mathbb{S}_{++}^n$, the Hermite Latent map $\Phi^H : (w, \mu, \Sigma) \mapsto \Lambda^H \in \ell^2$ is smooth (infinitely Fréchet differentiable).*

Proof. Each $c_{\mathbf{k}}^H$ is a smooth function of (w, μ, Σ) (it involves the Cholesky factor $L(\Sigma)$, which is smooth for $\Sigma \succ 0$, and exponentials). The ℓ^2 convergence of the derivatives follows from the same factorial decay. □

---

4. From Hermite Latent to CDF

This section separates two constructions. First (§4.1–4.3), the Hermite truncation S_K provides a polynomial approximation in Gaussian variables. Second (§4.4), the implemented pipeline evaluates $S(z) = \sum w_i \exp(\mu_i + \ell_i^T z)$ directly at quadrature nodes and bypasses the Hermite truncation. Convergence of the first construction does not establish convergence or conditioning of the second; the GH/COS evaluator is assessed only through its stated numerical choices and experiments.

4.1 The Characteristic Function

The CF of S is:

$$\phi_S(t) = E[e^{itS}] = E\left[\exp\left(it \sum_{\mathbf{k}} c_{\mathbf{k}}^H H_{\mathbf{k}}(Z)\right)\right]$$

This is the CF of a random variable expressed as a convergent Hermite series. For a finite truncation at total degree K:

$$S_K = \sum_{|\mathbf{k}| \le K} c_{\mathbf{k}}^H H_{\mathbf{k}}(Z)$$

the CF is computable because S_K is a polynomial in Gaussian variables — its distribution is determined by a finite set of operations.

4.2 The Preferred Method: Gauss-Hermite Quadrature as Finite Sum

For moderate n (say $n \le 8$), the CF of S_K is:

$$\phi_{S_K}(t) = \int_{\mathbb{R}^n} e^{it \cdot p_K(z)} \gamma_n(z)\, dz \approx \sum_{\ell=1}^{Q^n} q_\ell\, e^{it \cdot p_K(z_\ell)}$$

where (z_ℓ, q_ℓ) are n-dimensional Gauss-Hermite quadrature nodes and weights, $p_K(z) = \sum_{|\mathbf{k}| \le K} c_{\mathbf{k}}^H H_{\mathbf{k}}(z)$ is a polynomial, and γ_n is the standard Gaussian density.

This is a **finite sum of Q^n terms**, each involving the exponential of a known polynomial evaluated at a known quadrature node. It is a parameterized quadrature approximation: for fixed Q, it is a finite algebraic expression built from known quantities. It requires no Monte Carlo simulation or iterative solver for the Gaussian-contraction step, but its accuracy depends on the numerical precision choices.

Why boundedness helps. The modulus $|e^{it \cdot p_K(z)}| = 1$ for all z, so the integrand is bounded. This differs from the moment-based approach, where the real exponential integrand grows without bound. Boundedness prevents integrand blow-up, but by itself it does not prove a quadrature error rate or a condition-number bound.

Convergence rate. Gauss-Hermite with Q nodes per dimension integrates polynomials of degree $2Q - 1$ exactly. The integrand $e^{it \cdot p_K(z)}$ is not polynomial, but its Taylor expansion $\sum_{j=0}^{J} (it \cdot p_K(z))^j / j!$ at any truncation J IS polynomial (degree KJ), and Gauss-Hermite with $Q \ge (KJ+1)/2$

nodes integrates it exactly. This provides a heuristic explanation for potentially rapid convergence as Q increases, but the present paper does not give a uniform Gauss-Hermite error bound for the resulting oscillatory non-polynomial integrand.

For reference, $Q = 15$ nodes per dimension integrates polynomials of degree 29 exactly. This fact does not quantify the error for the non-polynomial CF integrand. End-to-end CDF accuracy also depends on the COS parameters $(N, [\alpha, \beta])$; see §7.5 for the empirical pipeline comparison.

Cost. The corresponding node counts are 225 for $(n, Q) = (2, 15)$, 759,375 for $(5, 15)$, and 10^8 for $(8, 10)$. Runtime is implementation- and hardware-dependent. The tensor rule becomes impractical as dimension increases; sparse-grid or conditional methods would require separate validation.

4.3 Alternative Methods

Method A: Hermite moment transfer. The moments of S_K are:

$$E[S_K^p] = \sum_{\mathbf{k}_1,\ldots,\mathbf{k}_p} \prod_{r=1}^{p} c^H_{\mathbf{k}_r} \cdot E\left[\prod_{r=1}^{p} H_{\mathbf{k}_r}(Z)\right]$$

The expectation is computed from Hermite product identities or, after polynomial expansion, Wick-/Isserlis contractions. These Hermite-derived moments can feed the Padé-COS chain of Nagy (2026, *The Exact Latent Distribution of Correlated Lognormal Sums*). However, the moments of S_K only approximate the raw moments m_p, and no improved Padé-conditioning guarantee follows. This auxiliary route is not used in the numerical validation.

Method B (preferred): Gauss-Hermite quadrature as described in Section 4.2.

Method C: Hybrid. Use the Hermite coefficients to evaluate $p_K(z)$ at quadrature nodes, compute the CF by quadrature (Method B), then apply COS inversion.

All three methods above use the Hermite truncation S_K. But $S(z) = \sum_{i=1}^n w_i \exp(\mu_i + \ell_i^T z)$ can be evaluated **exactly** at any z from the finite latent, making the Hermite expansion unnecessary for computation. Theorem 1 supports the auxiliary approximation $S_K \to S$; it supplies no convergence or conditioning result for the direct GH/COS pipeline introduced next.

4.4 The Exact Integral Representation

For strictly positive weights and $\Sigma \succ 0$, the distribution of $S = \sum_{i=1}^n w_i e^{Y_i}$, $Y \sim N(\mu, \Sigma)$, is determined by the finite generative latent $(w, \mu, \Sigma) \in \mathbb{R}^{n(n+5)/2}$. The CDF is:

$$\boxed{F_S(x) = \frac{1}{2} - \frac{1}{\pi}\int_0^\infty \operatorname{Im}\left[\frac{\tilde{\phi}(t)}{t}\, e^{-it\log x}\right] dt, \qquad x > 0}$$

where $\tilde{\phi}(t) = E[S^{it}] = E[e^{it\log S}]$ is the **log-space characteristic function** (Mellin–Fourier transform of the density):

$$\tilde{\phi}(t) = \langle S(\cdot)^{it},\, \mathbf{1}\rangle_{\gamma_n} = \int_{\mathbb{R}^n} S(z)^{it}\, \gamma_n(z)\, dz$$

At any state point z, the finite latent gives

$$S(z) = \sum_{i=1}^{n} w_i \exp(\mu_i + \ell_i^T z), \qquad L = \mathrm{chol}(\Sigma).$$

For a positive-semidefinite covariance of rank $r < n$, replace Lz by Bz with $\Sigma = BB^T$, $B \in \mathbb{R}^{n\times r}$, and $z \in \mathbb{R}^r$. The exact Gaussian contraction is then r-dimensional. The implemented and validated path below uses the full-rank SPD setup.

This evaluation is exact. Apply the standard Gil–Pelaez inversion (Gil-Pelaez, 1951) to $T = \log S$. The identities

$$F_S(x) = F_T(\log x), \qquad \tilde{\phi}(t) = \phi_T(t)$$

reduce the result to

$$F_T(y) = \frac{1}{2} - \frac{1}{\pi}\int_0^\infty \mathrm{Im}\left[\frac{\phi_T(t)e^{-ity}}{t}\right] dt, \qquad y = \log x.$$

The integral representation has **no calibration parameters** beyond the generative latent. Its numerical evaluator nevertheless has explicit precision and domain choices:

- No truncation order K (no Hermite expansion — $S(z)$ is evaluated directly)
- The domain $[\alpha, \beta]$ is set from the cumulants of $\log S$ (deterministic, not fitted)
- The COS order N and quadrature order Q are precision parameters analogous to machine epsilon in evaluating $N(x)$

The formula is a composition of four Latent algebra operations:

$$\underbrace{(w,\mu,\Sigma)}_{\text{finite latent}} \xrightarrow{\text{map}} \underbrace{S(\cdot) \in L^2(\gamma_n)}_{\text{grade-1}} \xrightarrow{\text{complexify}} \underbrace{S(\cdot)^{it}}_{\text{grade-1, complex}} \xrightarrow{\text{contract}_\gamma} \underbrace{\tilde{\phi}(t)}_{\text{grade-0}} \xrightarrow{\text{contract}_{dt}} F_S(x)$$

The first arrow is a **map**: the finite latent generates $S(z)$ by matrix-vector multiply (Lz, $O(n^2)$) followed by n exponentials and a weighted sum. The second arrow is **complexification**: $S^{it} = e^{it\log S}$ extends the distribution to $\mathbb{C}$ via the log-space (Mellin) transform. The third arrow is a **contraction** against the Gaussian measure — the inner product $\langle S(\cdot)^{it}, \mathbf{1}\rangle_{\gamma_n}$ that produces the log-space characteristic function $\tilde{\phi}(t)$. The fourth arrow is a **contraction** against Lebesgue measure on frequency space — the Fourier inversion integral.

The displayed arrows define the exact integral representation. Its Gauss-Hermite and COS implementations introduce the numerical choices Q, N, and $[\alpha, \beta]$; the reported positive-weight tests cover volatilities from $\sigma = 0.2$ through $\sigma = 2.0$.

The two contractions. The formula consists of two inner products composed:

Contraction	Domain	Against	Produces
$\langle S(\cdot)^{it}, \mathbf{1}\rangle_{\gamma_n}$	State space $\mathbb{R}^n$	Gaussian γ_n	$\tilde{\phi}(t)$

Contraction	Domain	Against	Produces
$\int_0^\infty \mathrm{Im}[\tilde{\phi}(t)e^{-it\log x}/t]\,dt$	Frequency space $\mathbb{R}_+$	Lebesgue dt	$F_S(x)$

The first integral is well-defined because $|S(z)^{it}| = 1$ for $S(z) > 0$. The second is interpreted as the Gil–Pelaez improper inversion limit and is valid at continuity points under that theorem's hypotheses. Although Riemann–Lebesgue gives $\tilde{\phi}(t) \to 0$ when $\log S$ has an integrable density, decay to zero by itself would not prove convergence of the displayed $1/t$ integral. The moment-generating route instead involves e^{tS}, whose expectation is infinite for $t > 0$ in the non-degenerate positive lognormal setting.

