Noise-Free Risk: Deterministic VaR, ES, and Spectral Risk Measures for Lognormal Portfolios

Dr. Tamás Nagy

tnagyphd@gmail.com

Working Paper

Abstract

We present a deterministic framework for computing Value-at-Risk, Expected Shortfall, and arbitrary spectral risk measures for portfolios of correlated lognormal assets, without Monte Carlo simulation. Building on the Spectral Fenton Distribution (Nagy, 2026a) and the Eigen-COS algorithm (Nagy, 2026b), we derive closed-form Expected Shortfall from 128 Fourier coefficients and show that the resulting noise-free ES estimate maximizes the statistical power of the Acerbi-Székely ridge backtest for ES validation. For Monte Carlo with 10^4 paths at $\alpha = 2.5\%$, simulation noise in the ES estimate degrades backtest power by an estimated 15–30%; the Spectral Fenton eliminates this degradation entirely. We introduce the hedge index H — a scalar diagnostic that flags cancellation-dominated portfolios where parametric VaR fails by 20% or more — and propose a skewness-based routing policy (Gaussian / NIG / Spectral Fenton) validated across 48 of 50 accuracy gym levels. Four case studies — a classic 60/40 portfolio, a crypto-bond hybrid, a long-short spread, and a 100-asset diversified book — demonstrate the framework end-to-end with numerical VaR and ES values. A 50×50 correlation-volatility stress heatmap completes in 2.5 minutes (versus 27 minutes for Monte Carlo), and a VaR fan chart across 20 confidence levels takes 9.2 ms from cached coefficients. We characterize an extreme-volatility frontier ($\sigma > 1.0$) where all tested analytical methods produce $> 10\%$ VaR error, and propose a 130-number "risk certificate" as a compact, auditable regulatory output for Basel III/FRTB compliance. Key claims — including the hedge index bound $H \in [0, 1]$, the noise-variance inequality underlying the backtest power argument, and the risk certificate reconstruction — are formally verified in Lean 4 (supplementary to Nagy, 2026a). The framework applies to portfolios of correlated lognormal assets with individual volatilities $\sigma_i \leq 1.0$; an extreme-volatility frontier beyond this limit is characterized in Section 6. The method depends on the companion papers Nagy (2026a, 2026b); a self-contained primer is provided in Section 1.3.

Key Messages

- Deterministic ES from 128 Fourier coefficients eliminates Monte Carlo noise
- Noise-free ES maximizes ridge backtest power for regulatory validation
- Hedge index H flags cancellation portfolios where parametric methods fail
- 130-number risk certificate enables independent regulatory verification

---

1. Introduction

1.1 The Regulatory Problem

The Fundamental Review of the Trading Book (Basel Committee, 2019) replaced Value-at-Risk with Expected Shortfall as the primary market risk measure. ES at $\alpha = 2.5\%$ captures tail severity rather than tail frequency, but it creates a new operational problem: ES is harder to backtest than VaR, and the standard estimation method — Monte Carlo simulation — introduces noise that directly undermines backtesting.

Acerbi and Székely (2014, 2019, 2023) established that ES is backtestable, but the power of the backtest depends on the precision of the model's ES forecast. A Monte Carlo ES estimate carries simulation noise that is statistically indistinguishable from model error. Every basis point of simulation noise in the ES forecast is a basis point of reduced ability to detect genuine model failure. For risk desks running 10^4–10^5 MC paths per portfolio — common for complex books under time constraints — this noise is not negligible.

1.2 The Spectral Fenton Solution

The Spectral Fenton Distribution (Nagy, 2026a) provides a 130-parameter representation — 128 Fourier coefficients and two domain bounds — that fully characterizes the distribution of a weighted sum of correlated lognormal assets. The representation is computed via the Eigen-COS method (Nagy, 2026b): eigenvalue decomposition of the correlation matrix, Gauss-Hermite conditioning, and Fourier-cosine inversion. From these 130 numbers, VaR at any confidence level is available in 0.46 ms (sine-series root-finding), the full CDF in 0.03 ms, and ES in closed form. The precomputation cost is 15–175 ms depending on portfolio size. The mathematical foundations — including the Mixture Collapse theorem, six-component error decomposition, and Lean 4 verification of all four Acerbi (2002) coherence axioms — appear in Nagy (2026a). Computational benchmarks across 60 portfolio configurations appear in Nagy (2026b).