4.5 Evaluators

The exact integral representation requires evaluating two integrals. Different evaluator choices produce different computational formulas, but the underlying algebraic object is the same.

Evaluator for the Gaussian contraction. Gauss-Hermite quadrature with Q nodes per dimension:

$$\tilde{\phi}(t) \approx \sum_{\ell=1}^{Q^n} q_\ell\, S(z_\ell)^{it}$$

where $S(z_\ell) = \sum_{i=1}^n w_i \exp(\mu_i + \ell_i^T z_\ell)$ is computed **directly** from the finite latent — no Hermite expansion, no truncation at degree K. The integrand $|S(z)^{it}| = 1$ is bounded for $S(z) > 0$, and the reported tests show rapid convergence as Q increases in the evaluated regimes; a uniform quadrature error bound is not established here. The validation script checks the analytical mean to machine precision, but it does not independently certify 12-digit CF accuracy at every frequency.

Evaluator for the Fourier inversion. The COS method (Fang and Oosterlee, 2008) discretizes the Fourier integral on a log-space domain $[\alpha, \beta]$ as a cosine series with N terms:

$$F_S(x) \approx u + \sum_{k=1}^{N} \frac{2}{k\pi} A_k \sin(k\pi u), \qquad u = \frac{\log x - \alpha}{\beta - \alpha}$$

where $A_k = \mathrm{Re}[\tilde{\phi}(k\pi/(\beta-\alpha)) \cdot e^{-ik\pi\alpha/(\beta-\alpha)}]$. The domain $[\alpha, \beta]$ and order N are evaluator parameters. The tracked implementation sets the domain deterministically to 12 weighted standard deviations on either side of the weighted mean of $\log S$ at the selected GH nodes. The caller supplies both Q and N explicitly; there is no automatic order selector or stopping rule. Section 7.6 gives a coupled-refinement diagnostic for choosing those inputs. Smoothness of the log-sum density supports spectral convergence on a well-chosen finite domain; the observed rate also depends on tail truncation and endpoint behavior, so no universal exponential rate is claimed.

Why log-space. Working with $T = \log S$ maps the heavy right tail of S to a lighter-tailed coordinate. In the reported comparisons, log-space COS substantially outperforms the reference direct-space implementation. The observed rate depends on domain and endpoint behavior; geometric convergence and a fixed sufficient order are not proved.

Why the integral is necessary (the coordinate mismatch). In y-space ($Y = \log V$), the generative structure is quadratic (the Gaussian density $e^{-(y-\mu)^T\Sigma^{-1}(y-\mu)/2}$), but the sum $S = w^T e^y$ is transcendental. In v-space ($V = e^Y$), the sum $S = w^T v$ is linear, but the measure involves $\log v$ (lognormal density). No single coordinate system makes both the measure and the integrand quadratic simultaneously — the nonlinearity is an invariant of the problem, preserved under smooth coordinate changes. The evaluator cost is the price of bridging this coordinate mismatch.

Dimension reduction caveat. The tensor-product GH evaluator costs Q^n nodes in the full-rank SPD setup. A genuinely rank-r covariance $\Sigma = BB^T$ reduces the exact Gaussian state dimension to r and the tensor rule to Q^r. By contrast, a factor-plus-idiosyncratic decomposition $\Sigma = BB^T + D$ with nonzero D retains the idiosyncratic variables. Independence does not factorize $E[(\sum_i w_i e^{Y_i})^{it}]$, and Cholesky triangularity alone does not justify an nQ^r evaluator. Sparse grids or conditional approximations require separate derivation and validation.

Combining both evaluators gives the computational formula:

$$F_S(x) \approx u + \sum_{k=1}^{N} \frac{2}{k\pi} \sin(k\pi u) \sum_{\ell=1}^{Q^n} q_\ell \cos\left(k\pi \cdot \frac{\log S(z_\ell) - \alpha}{\beta - \alpha}\right)$$

For fixed evaluator settings this is a double finite sum — no Monte Carlo sampling or nonlinear solve is used in CDF evaluation. Two numerical operators approximate the two contractions: the COS operator replaces Fourier inversion by a cosine sum, and the GH operator replaces the Gaussian integral by a weighted sum.

Component	Cost	Nature
Cholesky $L = \mathrm{chol}(\Sigma)$	$O(n^3)$, once	From latent
$S(z_\ell) = \sum w_i \exp(\mu_i + \ell_i^T z_\ell)$	$O(n^2)$ per node	Direct evaluation — no Hermite expansion
GH quadrature (Gaussian contraction)	Q^n nodes	Tensor evaluator for contraction 1
COS synthesis (Fourier contraction)	N terms	Evaluator for contraction 2
Total	$O(n^2 \cdot N \cdot Q^n)$	Baseline tensor implementation

The PDF follows by differentiating the COS synthesis:

$$f_S(x) = \frac{1}{x(\beta - \alpha)} \left(1 + \sum_{k=1}^{N} 2\, A_k \cos(k\pi u)\right)$$

The density is synthesized directly from the same grade-0 Latents $\{A_k\}$ — no additional contraction needed.

4.6 Truncation Order and Convergence for All Volatilities

For a single lognormal e^Y with $Y \sim N(\mu, \sigma^2)$, the Parseval identity gives:

$$E[(e^Y)^2] = e^{2\mu+\sigma^2} \cdot \sum_{k=0}^{\infty} \frac{\sigma^{2k}}{k!} = e^{2\mu+2\sigma^2}$$

The L^2 truncation error at degree K is:

$$\varepsilon_K^2 = e^{2\mu+\sigma^2} \cdot \sum_{k>K} \frac{\sigma^{2k}}{k!}$$

This is the tail of e^{σ^2}, which **converges for ALL** σ — the factorial $k!$ defeats the power σ^{2k} for $k > e\sigma^2$ (by Stirling's approximation).

Theorem 2 (L^2 convergence for finite lognormal sums). *Let S_K be the total-degree Hermite truncation and write $\sigma_i^2 = \|\ell_i\|_2^2$. Then*

$$E[(S - S_K)^2] \le n \sum_{i=1}^{n} w_i^2 e^{2\mu_i+\sigma_i^2} \left(\sum_{m>K} \frac{\sigma_i^{2m}}{m!} \right).$$

Consequently $S_K \to S$ in L^2 for every fixed finite latent.

Proof. The degree-m Hermite energy of component i is $w_i^2 e^{2\mu_i+\sigma_i^2} \sigma_i^{2m}/m!$ by the multinomial theorem. The squared norm of the sum of the n component remainders is at most n times the sum of their squared norms. Summing over $m > K$ gives the bound. □

Concrete single-asset truncation errors. For $n = 1$, the relative root-mean-square error is exactly

$$\frac{\|S - S_K\|_2}{\|S\|_2} = \left(e^{-\sigma^2} \sum_{k>K} \frac{\sigma^{2k}}{k!} \right)^{1/2}.$$

Direct evaluation gives:

σ	Regime	$K = 10$	$K = 15$	$K = 20$
0.3	Normal equity	2.69×10^{-10}	9.02×10^{-16}	1.40×10^{-21}
0.5	Volatile equity	6.89×10^{-8}	2.97×10^{-12}	5.92×10^{-17}
0.8	High volatility	1.02×10^{-5}	4.56×10^{-9}	9.51×10^{-13}
1.0	Higher volatility	1.00×10^{-4}	1.37×10^{-7}	8.69×10^{-11}
1.2	Extreme test regime	6.10×10^{-4}	2.06×10^{-6}	3.24×10^{-9}
2.0	Hypothetical extreme	5.33×10^{-2}	2.21×10^{-3}	4.39×10^{-5}

These values concern approximation of the random variable by its Hermite truncation, not end-to-end CDF error. L^2 convergence implies convergence in probability and hence convergence of $F_{S_K}(x)$ to $F_S(x)$ at continuity points, but no uniform finite-K CDF bound is derived here. The implemented Method B evaluates $S(z)$ directly and bypasses K; its CDF accuracy instead depends on $(Q, N, [\alpha, \beta])$ and is assessed separately in §7.5.

4.7 Multi-Asset Truncation

For n assets with truncation degree K, the number of multi-indices is $\binom{K+n}{n}$:

n	K	$\binom{K+n}{n}$	GH quadrature cost ($Q = 15$)
2	20	231	225
5	20	53,130	759,375
5	25	142,506	759,375
8	15	490,314	10^8

The Hermite coefficient enumeration grows combinatorially with n, but the GH quadrature cost (Q^n) is independent of K. For the preferred Method B pipeline, the dominant cost is the Q^n quadrature evaluations, not the coefficient enumeration.

---

5. Grade-3 Latents

5.1 The Graded Latent Hierarchy

We now have three levels of Latent structure:

Grade	Object	Space	Role
0	Parameters (w, μ, Σ)	$\mathbb{R}^{n+n+n(n+1)/2}$	Problem definition
1	Distribution Latent Λ	Basis-dependent	What the distribution IS
2	Extraction Latent α^*	$\mathbb{R}_{>0}$	How hard to extract (condition number)
3	**Representation Latent $\mathcal{B}^*$**	**Basis space**	**Which representation to use**

The grade-3 Latent $\mathcal{B}^*$ denotes a structurally motivated basis choice. It informs the grade-2 Latent (extraction difficulty) and hence the practical feasibility of the entire chain.

5.2 The Representation-Structure Matching Heuristic

The following qualitative heuristic synthesizes a grade-3 selection idea from classical results.

Heuristic 3 (Representation-Structure Matching). When a random input has a distinguished generative measure or geometry, begin numerical design with basis families adapted to that structure. For example, Hermite polynomials are natural candidates for Gaussian inputs, Fourier modes for periodic variables, and Chebyshev or COS bases for suitable bounded intervals.

This is a search rule, not a quantified theorem. Matching alone does not imply factorial coefficient decay, uniform conditioning, minimax optimality, or divergence of every mismatched basis. Any such conclusion requires problem-specific hypotheses on regularity, complex-domain growth, truncation, and the numerical operator.

For the lognormal sum treated here, the explicit coefficient formula in §2 proves factorial decay for the Hermite expansion under the stated finite-dimensional Gaussian model. The moment-growth calculation in §1 separately explains the difficulty of the formal MGF coordinates in that model. These two calculations motivate the heuristic but do not establish a universal representation-selection law.

Formalization note. The repository file fenton_hermite_p3.py supplies only auxiliary arithmetic regression checks. It does not formalize this heuristic, the Hermite expansion, Theorems 1–2, or the GH/COS convergence chain.

Examples motivating the heuristic:

System	Generative measure μ	Matched basis $\mathcal{B}^*$	Coefficient behavior	Limitation or comparator
Lognormal sums	$N(0,I)$ (Gaussian)	Hermite $\{H_{\mathbf{k}}(Z)\}$	$\sigma^k/k!$ (factorial)	Moments: $e^{k^2\sigma^2/2}$
General PCE	Distribution of X	Askey polynomial for X	Depends on analyticity	Wrong Askey polynomial: Runge-type blowup
Bounded-domain pricing	Uniform on $[a,b]$	Chebyshev/COS	Geometric/exponential	Hermite: poor on bounded domain
Square-integrable random element	Covariance operator C	Leading KL/PCA eigenfunctions	Rank-m mean-square error $\sum_{j>m}\lambda_j$	Optimal only for linear orthogonal rank-m reconstruction

The last row uses the precise Karhunen–Loève/PCA optimality statement. Let X be a centered square-integrable random element in a separable Hilbert space, and suppose its covariance operator C is positive, self-adjoint, and trace class, with eigenvalues $\lambda_1 \geq \lambda_2 \geq \cdots \geq 0$. For every rank-$m$ orthogonal projection P,

$$E\|X - PX\|^2 = \mathrm{tr}(C) - \mathrm{tr}(PC) \geq \sum_{j>m} \lambda_j.$$

Equality is attained by projection onto a leading m-dimensional covariance eigenspace. The optimizer need not be unique when eigenvalues are tied. This theorem concerns mean-square error within the class of linear orthogonal rank-m reconstructions. It does not rank wall-clock speed, conditioning of a downstream numerical operator, or nonlinear, adaptive, overcomplete, or task-specific representations.

5.2.1 Relation to Existing Results

Heuristic 3 is an interpretive **synthesis** of results scattered across five fields:

Field	Known result	Role in Heuristic 3
Probability (Cameron-Martin 1947)	Hermite expansion of $L^2(\gamma)$	Motivates the grade-3 interpretation
UQ (Xiu-Karniadakis 2002)	Askey scheme for polynomial chaos	Motivates the matching principle
Approximation theory (Trefethen 2013)	Basis choice affects convergence rate	Motivates comparison of representation-dependent rates
Orthogonal polynomials (Favard 1935)	Three-term recurrence positive measure	Gives the structural characterization of μ from the basis
Functional analysis (KL/PCA)	KL modes minimize rank-m linear mean-square error	Supplies one metric-specific optimality theorem

The purpose of this section is synthesis: it records an interpretive connection between these cases and the grade hierarchy. It does not establish a single universal optimality theorem or a termination proof.

5.2.2 The Matching Principle as a Theory of Representability

Heuristic 3 can be read as a **perspective on representability**: it suggests when a system may admit an efficient finite representation and why.

The Latent Theorem (Nagy, 2026, *The Latent: Finite Sufficient Representations of Smooth Systems*, Theorem 1) assumes exponential compressibility in a specified extraction basis $\mathcal{B}$. Here $\rho_{\mathcal{B}} > 1$ denotes the basis-dependent coefficient-decay parameter: coefficients of total degree K are bounded geometrically in K at that rate. For a grade-r object, its upper bound is

$$\|\Lambda - \Lambda_N\|^2 \leq C_1 N^{r-1} \rho_{\mathcal{B}}^{-2N} \leq C_1' \rho'^{-2N}, \qquad 1 < \rho' < \rho_{\mathcal{B}}.$$

The polynomial prefactor comes from multi-index multiplicity and can be absorbed only by the arbitrarily small rate loss from $\rho_{\mathcal{B}}$ to ρ'. Thus logarithmic mode count follows from exponential compressibility; the theorem does not derive exponential compressibility from generic smoothness.

The distinctions are material. A C^∞ object has derivatives of every finite order, but need not be analytic. Analytic continuation to a suitable complex neighborhood often implies exponential spectral decay in an adapted classical basis, while the coefficient condition $\rho_{\mathcal{B}} > 1$ states that decay directly. Neither analyticity nor a fixed geometric coefficient rate follows from C^∞ smoothness alone. The following schematic records the heuristic interpretation:

$$\text{Latent existence statement} \quad \xrightarrow{\text{structural heuristic}} \quad \text{choose a candidate basis adapted to the generative measure}$$

Together, they answer the three fundamental questions of representability:

Question	Answer	Theorem
Does an exponentially accurate truncation exist?	Yes when $\rho_{\mathcal{B}} > 1$ in a specified extractable basis	Latent Theorem (Nagy, 2026, Thm. 1)
How large is it?	Logarithmic mode count after the polynomial-prefactor rate loss $\rho' < \rho_{\mathcal{B}}$	Latent Theorem
Which representation may be effective?	**One whose orthogonality measure matches the generative structure**	**Heuristic 3 (this paper)**

This is analogous to a construction heuristic in logic: the Latent Theorem is an existence statement, while Heuristic 3 suggests how a basis might be selected by matching structure.

Under its hypothesis $\rho_{\mathcal{B}} > 1$, the Latent framework gives an exponentially accurate finite truncation without identifying the best basis. Heuristic 3 adds the practical suggestion that generative structure can guide basis selection; it does not establish exponential decay, uniqueness, or optimality.

For lognormal sums, the Gaussian input makes Hermite coordinates a natural candidate. This observation motivates which representation to test; it does not prove that the basis is universally optimal.

5.3 Grade-3 as a Basis-Selection Map

The grade-3 Latent can be formalized as a map from system parameters to basis choice:

$$\mathcal{B}^* : (\text{system class}) \to (\text{basis family})$$

For parametric families (like lognormal sums), the Hermite basis is a structurally motivated choice across the parameter values considered here. This is an interpretive principle, not a proved universal optimum.

This contrasts with the grade-2 Latent α^*, which varies continuously with (w, μ, Σ). The hierarchy:

- Grade 3 is **discrete** (basis selection) — chosen in response to the problem class.
- Grade 2 is **continuous** (resolution) — determined by parameter values.
- Grade 1 is **functional** (the distribution) — determined by everything.

5.4 Grade-3 Interpretation and the Hierarchy

The grade-3 Latent is described here as a useful structural choice. The following criterion is a heuristic interpretation rather than an established minimax result.

A heuristic criterion. For each candidate representation, one should specify the complete numerical operator and measure its stability over the parameter regime of interest. The present evidence is narrower: the moment coefficients have the growth described in §1, the referenced Padé implementation fails in selected high-volatility tests, and the direct GH/COS evaluator has a bounded integrand and the empirical errors reported in §7. None of these facts supplies a general condition-number formula for either representation.

Why the hierarchy terminates.

Each grade answers a different question:

Grade	Question	Type	Determined by
0	What are the parameters?	**Data**	Given by the user
1	What is the distribution?	**WHAT**	Parameters + basis
2	How hard is it to extract?	**HOW**	Parameters + basis + precision
3	In what language do you describe it?	**WHERE**	A basis selected using the generative structure

The *generative structure* is given, but the grade-3 representation selected in response to it remains a choice. For the present problem, Gaussian input motivates testing Hermite coordinates; it does not logically force that choice or exclude competing representations.

As a design heuristic, a possible grade-4 question is: "given a family of problem classes, which grade-3 representation should be tried first?" The **Representation–Structure Matching Heuristic** (§5.2) recommends testing bases whose measure or geometry reflects the generative structure:

Problem class	Generative structure	$\mathcal{B}^*$ (grade-3)	Matching principle (grade-4)
Lognormal sums	Gaussian $\to$ exponential	Hermite	Orthogonality measure = Gaussian weight
Polynomial chaos (general X)	Distribution of X	Askey scheme	Orthogonality measure = distribution of X
PDE with analytic data	Smoothness/analyticity	Chebyshev/Fourier	Orthogonality measure matches domain

The repeated rule is: *prefer a basis adapted to the probability measure, periodicity, domain, and regularity.* Orthogonality facts such as Hermite/Gaussian matching are theorems; the stronger claim that the matched basis is computationally optimal is not. Basis choice, numerical method, and precision therefore remain coupled design decisions.

The sequence of grades is:

$$\underbrace{\text{Structural heuristic}}_{\text{problem-adapted}} \to \underbrace{\text{Basis choice}}_{\text{grade 3}} \to \underbrace{\text{Numerical behavior}}_{\text{grade 2, method-dependent}} \to \underbrace{\text{Latent}}_{\text{grade 1}} \to \underbrace{F_S(x)}_{\text{observable}}$$

Grade 3 is often a high-leverage choice because it strongly influences the downstream numerical problem. The hierarchy is interpretive: it organizes design choices but is not a theorem that representation alone determines condition numbers or feasibility.

5.5 Why Moments Were Used (Historical)

The moment representation was a natural starting point because moments are standard distributional summaries and admit a closed formula here (Nagy, 2026, *The Exact Latent Distribution of Correlated Lognormal Sums*, Theorem 1). They do not uniquely characterize the lognormal law among all positive distributions. For the Padé evaluation chain studied here, the moment basis is also poorly matched to the generative Gaussian structure.

The Smooth Latent Operator (Nagy, 2026, *The Smooth Latent Operator: Parameter-Free Distributional Representations via Kernel Moment Recovery*) attempted to fix the grade-2 problem (extraction difficulty) without addressing the grade-3 problem (basis choice). This is analogous to using adaptive-precision arithmetic to evaluate a poorly converging Taylor series instead of switching to the Fourier representation — optimizing HOW you compute in the wrong language, instead of switching to the right language.

---

6. Comparison: Moment Latent vs. Hermite Latent

6.1 Theoretical Comparison

The preceding construction can now be compared directly with the moment representation.

Property	Moment Latent Λ^M	Hermite Latent Λ^H
Space	$\mathcal{H}_{\beta_{\mathrm{wt}}}$ (Gaussian weight required)	ℓ^2 (no weight)
Coefficient growth	$e^{\sigma^2 k^2/2}$ (super-exponential)	$\sigma^k/k!$ (factorial decay)
CF recovery	Padé resummation (divergent $\to$ rational)	Inner product or quadrature (numerically tested)
Conditioning	Severe in referenced Padé tests	Bounded integrand; full conditioning not proved
High-σ regime	Fails in the referenced examples for $\sigma > 0.8$	Tested in selected cases
Negative weights	Empirical only	Outside the positive-sum evaluator
Computational cost	$O(N_P^3)$ (dense Toeplitz solve)	$O(NQ^n)$ nodes/terms for direct tensor GH-COS, excluding per-node work

6.2 The Cost: Dimensionality

Naturally, the coin has two sides — the Hermite Latent's weakness is its **dimensionality**. While the moment Latent is indexed by a single integer k (a 1D sequence), the Hermite Latent is indexed by a multi-index $\mathbf{k} \in \mathbb{N}_0^n$ (an n-dimensional lattice).