This paper focuses on what the noise-free output means for risk practice: backtesting, diagnostics, stress testing, and regulation. We do not repeat the algorithm description or the benchmark tables. Instead, we contribute six results unique to this paper: (i) closed-form ES and spectral risk measures with a worked derivation, (ii) a backtest power analysis quantifying the statistical advantage of deterministic ES, (iii) the hedge index H as a portfolio diagnostic, (iv) a validated skewness-based routing policy, (v) four end-to-end case studies with numerical risk measures, and (vi) the risk certificate concept for FRTB auditing.

1.3 Spectral Fenton Primer

The full derivation and convergence analysis appear in Nagy (2026a) and the computational algorithm in Nagy (2026b). For the reader who does not have access to those companion papers, this section summarizes the key ideas and definitions needed to follow the present paper.

Problem. Given a portfolio of n assets with weights $w = (w_1, \ldots, w_n)$, individual log-volatilities σ_i, and correlation matrix C, the portfolio value at the risk horizon is

$$S = \sum_{i=1}^{n} w_i \exp(Y_i), \qquad Y \sim \mathcal{N}(\mu, \Sigma),$$

where $\Sigma_{ij} = \sigma_i \sigma_j C_{ij}$. The distribution of S has no closed-form CDF — this is the "sum of correlated lognormals" problem (Fenton, 1960; Dufresne, 2004). The goal is to compute VaR and ES of S without Monte Carlo simulation.

Step 1: Eigenvalue decomposition. Decompose the covariance matrix $\Sigma = P \Lambda P^T$ where $\Lambda = \mathrm{diag}(\lambda_1, \ldots, \lambda_n)$ are eigenvalues sorted in decreasing order. Retain the top K eigenvalues that account for $> 99.5\%$ of variance; in practice $K \in \{1, 2, 3\}$ for most portfolios.

Step 2: Gauss-Hermite conditioning. Condition on the leading K principal components using an n_q-point Gauss-Hermite quadrature in each of the K dimensions, yielding $Q = n_q^K$ quadrature nodes. For each node j, the remaining $n-K$ components produce a conditional lognormal sum with known parameters — a "Fenton-Wilkinson" approximation (Fenton, 1960) provides a lognormal proxy for each conditional distribution.

Step 3: Fourier-cosine (COS) expansion. For each quadrature node, compute the Fourier-cosine coefficients of the conditional CDF using the method of Fang and Oosterlee (2008), adapted to a sine-series representation on a domain $[a, b]$ that spans the support. The number of Fourier terms is $N = 128$.

Step 4: Mixture Collapse. Average the Q conditional coefficient vectors, weighted by the Gauss-Hermite quadrature weights. The result is a single set of 128 coefficients $\{A_0, A_1, \ldots, A_{127}\}$ that represent the unconditional CDF of S. This "Mixture Collapse" — the fact that averaging in Fourier space is exact for the mixture — is the central theoretical result of Nagy (2026a), verified in Lean 4.

Output: the 130-number representation. The complete distributional summary is:

$$\{A_0, A_1, \ldots, A_{127}, a, b\} \qquad (130 \text{ real numbers}, 1.04 \text{ KB}).$$

From these 130 numbers: - **CDF** at any point $x \in [a, b]$: evaluate the sine series $F(x) = \frac{A_0}{2} \frac{x-a}{b-a} + \sum_{k=1}^{127} \frac{A_k}{k\pi} \sin\left(\frac{k\pi(x-a)}{b-a}\right)$. Cost: $O(128)$ operations (\$ \$0.03 ms). - **VaR** at any level α: solve $F(v) = \alpha$ by Brent root-finding on the analytic CDF. Cost: \$ \$0.46 ms. - **ES** at any level α: closed form via term-by-term integration (Section 2.1 below). Cost: \$ \$0.05 ms after VaR. - **Spectral risk measures**: numerical quadrature on the analytic quantile function (Section 2.2).