For n assets and a truncation degree K:

- The moment Latent requires $2N_P$ coefficients.
- The Hermite Latent requires $\binom{K+n}{n}$ coefficients.

n	K	$\binom{K+n}{n}$
2	20	231
5	20	53,130
10	15	3,268,760
20	10	30,045,015
50	8	1,916,797,311

For $n > 10$, the multi-index enumeration becomes expensive. This is the "curse of dimensionality" inherent to polynomial chaos methods.

6.3 Sparse Hermite Latent

The curse can sometimes be mitigated by exploiting sparsity in the chosen factor L. Sparse Σ does not by itself imply sparse Cholesky factors because elimination can create fill-in.

Proposition 2 (factor-support sparsity). *Suppose each row ℓ_i of L has support size at most s. A component contribution to $c^H_{\mathbf{k}}$ can be nonzero only when*

$$\operatorname{supp}(\mathbf{k}) \subseteq \operatorname{supp}(\ell_i)$$

for some i. Hence the number of potentially nonzero coefficients at exact degree K is at most

$$n\binom{K+s-1}{s-1}.$$

This is an upper bound; overlapping supports and cancellation can reduce the count further.

For a genuinely rank-r covariance $\Sigma = BB^T$, the Hermite expansion can be written in r factor variables. For the common factor-plus-idiosyncratic form $\Sigma = BB^T + D$ with nonzero diagonal D, the idiosyncratic variables remain; reduction to r dimensions then requires an additional conditional approximation and is not automatic.

---

7. The CF Computation: Detailed Methods

7.1 Method A: Hermite Moment Transfer

Given the Hermite coefficients $\{c^H_{\mathbf{k}}\}$, the moments of S_K are:

$$E[S_K^p] = \sum_{\mathbf{k}_1,\ldots,\mathbf{k}_p} \prod_{r=1}^{p} c^H_{\mathbf{k}_r} \cdot E\left[\prod_{r=1}^{p} H_{\mathbf{k}_r}(Z)\right]$$

The Gaussian expectations of products of Hermite polynomials can be computed by the Hermite product formula, equivalently by expanding the finite polynomial and applying Wick/Isserlis contractions to its Gaussian monomials. They do not reduce, for products of three or more Hermite factors, to pairwise equality indicators between the multi-indices.

The resulting moments of S_K can feed the Padé-COS chain of Nagy (2026, *The Exact Latent Distribution of Correlated Lognormal Sums*). Factorial decay of the input Hermite coefficients does not imply that the derived high-order raw moments or their Padé Toeplitz system are well-conditioned; no such guarantee is claimed for this auxiliary route.

7.2 Method B: Gauss-Hermite Quadrature

For moderate n (say $n \leq 8$):

$$\phi_{S_K}(t) = \int_{\mathbb{R}^n} e^{it\, p_K(z)} \gamma_n(z) dz \approx \sum_{\ell=1}^{Q^n} q_\ell \, e^{it\, p_K(z_\ell)}$$

where (z_ℓ, q_ℓ) are n-dimensional Gauss-Hermite quadrature nodes and weights. Because the integrand is smooth and the Gaussian weight is the natural measure, the quadrature is typically observed to converge rapidly as the number of nodes per dimension Q increases; the present paper does not establish a uniform exponential rate for this non-polynomial oscillatory integrand.

For $n = 5$ and $Q = 15$, the tensor rule uses $15^5 = 759{,}375$ evaluations. Runtime depends on implementation, vectorization, and hardware; no timing claim is made here.

This is technically a numerical quadrature, but it is:

1. **Rapidly convergent in selected tests** (rather than sampling-noise-limited like Monte Carlo).
2. **Deterministic** (no sampling noise).
3. **Structured** (the nodes and weights are precomputed).

Gauss-Hermite is exact only for polynomial integrands up to degree $2Q - 1$. The oscillatory integrand $e^{itp_K(z)}$ is not a polynomial, so the displayed finite sum remains a quadrature approximation even though finite Taylor truncations can be integrated exactly at sufficiently high order.

7.3 Method C: Hybrid (Hermite → Padé → COS)

The most elegant approach combines the Hermite Latent with the Padé-COS chain:

1. **Compute** the Hermite coefficients $\{c_{\mathbf{k}}^H\}_{|\mathbf{k}| \leq K}$ (closed form).
2. **Derive** the moments of S_K from the Hermite coefficients using the Wick theorem (Method A).
3. **Apply** the Padé-COS chain (Nagy, 2026, *The Exact Latent Distribution of Correlated Lognormal Sums*) to these Hermite-derived moments.

The key insight: the Hermite-derived moments of S_K are NOT the same as the raw moments $m_k = E[S^k]$. They are:

For each fixed p for which the required uniform-integrability estimate is controlled, $E[S_K^p] \to E[S^p]$ as $K \to \infty$. The L^2 estimate in Theorem 2 alone does not provide the displayed factorial error rate for arbitrary p. Nor does convergence of these moments imply that a Padé Toeplitz system built from them is well-conditioned.

7.4 Three Routes from Latent to CDF Formula

The Hermite Latent Λ^H determines S as an $L^2(\gamma_n)$ random variable and therefore determines its distribution. Pointwise CDF evaluation is not, however, a continuous functional on generic L^2 random variables or on the corresponding ambient coefficient space: $X_r = 1/r$ converges to $X = 0$ in L^2, while $P(X_r \le 0) = 0$ and $P(X \le 0) = 1$.

A finite truncation has the weaker regularity actually needed here. Fix a finite index set and define

$$p_c(Z) = \sum_{\mathbf{k} \in \mathcal{I}_K} c_{\mathbf{k}} H_{\mathbf{k}}(Z).$$

Suppose the coefficient vectors converge to c in the Euclidean topology and the induced law has no atom at x. The corresponding polynomials then converge in L^2, hence in probability, and

$$P(p_{c^{(r)}}(Z) \le x) \longrightarrow P(p_c(Z) \le x)\,.$$

Thus the finite-truncation CDF map is continuous at coefficient vectors for which x is a continuity point of the induced law; no Fréchet smoothness on ℓ^2 is asserted. Separately, for fixed (N, Q) the *unclipped* finite evaluator $\mathcal{H}^n_{N,Q}$ is C^∞ in its finite inputs on the open domain

$$x > 0, \qquad \beta > \alpha, \qquad w_i > 0, \qquad S(z_\ell) > 0 \text{ for every quadrature node.}$$

Here the variables are $(x, w, \mu, L, \alpha, \beta)$ with their Euclidean topology. If $L = \mathrm{chol}(\Sigma)$, the same conclusion holds with Σ restricted to the open positive-definite cone because the Cholesky map is smooth there. Output clipping and data-dependent domain-selection rules can introduce nonsmooth switching points. This section now examines three routes from the representation to a CDF approximation.

Route A: Direct Perturbation (Edgeworth-type) — NOT JUSTIFIED HERE

The natural idea: decompose $S = S_1 + R$ where $S_1 = \mu_S + \sigma_S \tilde{Z}$ is the Gaussian part (degree-1 Hermite) and $R = \sum_{|\mathbf{k}|\ge 2} c^H_{\mathbf{k}} H_{\mathbf{k}}(Z)$ is the correction. Then Taylor-expand $F_S(x) = E[N(z - R/\sigma_S)]$ where $z = (x - \mu_S)/\sigma_S$:

$$F_S(x) \stackrel{?}{=} N(z) - \phi(z) \sum_{j=1}^{\infty} \frac{1}{j!} He_{j-1}(z)\, E\left[\left(\frac{R}{\sigma_S}\right)^j\right]$$

The correction R is a polynomial of degree $K \ge 2$ in Gaussians. Hypercontractivity (Nelson, 1973) gives the upper bound

$$E\left[\left|\frac{R}{\sigma_S}\right|^j\right] \le (j-1)^{jK/2} \left(\frac{\mathrm{Var}(R)}{\sigma_S^2}\right)^{j/2}$$

Because this is an upper bound, its growth cannot prove divergence. It does show that hypercontractivity alone is insufficient to justify termwise expectation or convergence of the proposed expansion. Classical Edgeworth expansions are generally asymptotic rather than globally convergent, and the rapidly growing projected moments of lognormal sums provide an additional warning. This route is therefore not used by the evaluator.

Lesson: The multi-dimensional Hermite basis is useful for representing and analysing the Gaussian latent, while the one-dimensional Fourier-cosine basis is useful for numerical inversion in log-space. The implemented route evaluates the positive function $S(z)$ directly at GH nodes; it does not take the logarithm of a finite Hermite truncation, which need not remain positive.

Route B: The Direct Log-CF COS Evaluator

For fixed numerical settings define the finite evaluator $\mathscr{H}^n_{N,Q}$:

$$\mathscr{H}^n_{N,Q}(x;\ w,\mu,L,[\alpha,\beta]) := u + \sum_{k=1}^{N} \frac{2}{k\pi}\sin(k\pi u) \times \sum_{\ell=1}^{Q^n} q_\ell \cos\left(k\pi \frac{\log S(z_\ell)-\alpha}{\beta-\alpha}\right),$$

where $u = (\log x - \alpha)/(\beta - \alpha)$, q_ℓ are the GH weights, and $S(z_\ell) = \sum_i w_i \exp(\mu_i + \ell_i^T z_\ell) > 0$.

The numerical approximation used in the experiments is

$$\boxed{F_S(x) \approx \mathscr{H}^n_{N,Q}(x;\ w,\mu,L,[\alpha,\beta]).}$$

This is a deterministic finite formula for fixed settings, not an exact finite CDF. Its errors include Gaussian quadrature, COS truncation, and finite-domain truncation; the present paper does not establish a universal end-to-end bound.

The log-transform maps the positive heavy-tailed variable to a lighter-tailed coordinate and performs well in the reported tests, including cases with $\sigma > 1$. This is an empirically supported design choice rather than a proved necessity or universal convergence statement.

Comparison with standard special functions:

Special function	Definition	Elementary?	"Closed-form"?
$N(x) = \frac{1}{\sqrt{2\pi}}\int_{-\infty}^{x} e^{-t^2/2}dt$	Integral	No (Liouville)	Yes (by convention)
$\chi^2_\nu(x) = \frac{\gamma(\nu/2,x/2)}{\Gamma(\nu/2)}$	Incomplete gamma	No	Yes
$\Phi_n(x;\mu,\Sigma)$ (multivariate normal CDF)	n-fold integral	Not at issue here	Standard numerically evaluated function
$\mathscr{H}^n_{N,Q}(x;w,\mu,L,[\alpha,\beta])$	Finite numerical sum of log, sin, cos	**Yes, for fixed settings**	Approximation

The evaluator is a finite expression built from elementary functions after its numerical settings are fixed. This does not establish an elementary expression for the underlying exact CDF or remove the approximation error.

Route C: The Quadratic Special Case ($K = 2$)

For $K = 2$, the truncation S_2 is a quadratic form in Gaussian variables plus a linear term:

$$S_2 = \mu_0 + \mathbf{a}^T Z + Z^T A Z$$

where:

- $\mu_0 = c^H_{\mathbf{0}} - \sum_{j=1}^{n} c^H_{2e_j}$ (constant)
- $a_j = c^H_{e_j}$ (linear coefficients)
- $A_{jj} = c^H_{2e_j}$, $A_{jk} = c^H_{e_j+e_k}/2$ for $j \neq k$ (quadratic form matrix)

Eigendecompose $A = UDU^T$, set $W = U^T Z \sim N(0, I)$, and put $b = U^T \mathbf{a}$. For every nonzero eigenvalue,

$$d_j W_j^2 + b_j W_j = d_j \left(W_j + \frac{b_j}{2d_j} \right)^2 - \frac{b_j^2}{4d_j}.$$

Zero-eigenvalue directions retain a linear Gaussian term. Thus S_2 is a generalized noncentral Gaussian quadratic form. Its distribution can be evaluated by standard characteristic-function inversion methods such as Imhof's; a weighted sum with unequal or negative d_j is not, in general, one ordinary noncentral chi-squared variable.

For a single asset, the exact relative root-mean-square truncation error at $K = 2$ is

$$\left(e^{-\sigma^2} \sum_{k>2} \frac{\sigma^{2k}}{k!} \right)^{1/2}.$$

It equals approximately 0.0107, 0.0465, 0.165, and 0.420 for $\sigma = 0.3$, 0.5, 0.8, and 1.2, respectively. These are random-variable errors, not CDF-error estimates; no universal volatility threshold for adequacy is asserted.

Summary of Routes

Route	Formula type	Convergent?	Practical?
A: Edgeworth	$N(z) + \phi(z) \sum \Delta_p$	Not established	Not used
B: Direct log-CF COS $\mathcal{H}^n_{N,Q}$	Double finite sum	Tested numerically	Selected positive-weight cases
C: Quadratic ($K = 2$)	Generalized quadratic-form inversion	Exact for S_2	Low-order diagnostic

Route B is the proposed practical route. It avoids the moment-series route used by Edgeworth-type approximations and evaluates $S(z)$ directly. In the reported two-asset high-volatility tests, log-space COS achieves 10^{-4}-scale discrepancies against Monte Carlo with Q between 60 and 120 nodes per dimension.

Route C could provide an independent low-order diagnostic if a generalized quadratic-form inversion is implemented; that comparison is not part of the current validation script.

7.5 Numerical Validation

The direct log-CF COS evaluator ($\mathcal{H}^n_{N,Q}$, using direct $S(z)$ evaluation at GH nodes) is compared with Monte Carlo generated with NumPy seed 42 across eight test cases. Cases 1–7 use 5M samples and Case 8 uses 2M. All computations use standard float64 arithmetic. Cases 1–5 and 8 have positive weights and lie within the stated method; Cases 6–7 are explicitly outside-scope signed experiments.

Case	σ	Weights	(Q, N)	Nodes	Max $\lvert F - F_{\mathrm{MC}}\rvert$
1. Baseline ($\rho = 0.3$)	0.2	$[0.5, 0.5]$	$(60, 20)$	3600	9.14×10^{-4}
2. High corr ($\rho = 0.95$)	0.2	$[0.5, 0.5]$	$(60, 20)$	3600	9.03×10^{-4}
3. High volatility	0.8	$[0.5, 0.5]$	$(80, 26)$	6400	2.64×10^{-4}
4. Extreme test regime	1.2	$[0.3, 0.7]$	$(100, 33)$	10000	4.07×10^{-4}
5. Ultra-high volatility	2.0	$[0.5, 0.5]$	$(120, 40)$	14400	2.04×10^{-4}
6. Mixed-sign experiment (outside scope)	0.2	$[1, -0.5]$	$(60, 20)$	3600	2.88×10^{-3}
7. Near-cancellation experiment (outside scope)	0.2	$[1, -0.99]$	$(60, 20)$	3600	2.20×10^{-3}
8. Five-asset ($n = 5$)	0.25	$[0.2]^5$	$(12, 166)$	248832	6.34×10^{-4}

The reported GH nodes reproduce the analytical mean $E[S] = \sum w_i e^{\mu_i + \sigma_i^2/2}$ to machine precision ($\sim 10^{-16}$) in these tests.

Monte Carlo uncertainty. The MC standard error for CDF estimation with N samples is approximately $\sqrt{p(1-p)/N}$, which for $p \approx 0.5$ and $N = 5\text{M}$ gives $\sigma_{\mathrm{MC}} \approx 2.2 \times 10^{-4}$. The smallest reported errors (Cases 3, 5: $\sim 2 \times 10^{-4}$) are comparable to this noise floor; the agreement demonstrates consistency rather than demonstrated sub-10^{-4} accuracy. The larger errors in Cases 1–2 ($\sim 9 \times 10^{-4}$) and Cases 6–7 ($\sim 10^{-3}$) are well above the MC uncertainty and reflect genuine approximation error from truncation and quadrature parameters.

Key observations:

1. **Cases 3–5 (high-volatility regime)**: In these particular experiments, the high-σ cases show lower discrepancies against Monte Carlo than the low-σ baseline. Because the discrepancies

are partly at the Monte Carlo noise floor and the evaluator settings differ, this is evidence of robustness in the tested cases, not proof that error decreases with volatility.

2. **Log-space performed better in these tests.** The same GH tensor evaluated through direct-space COS gave max errors above 10^{-2} in the reported comparisons, while log-space COS reached approximately 10^{-4} to 10^{-3}. This experiment does not establish that every direct-space implementation must fail.

3. **Coupled convergence in Q and N.** At volatility 1.2, the observed transition is

$$(Q, N) : (40, 13) \longrightarrow (60, 20), \qquad \max |F - F_{\mathrm{MC}}| : 1.7 \times 10^{-2} \longrightarrow 8.8 \times 10^{-4}.$$

   The next three pairs—$(80, 26)$, $(100, 33)$, and $(120, 40)$—plateau near 2×10^{-4}. This is an empirical coupled-refinement observation, not a pure-Q convergence theorem.

4. **Hedge portfolios (Cases 6–7).** The displayed log-CF construction requires $S > 0$. For mixed-sign weights, a globally finite shift based on $\min_z S(z)$ is generally unavailable, so signed portfolios are an open separate problem and the reported shifted results are not covered by the positive-sum evaluator.

5. **All computations in float64.** No arbitrary precision arithmetic is needed, unlike the moment-Padé approach which required 50–500 digit precision for $\sigma > 0.5$.

One-asset exact diagnostic. For $n = 1$, $w = 1$, and $\mu = 0$, the identity $\log S = \sigma Z$ makes the exact CDF available. With $Q = 60$, the maximum error over the same nine probabilities is independent of σ in the tested set $\{0.2, 0.8, 1.2, 2.0\}$:

N	8	20	30	40
Max CDF error	5.06×10^{-2}	6.79×10^{-4}	3.00×10^{-6}	3.89×10^{-9}

This rejects the former default $N = 8$ even in the simplest exact case. It also shows why increasing N without checking whether Q resolves the required frequencies is unsafe.

The tracked reproduction script is forge/fin_fenton_solved/demo_hermite_cos.py (SHA-256 ff3f4a8016fbd8b7598fde19241aea5eb0c28302f4bbfab92fcc9ad2226bdcdd). It contains the evaluator, all eight case definitions, the exact one-asset diagnostic, coupled convergence studies, and the Monte Carlo harness, and requires NumPy and SciPy.

7.6 Implementation Recipe

The following pseudocode is a recommended coupled convergence diagnostic. The tracked script supplies explicit orders and reports fixed convergence grids; it does not implement this protocol as an automatic selector or stopping rule.

```
ALGORITHM: Log–CF/COS CDF with coupled refinement
INPUT:   w      ,        , Σ     S, diagnostic grid G     ,
         candidate orders Q<...<Q_J and N<...<N_H, tolerance
OUTPUT:  stabilized CDF values on G, or "unresolved"

1. CHOLESKY:   L ← cholΣ()                          // O(n³), once
```

```
2. FOR each Q_j:
     construct the Q_j^n GH nodes and evaluate S_ and X_=log S_
     set [_j ,_j] from the declared domain rule
3. FOR each N_h at this fixed Q_j:
     compute A_k for k=0,...,N_h and raw COS values H_{Q_j,N_h} on G
     record Δ_COS = max_G |H_{Q_j,N_h}−H_{Q_j,N_{h−1}}|
4. When the same N_h is available at Q_{j−1}, recordΔ
    _GH = max_G |H_{Q_j,N_h}−H_{Q_{j−1},N_h}|
5. ACCEPT a pair only after both Δ _COS and Δ _GH on successive refinements;
     also inspect range, monotonicity, and density−sign diagnostics
6. If the checks fail or later orders destabilize, enlarge/revise the order grid
     and report "unresolved" rather than returning a certified value
7. Clamp only for presentation after the raw convergence diagnostics
```

The two successive differences are diagnostics, not rigorous error estimators: simultaneous aliasing can create a false plateau. A production implementation should enlarge the diagnostic grid, include the intended tail region, and compare against an independent method when material decisions depend on the result. No universal stopping guarantee is claimed.

Settings used in the validation. The eight explicit (Q, N) pairs are listed in §7.5 and hard-coded in the tracked script. The one-asset exact diagnostic uses $Q = 60$ and $N \in \{8, 20, 30, 40\}$. Errors are evaluated at probabilities 0.01, 0.05, 0.10, 0.25, 0.50, 0.75, 0.90, 0.95, and 0.99, using seed 42 for Monte Carlo cases. These reproduce the tables; they are not universal defaults.

The tensor implementation scales exponentially in n and is not validated beyond the displayed five-asset case. Sparse-grid or conditional factor methods are possible extensions, but no Cholesky-sequential complexity reduction is established here. The present log-CF evaluator is restricted to strictly positive weights; signed portfolios require separate treatment.

---

8. VaR, ES, and Risk Measures

The CDF approximation (§4.4) provides a deterministic numerical route to the two risk measures considered here.

8.1 Value-at-Risk

The VaR at confidence level α is the quantile:

$$\mathrm{VaR}_\alpha(S) = F_S^{-1}(\alpha) = \inf\{x > 0 : F_S(x) \geq \alpha\}$$

Because F_S is continuous and strictly increasing on $(0, \infty)$ for the non-degenerate positive-weight setting, the inverse exists and is unique. The computation evaluates the approximation on a grid, brackets the root, and refines it using bisection or Newton's method. Each baseline tensor evaluation costs $O(NQ^n)$ after node values are prepared; vectorized reuse can lower constants when evaluating many x values.

For a bracketed quantile, bisection adds a logarithmic number of CDF evaluations to reach a requested x-tolerance. Concrete runtime depends on vectorization, reuse of coefficients, dimension, and hardware; no universal sub-millisecond claim is made.

8.2 Expected Shortfall

The ES at level α is:

$$\mathrm{ES}_\alpha(S) = \frac{1}{1-\alpha}\int_\alpha^1 \mathrm{VaR}_u(S)\,du = \frac{1}{1-\alpha}\int_{\mathrm{VaR}_\alpha}^\infty x\,f_S(x)\,dx$$

Using the COS density $f_S(x)$ from Section 4.5 and defining $q_\alpha := \mathrm{VaR}_\alpha(S)$, we have:

$$\mathrm{ES}_\alpha(S) = \frac{1}{1-\alpha}\left(E[S] - \int_0^{q_\alpha} x\,f_S(x)\,dx\right)$$

Here, $E[S] = \sum_{i=1}^n w_i e^{\mu_i+\sigma_i^2/2}$ is known exactly in closed form. The integral is then evaluated by Gauss-Legendre quadrature on the interval $[0, q_\alpha]$ using the COS density. This represents a single additional quadrature step atop the CDF computation.

No direct COS partial-moment formula is used in the reported experiments. Such a formula must preserve the density normalization, including the factor of two on nonzero cosine modes. Because ES was not benchmarked here, the paper retains the explicit Gauss–Legendre route above as a computational proposal rather than claiming a validated direct-COS ES evaluator.

8.3 Practical Accuracy

CDF approximation error propagates to VaR and ES, but those risk measures were not independently benchmarked in §7.5. A local first-order quantile perturbation is

$$|\mathrm{VaR}_\alpha^{\mathrm{exact}} - \mathrm{VaR}_\alpha^{\mathrm{approx}}| \approx \frac{|F_S - \hat{F}_S|}{f_S(\mathrm{VaR}_\alpha)} \sim \frac{10^{-4}}{f_S(\mathrm{VaR}_\alpha)}$$

The conversion from CDF error to quantile error depends on the local density and can deteriorate in a thin tail. The experiments in §7.5 do not establish a regulatory accuracy guarantee for VaR or ES. ES error likewise requires a separate tail-integral analysis and is not automatically smaller than pointwise CDF error.

---

9. When Hermite Fails: The Curse of Dimensionality

9.1 The Fundamental Tradeoff

In the reported experiments, the Hermite route replaces the observed instability of the moment-Padé route with a tensor-dimensionality cost. A practical comparison must account for both:

$$\binom{K+n}{n} \ll \text{precision required for moment Padé}$$

For $n = 5$, a total-degree truncation at $K = 25$ contains $\binom{30}{5} = 142{,}506$ coefficients. This is a combinatorial illustration only. No necessity claim is made because a required K depends on weights, means, covariance, target norm or observable, and tolerance.

For $n = 50$, even $K = 8$ contains $\binom{58}{8} = 1{,}916{,}797{,}311$ coefficients. This makes unstructured enumeration impractical, without establishing that a particular competing method is uniformly preferable.

9.2 Factor-Model Reduction

For large n, the covariance Σ typically has low-rank structure (factor model):

$$\Sigma = BB^T + D, \quad B \in \mathbb{R}^{n\times r}, \quad D = \mathrm{diag}$$

with $r \ll n$ common factors:

$$S = \sum_i w_i e^{\mu_i + d_i \eta_i + (B^T)_i \cdot F}$$

where $F \sim N(0, I_r)$ are common factors and the η_i are independent idiosyncratic normals. Conditional on F, the terms are independent, but the distribution of their sum still requires computation. Expanding only in F does not eliminate the idiosyncratic integrations. A rigorous and efficient conditional factor evaluator is therefore an extension, not a result of the current tensor implementation.

9.3 Candidate Representation Choices

The grade-3 decision is now nuanced:

Regime	Candidate representation	Rationale
n small, σ large	Hermite Latent	Bounded integrand; selected cases tested
n large, σ small	Moment Latent + Padé	Low-dimensional coefficient sequence; selected cases tested
n large, σ large, low-rank Σ	Conditional factor method	Potential dimension reduction; not validated here
n large, σ large, full-rank Σ	Open problem	Neither method scales well

This table is a qualitative decision guide, not an optimality theorem. Dimension, volatility, covariance structure, target accuracy, and implementation all affect the choice.

10. Implications for the Latent Framework

10.1 Representation Is Not Neutral

The original Latent framework (Nagy, 2026, *The Latent: Finite Sufficient Representations of Smooth Systems*) treats the Latent as a fundamental property of the system — "the" finite representation. However, the present paper shows that the representation is, in fact, a CHOICE. Furthermore, different choices lead to radically different practical performance.

The Latent Theorem applies only after exponential compressibility has been established in a specified extraction basis. The two coordinate systems here lie on opposite sides of that hypothesis:

- The scaled moment coefficients do not decay; they grow at a quadratic-exponential rate. The theorem's condition $\rho_{\mathcal{B}} > 1$ does not apply to these formal MGF coordinates.
- The Hermite coefficients decay factorially and hence supergeometrically: for every fixed finite $\rho > 1$, they eventually satisfy a bound of order ρ^{-K}. Assigning one finite "$\rho \gg 1$'' is therefore unnecessary.

Consequently, practical truncation behavior can depend strongly on basis choice. This contrast does not turn generic C^∞ smoothness into exponential decay, and no formal optimization over all bases is carried out.

10.2 The Hierarchy Principle

The graded Latent hierarchy is:

$$\mathcal{B}^* \quad \xrightarrow{\text{determines}} \quad \alpha^*(\mathcal{B}^*) \quad \xrightarrow{\text{determines}} \quad \Lambda(\alpha^*, \mathcal{B}^*) \quad \xrightarrow{\text{determines}} \quad F_S(x)$$

Each grade determines the next. The total system is:

$$F_S(x) = \mathrm{COS}\left(\mathrm{CF}\left(\Lambda_{\alpha^*}^{\mathcal{B}^*}\right)\right)$$

The grade-3 choice is the highest-leverage decision: it determines whether the entire chain is feasible.

10.3 Analogy to Coordinate Selection in Physics

The grade-3 Latent is analogous to the choice of coordinates in physics:

- Cartesian coordinates for the two-body problem $\to$ complicated.
- Polar coordinates $\to$ separable, exactly solvable.
- The right coordinates make the physics transparent.

Similarly:

- Moments (monomial coordinates) for lognormal sums $\to$ divergent.
- Hermite coordinates $\to$ convergent, tractable.
- The right representation makes the distribution transparent.

The Latent framework uses this as an organizing heuristic: candidate coordinates can be proposed from the generative structure, then accepted or rejected by problem-specific analysis and numerical evidence.

10.4 Related Work and Novelty Delineation

The evaluator combines established components. A targeted OpenAlex/Crossref search found close quadrature-based lognormal-sum methods, so this section records scope and overlap rather than claiming a new implementation class.

Known components (prior art):

Component	Origin	Used how
COS method (CF $\to$ CDF via Fourier-cosine)	Fang & Oosterlee (2008)	Standard step 3 of our chain
Hermite polynomial chaos (PCE)	Wiener (1938), Cameron-Martin (1947), Xiu-Karniadakis (2002)	Expansion of e^Y in Hermite polynomials of Z
Gauss-Hermite quadrature for lognormal CFs	Gubner (2006)	Direct precedent for GH evaluation of a lognormal characteristic function
Quadrature CF/CDF approximations for independent and correlated lognormal sums	Mahmoud (2010)	Direct GH precedent for correlated-sum CF/CDF approximation
High-accuracy CF/MGF computation and CDF inversion for independent lognormal sums	Senaratne & Tellambura (2009, 2010)	Alternative transform-evaluation and inversion route
Closed-form lognormal moments	Schwartz-Yeh (1982), Fenton (1960)	Not used here (replaced by Hermite)
Padé approximation for CF	Baker-Graves-Morris (1996)	Not used here (replaced by GH)

Scope evaluated in this paper:

1. **An explicit Hermite-chaos analysis for correlated lognormal sums.** The closed-form coefficients make the Gaussian-latent structure explicit and provide useful L^2 truncation estimates; this is an analytical specialization, not a claim to have invented Hermite chaos.

2. **A tested log-coordinate evaluator.** The studied path evaluates the CF of $T = \log S$ as $E[S^{it}]$ at tensor GH nodes and applies COS to T, avoiding moment resummation and finite Hermite truncation in this implementation.

3. **Monte Carlo comparison in stated low-dimensional cases.** The seeded comparisons cover positive-weight two-asset volatilities through $\sigma = 2$ and one positive-weight five-asset case, plus two outside-scope signed experiments.

4. **A representation-matching interpretation.** The graded Latent discussion frames basis selection as a high-level numerical design decision.

These items describe the paper's analytical specialization, implementation, and interpretation. Mahmoud (2010) is direct prior art for quadrature-based CF and CDF approximation of correlated lognormal sums. The present study does not implement or benchmark Mahmoud's method, so it makes no empirical superiority or distinctiveness claim against that predecessor. It also makes no claim that GH-based lognormal-sum CDF evaluation, transform inversion, the log-coordinate composition, or any individual numerical ingredient was previously unknown.

What this paper is NOT claiming:

- We do not claim to have invented PCE, COS, or GH quadrature.
- The COS inversion is the Fang-Oosterlee (2008) method, applied in log-coordinates.
- The Hermite coefficients follow from Cameron-Martin (1947) theory, specialized to the lognormal case.
- The numerical evidence establishes only the reported evaluator-versus-Monte-Carlo discrepancies; it does not establish an improvement over Mahmoud (2010) or another literature method.

---

11. Conclusion

This paper provides three concrete deliverables for the lognormal sum problem:

1. **A deterministic log-CF quadrature/COS scheme** (§4.4): the log-space characteristic function $\tilde{\phi}(t) = E[S^{it}]$ is approximated by Gauss–Hermite quadrature and inverted by Fourier-cosine synthesis. For fixed evaluator settings, the resulting approximation is a double finite sum of elementary functions.

2. **Numerical evidence** (§7.5): in the six positive-weight, positive-definite cases within the method's stated domain, maximum CDF discrepancies against Monte Carlo range from 2.04×10^{-4} to 9.14×10^{-4}, including one case with $\sigma = 2.0$. The two outside-scope signed experiments have discrepancies of 2.88×10^{-3} and 2.20×10^{-3}. These results use explicit case-specific (Q, N) pairs and provide limited empirical validation rather than a universal accuracy guarantee.

3. **A structural interpretation** (§5): matching a representation basis to the problem's generative structure is a useful heuristic. For lognormal sums (Gaussian $\to$ exponential $\to$ sum), Hermite is a natural candidate.

A major obstruction in the classical moment/MGF evaluation route is representational: it uses the formal Taylor data associated with $M(z) = E[e^{zS}]$, whose positive-order coefficients grow at a quadratic-exponential rate and whose Taylor series has zero radius. The complete moment sequence also does not uniquely determine a lognormal law among all positive distributions. The Hermite-chaos expansion instead uses coordinates adapted to the underlying Gaussians, yielding coefficients that decay factorially ($\sigma^k/k!$). In the computational pipeline, the log-space complexification S^{it} (bounded: $|S^{it}| = 1$ for $S > 0$) replaces the real exponential e^{tS} (unbounded), converting a divergent contraction into a convergent one. This avoids that specific obstruction; it does not remove all numerical difficulty in the full evaluator.

The practical consequence is that the direct log-CF COS pipeline performs well in the reported high-volatility test cases ($\sigma > 0.8$), where the referenced moment-Padé implementation encounters singular matrices. Whether this advantage extends across broader portfolios, dimensions, and evaluator settings remains open.

11.1 Is This an Analytical Solution?

The implemented direct log-CF COS evaluator $\mathcal{H}^n_{N,Q}$ involves the precision parameters N and Q and a log-space domain $[\alpha, \beta]$. It is a deterministic numerical scheme built from an exact integral representation.

Four standards of "analytical":

Standard	Criterion	$\mathcal{H}^n_{N,Q}$?
Elementary closed-form	Finite composition of $+, -, \times, /, \exp, \log, \sin$ with no free parameters	Not supplied
Named special functions	Expressed via existing $N(x), \Gamma, {}_2F_1, \ldots$	Not supplied
Deterministic numerical representation	Explicit parameterized quadrature/COS evaluation	**Yes, in the tested setting**
Fixed-setting finite evaluator	Finite arithmetic expression after Q, N, and the log-domain are supplied	**Yes**

Comparison. The standard normal CDF $N(x) = \frac{1}{\sqrt{2\pi}} \int_{-\infty}^{x} e^{-t^2/2} dt$ is defined by an integral with no elementary expression and is evaluated numerically in practice (Abramowitz and Stegun, 1964). Likewise, $\mathcal{H}^n_{N,Q}$ is a finite numerical approximation after its settings are fixed. Its elementary finite-sum form neither proves nor disproves the existence of a different exact closed form for the target CDF.

What the scheme achieves:

- A tested log-space GH-COS evaluation route for selected positive-weight, low-dimensional examples.
- A finite-sum numerical approximation for fixed evaluator settings.
- 10^{-4}-scale agreement with Monte Carlo in the reported Bitcoin-regime examples.
- Each CDF evaluation is $O(N \times Q^n)$ arithmetic operations; numerical VaR still uses root-finding and ES may use additional quadrature.

What the scheme does not achieve:

- No elementary or standard-special-function closed form is supplied. The lognormal-sum literature generally relies on approximations, bounds, quadrature, or transform inversion (for example, Fenton, 1960; Senaratne and Tellambura, 2009, 2010; Mahmoud, 2010). This state of the literature is not an impossibility theorem.
- A proved parameter-selection or evaluator-error guarantee: the log-domain $[\alpha, \beta]$ and orders must be chosen heuristically, and the stated coverage heuristic has not been established as a universal bound here.
- Cost Q^n for large n. For $n > 8$, factor-model reduction or sparse-grid quadrature is required.
- A signed-portfolio extension: mixed-sign weights are outside the positive-sum log-CF construction.

Verdict. The Hermite-COS construction is a deterministic numerical evaluation scheme built from an exact log-CF integral representation. It is promising on the reported positive-weight, low-dimensional tests, but rigorous evaluator bounds, signed-portfolio treatment, and broad comparative benchmarking remain open.

11.2 The Direct Log-CF COS Evaluator $\mathcal{H}$

The integral representation now yields a precise definition of the finite evaluator.

Definition. For $n, N, Q \geq 1$, positive weights, positive-definite $\Sigma = LL^T$, and a finite log-domain $[\alpha, \beta]$, define:

$$\mathcal{H}^n_{N,Q}(x; w, \mu, L, [\alpha, \beta]) := u + \sum_{k=1}^{N} \frac{2}{k\pi} \sin(k\pi u) \times \sum_{\ell=1}^{Q^n} q_\ell \cos\left(k\pi \frac{\log S(z_\ell) - \alpha}{\beta - \alpha}\right),$$

where $u = (\log x - \alpha)/(\beta - \alpha)$, (z_ℓ, q_ℓ) are tensor-product GH nodes and normalized weights, and $S(z) = \sum_i w_i \exp(\mu_i + \ell_i^T z)$.

Type: $\mathcal{H}^n_{N,Q} : \mathbb{R}_{>0} \to \mathbb{R}$ on the reconstruction domain. It is a parameterized numerical approximation to the CDF.

Parameters:

Parameter	Role	Analogous to
N	COS order (CDF reconstruction accuracy)	Terms in Fourier series
Q	Quadrature nodes per dimension (CF accuracy)	Quadrature order when computing $N(x)$
$[\alpha, \beta]$	Log-space domain of reconstruction	$(-\infty, \infty)$ for $N(x)$, in practice finite
(w, μ, L)	Generative latent evaluated at nodes	Distribution parameters

Properties (observations and open gaps):

1. *Well-posedness:* Positive weights make $S(z_\ell) > 0$, so every logarithm is defined.
2. *Range / monotonicity:* For finite (N, Q), the raw cosine sum is not guaranteed to lie in $[0, 1]$ or be monotone; the reference code clips its output to $[0, 1]$.
3. *Convergence:* The paper does not establish an end-to-end theorem as $N, Q \to \infty$ with an expanding domain, nor a uniform rate for the GH/COS stages. Rapid convergence is observed in selected tests.
4. *Differentiability:* The unclipped $\mathcal{H}$ is infinitely differentiable in x for $x > 0$ (finite sum of $\sin(\cdot \log x)$). Its derivative gives the associated finite COS density approximation, not an exact density theorem.
5. *Algebraic structure:* For fixed (N, Q) and latent parameters, $\mathcal{H}$ is a finite composition of exponential, logarithmic, trigonometric, and arithmetic operations.

Comparison with standard special functions:

Function	Definition	Structure	Elementary?	Parameters
$N(x) = \frac{1}{\sqrt{2\pi}} \int_{-\infty}^{x} e^{-t^2/2} dt$	Integral	Non-elementary integral	**No** (Liouville 1835)	None
$\mathrm{erf}(x) = \frac{2}{\sqrt{\pi}} \int_{0}^{x} e^{-t^2} dt$	Integral	Non-elementary integral	**No**	None
$\gamma(s, x) = \int_0^x t^{s-1} e^{-t} dt$	Integral	Non-elementary for general s	**No**	s
$\Phi_n(x; \mu, \Sigma)$	n-fold integral	Standard numerically evaluated function	Not classified here	μ, Σ
$\mathcal{H}^n_{N,Q}(x; w, \mu, L, [\alpha, \beta])$	**Finite numerical sum**	exp, log, sin, cos	**For fixed settings**	$N, Q, [\alpha, \beta]$

11.3 Convergence and the Latent Algebra

The pipeline is a well-defined composition of exact maps and integrals, followed by numerical evaluators. The reported experiments support those evaluators in selected regimes; a general convergence theorem for their composition is not supplied.

The core distinction is convergence, not elementariness. For any fixed precision parameters (N, Q), the Hermite-COS formula $\mathcal{H}^n_{N,Q}(x)$ is a finite sum of sines, cosines, and logarithms. Its elementary character is a property of the *finite truncation*, not a classification of the exact CDF. The cited literature provides approximations and numerical integral or transform methods rather than a generally available elementary closed form, but this paper proves no Liouville–Risch non-elementarity result for correlated lognormal sums. The relevant comparison is therefore between evaluation chains: for the same target CDF, what convergence and error evidence does each numerical chain provide?

The Latent algebra perspective. The exact integral representation (§4.4) is a composition of four algebraic operations:

1. **Map** (finite latent $\to$ grade-1 Latent): $(w, \mu, \Sigma) \mapsto S(\cdot) \in L^2(\gamma_n)$, where $S(z) = \sum w_i \exp(\mu_i + \ell_i^T z)$.
2. **Complexification** (grade-1, $\mathbb{R} \to$ grade-1, $\mathbb{C}$): $S \mapsto S^{it} = e^{it \log S}$. The log-space Mellin transform extends the real random variable to the complex plane.
3. **Gaussian contraction** (grade-1 $\to$ grade-0): $\tilde{\phi}(t) = \langle S(\cdot)^{it}, \mathbf{1} \rangle_{\gamma_n}$. The inner product against Gaussian measure produces the log-space characteristic function.
4. **Fourier contraction** (grade-0 $\to$ function). The final step uses Gil–Pelaez inversion for $T = \log S$:

$$F_S(x) = \frac{1}{2} - \frac{1}{\pi} \int_0^\infty \mathrm{Im}\left[\frac{\tilde{\phi}(t) e^{-it \log x}}{t} \right] dt, \qquad \tilde{\phi}(t) = E[S^{it}].$$

This is not the level-space characteristic function $\phi_S(t) = E[e^{itS}]$. Integration against Lebesgue measure dt recovers the CDF.

The arrows define the exact integral representation for positive S. The numerical evaluator introduces Q, N, and $[\alpha, \beta]$ as explicit precision and domain choices; the reported heuristics do not provide a proved universal selection or error guarantee.

Why the MGF route diverges. The moment-based route involves $E[e^{tS}]$ instead of $E[S^{it}]$. Along directions where one lognormal component grows, $S(z)$ itself grows exponentially in $\|z\|$, so $e^{tS(z)}$ has double-exponential growth and defeats Gaussian decay for $t > 0$. The log-space complexification satisfies $|S^{it}| = 1$ for $S > 0$, making the Gaussian expectation finite.

	Moment route	CF route
Complexification	e^{tS} (real, $t > 0$)	$S^{it} = e^{it \log S}$ (complex, Mellin)
Integrand	$e^{tS(z)}$ (double-exponential along growing latent directions)	$\lvert S(z)^{it}\rvert = 1$ (bounded)
Gaussian contraction	Divergent	Numerically approximated by GH quadrature
Failure mode	Severe Padé conditioning in referenced tests; no general asymptotic proved	Evaluator accuracy and conditioning depend on numerical settings

The Fenton problem was not about the algebraic complexity of the distribution. The finite generative latent (w, μ, Σ) always determined the distribution — the algebra was always exact. The decades-long difficulty was the choice of transform: the real MGF route uses e^{tS} with $t > 0$, whose Gaussian integral generically diverges for lognormal sums, whereas the CF / Mellin pairings used here employ a purely imaginary exponent (e^{itS} or $S^{it} = e^{it \log S}$), giving a bounded oscillatory factor on the state space and a convergent Gaussian contraction.

The complexification. The level-space characteristic function $\phi_S(t) = E[e^{itS}]$ is the standard complexification of S.

The implemented log-space COS pipeline (§4.4–4.5) instead uses the characteristic function of $T = \log S$:

$$\tilde{\phi}(t) = E[S^{it}] = E[e^{it \log S}].$$

This matches Gil–Pelaez inversion in the log coordinate, with $x = e^y$. The cosine coefficients are

$$A_k = \mathrm{Re}\left[\tilde{\phi}\left(\frac{k\pi}{\beta - \alpha}\right) e^{-ik\pi\alpha/(\beta-\alpha)}\right].$$

Under standard regularity (smooth density of $\log S$ on the truncated interval, COS assumptions as in Fang and Oosterlee, 2008), the cosine expansion is expected to converge rapidly; truncation at N terms remains a numerical approximation whose end-to-end error requires analysis.

This is the same structural move as the complexification in the finiteness proof for central configurations (Smale's 6th Problem; Nagy, 2026, *The Exact Latent Solution of the Gravitational Three-Body Problem*): extend the real CC equations $F(\mathbf{q}) = \lambda$ to $\mathbb{C}$, revealing branch points and monodromy invisible in the real domain. In both cases, the pattern is: a real-domain problem that appears intractable → complexification reveals algebraic structure → the complex structure solves or constrains the real problem. Euler discovered this principle for trigonometric identities; the Latent framework applies it to distribution theory and algebraic geometry as two instances of the same operation.

---

12. Formal Verification

12.1 Statement-Level Proof Map

No proof package is claimed for this paper. Theorem 1, Proposition 1, Theorem 2, and Proposition 2 retain ordinary paper-proof provenance: each has explicit hypotheses and a manuscript derivation, but no formal companion statement. Their absence from a machine-checkable statement layer is not a mathematical defect and is not concealed by synthetic annotations.

Executing fenton_hermite_p3.py reports 20/20 passing regression tests. The current tracked manifest fenton_hermite_p3_manifest.yaml, however, records five checked kernel theorems, all five unused, with no dependency edges or terminal proof outputs. These are distinct facts: the executable regression count is 20, while the manifest's checked theorem inventory is 5. Neither provides a formal proof of a manuscript theorem.

Manuscript claim	Formal status	Evidence used here
Theorem 1: explicit multivariate Hermite coefficients and coefficient bound	Not kernel-formalized	Conventional derivation in §2
Theorem 2: L^2 convergence and component-tail bound	Not kernel-formalized	Parseval and multinomial argument in §4
Exact log-CF/Gil–Pelaez integral representation	Not kernel-formalized	Standard probability theorems and the proof in §4
Finite GH/COS evaluator and parameter heuristics	Not kernel-formalized	Executable numerical validation in §7
Unit-modulus arithmetic after assuming a unit modulus	Executable regression test only; not among the five current manifest declarations	fenton_hermite_p3.py; does not prove the complex modulus identity itself
Generic growth and monotonicity implications	Five checked/unused manifest theorems; 20 executable tests	fenton_hermite_p3.py; not an end-to-end convergence theorem

The failing fenton_copula_newton.py entry point concerns copulas, CDO tranches, and portfolio-loss DSL statements. The failing spectral_fenton_proof.py entry point concerns variance inflation and spectral reduction. Neither supplies a premise or conclusion used in Theorems 1–2, the log-CF representation, or the numerical evaluator. They are therefore not load-bearing for this manuscript and are excluded from its proof provenance. This exclusion does not upgrade the main claims to formal verification: the table above is the complete statement-level boundary.

Declaration of Generative AI and AI-assisted technologies in the writing process

During the preparation of this work the author used Cursor and its integrated AI models (Claude and GPT families) in order to assist with manuscript drafting, structural editing, and language polishing. After using this tool/service, the author reviewed and edited the content as needed and takes full responsibility for the content of the publication.

References

- Nagy, T. (2026). The Latent: Finite Sufficient Representations of Smooth Systems. *Zenodo*. DOI: 10.5281/zenodo.19101209
- Nagy, T. (2026). The Exact Latent Distribution of Correlated Lognormal Sums. *Working paper.*
- Nagy, T. (2026). The Smooth Latent Operator: Parameter-Free Distributional Representations via Kernel Moment Recovery. *Working paper.*
- Nagy, T. (2026). The Exact Latent Solution of the Gravitational Three-Body Problem. *Zenodo*. DOI: 10.5281/zenodo.19101229
- Cameron, R.H. and Martin, W.T. (1947). The orthogonal development of non-linear functionals in series of Fourier-Hermite functionals. *Annals of Mathematics*, 48(2), 385–392. DOI: 10.2307/1969178
- Wiener, N. (1938). The homogeneous chaos. *American Journal of Mathematics*, 60(4), 897–936. DOI: 10.2307/2371268
- Xiu, D. and Karniadakis, G.E. (2002). The Wiener–Askey polynomial chaos for stochastic differential equations. *SIAM Journal on Scientific Computing*, 24(2), 619–644. DOI: 10.1137/S1064827501387826
- Baker, G.A. and Graves-Morris, P. (1996). *Padé Approximants*. Cambridge University Press.
- Gonnet, P., Güttel, S., and Trefethen, L.N. (2013). Robust Padé approximation via SVD. *SIAM Review*, 55(1), 101–117. DOI: 10.1137/110853236
- Ernst, O.G., Mugler, A., Starkloff, H.J., and Ullmann, E. (2012). On the convergence of generalized polynomial chaos expansions. *ESAIM: Mathematical Modelling and Numerical Analysis*, 46(2), 317–339. DOI: 10.1051/m2an/2011045
- Nobile, F., Tempone, R., and Webster, C.G. (2008). A sparse grid stochastic collocation method for partial differential equations with random input data. *SIAM Journal on Numerical Analysis*, 46(5), 2309–2345. DOI: 10.1137/060663660
- Fenton, L. F. (1960). The Sum of Log-Normal Probability Distributions in Scatter Transmission Systems. *IRE Transactions on Communications Systems*, CS-8, 57-67. DOI: 10.1109/tcom.1960.1097606

- Gubner, J. A. (2006). A New Formula for Lognormal Characteristic Functions. *IEEE Transactions on Vehicular Technology*, 55(5), 1668–1671. DOI: 10.1109/tvt.2006.878610
- Mehta, N. B., Wu, J., Molisch, A. F., and Zhang, J. (2007). Approximating a Sum of Random Variables with a Lognormal. *IEEE Transactions on Wireless Communications*, 6(7), 2690–2699. DOI: 10.1109/twc.2007.051000
- Senaratne, D. and Tellambura, C. (2009). Numerical Computation of the Lognormal Sum Distribution. *IEEE GLOBECOM 2009*, 1–6. DOI: 10.1109/glocom.2009.5425708
- Tellambura, C. and Senaratne, D. (2010). Accurate Computation of the MGF of the Lognormal Distribution and Its Application to Sum of Lognormals. *IEEE Transactions on Communications*, 58(5), 1568–1577. DOI: 10.1109/tcomm.2010.05.080640
- Mahmoud, A. S. H. (2010). New Quadrature-Based Approximations for the Characteristic Function and the Distribution Function of Sums of Lognormal Random Variables. *IEEE Transactions on Vehicular Technology*, 59(7), 3364–3372. DOI: 10.1109/tvt.2010.2051467
- Lo, C.-F. (2013). WKB Approximation for the Sum of Two Correlated Lognormal Random Variables. *Applied Mathematical Sciences*, 7, 6277–6289. DOI: 10.12988/ams.2013.39511
- Berggren, F. (2024). Bounds and Approximations for the Distribution of a Sum of Lognormal Random Variables. *arXiv:2405.09283*. DOI: 10.48550/arXiv.2405.09283
- Fang, Fang and Oosterlee, Cornelis W. (2008). A Novel Pricing Method for European Options Based on Fourier-Cosine Series Expansions. *SIAM Journal on Scientific Computing*, 31(2), 826-848. DOI: 10.1137/080718061
- Imhof, J.P. (1961). Computing the distribution of quadratic forms in normal variables. *Biometrika*, 48(3/4), 419–426. DOI: 10.2307/2332763
- Schwartz, S.C. and Yeh, Y.S. (1982). On the Distribution Function and Moments of Power Sums with Lognormal Components. *Bell System Technical Journal*, 61, 1441–1462. DOI: 10.1002/j.1538-7305.1982.tb04353.x
- Liouville, J. (1835). Mémoire sur l'intégration d'une classe de fonctions transcendantes. *Journal für die reine und angewandte Mathematik*, 13, 93–118. DOI: 10.1515/crll.1835.13.93
- Risch, R.H. (1969). The problem of integration in finite terms. *Transactions of the American Mathematical Society*, 139, 167–189. DOI: 10.1090/S0002-9947-1969-0237477-8
- Nelson, E. (1973). The free Markoff field. *Journal of Functional Analysis*, 12(2), 211–227. DOI: 10.1016/0022-1236(73)90025-6
- Favard, J. (1935). Sur les polynômes de Tchebicheff. *Comptes Rendus de l'Académie des Sciences*, 200, 2052–2053.
- Abramowitz, M. and Stegun, I.A. (1964). *Handbook of Mathematical Functions*. National Bureau of Standards.
- Trefethen, L.N. (2013). *Approximation Theory and Approximation Practice*. SIAM. DOI: 10.1137/1.9781611975949
- Gil-Pelaez, J. (1951). Note on the inversion theorem. *Biometrika*, 38(3/4), 481–482. DOI: 10.1093/biomet/38.3-4.481