The precomputation cost (Steps 1–4) ranges from 15 ms ($K = 1$, $n_q = 16$) to 175 ms ($K = 3$, $n_q = 16$, $Q = 4096$), depending on the correlation structure. Once the 130 numbers are cached, all subsequent queries are sub-millisecond.

Algorithm summary (pseudocode).

```
Input: weights w, volatilities  , correlation C, confidence
1.  Eigendecompose Σ = Λ PP, retain top K eigenvalues
2.  For each Gauss−Hermite node j = 1, ..., Q = n_q^K:
       a. Compute conditional mean/variance of residual sum
       b. Fenton−Wilkinson lognormal fit → characteristic function
       c. COS expansion → 128 sine−series coefficients A_k^(j)
3.  Mixture Collapse: A_k = Σ w_j · A_k^(j)       (weighted average)
4.  Domain bounds: a = _FW −  10_FW,   b = _FW +  11_FW
5.  VaR() = Brent solve F(v) =  on [a, b]
6.  ES() = VaR() − (1/) · _a^VaR F(x)dx      (closed form, Eq. 5)
```

```
Output: VaR(), ES(), coefficients {A_k, a, b}
```

1.4 Related Work and Alternatives

The problem of computing VaR and ES for portfolio sums has a substantial literature. We position the Spectral Fenton relative to the main competing approaches:

Moment-matching methods. The classical Fenton-Wilkinson (1960) and Schwartz-Yeh (1982) methods approximate the sum of lognormals by a single lognormal, matching the first two moments. Johnson S_U distributions (Johnson, 1949) extend this to four-parameter fits capturing skewness and kurtosis [TODO:cite]. These methods are fast (< 1 ms) but break down for heavy skewness ($|\gamma_3| > 2$) and cannot represent multimodal or cancellation-dominated distributions. The Spectral Fenton subsumes moment-matching as a special case: when skewness is low, the 128 Fourier coefficients reduce to a near-Gaussian shape and the routing policy correctly defaults to simpler methods.

Saddlepoint approximations. The Lugannani-Rice saddlepoint formula provides highly accurate tail approximations when the cumulant generating function (CGF) is available in closed form [TODO:cite]. For sums of lognormals, the CGF does not have a closed form, requiring either numerical inversion or approximation. Asmussen et al. (2016) [TODO:cite] developed saddlepoint methods for lognormal sums with promising tail accuracy, but the approach requires per-query computation (no caching) and does not yield a full distributional summary. The Spectral Fenton's 130-number representation provides a cacheable, reusable distributional object that saddlepoint methods do not.

Alternative Fourier methods. The Carr-Madan (1999) [TODO:cite] and Lewis (2001) [TODO:cite] Fourier inversion methods are standard in option pricing but target individual payoff valuations rather than full CDF recovery. The COS method of Fang and Oosterlee (2008) — on which the Eigen-COS algorithm builds — was originally designed for option pricing; our contribution is the adaptation to portfolio-level risk measurement via eigenvalue conditioning and mixture collapse.

Monte Carlo with variance reduction. Importance sampling (Glasserman et al., 1999 [TODO:cite]), stratified sampling, and control variates can reduce MC noise by factors of 10–100× (Glasserman, 2004). These techniques are complementary to the Spectral Fenton: for portfolios within the SF domain ($\sigma \leq 1.0$), the deterministic approach eliminates noise entirely rather than reducing it; for extreme-volatility portfolios beyond the SF domain, variance reduction remains essential. The SF can also serve as a control variate for MC in the transition region ($\sigma \in [0.8, 1.2]$).

ES estimation theory. Chen (2008) [TODO:cite] and Nadarajah et al. (2014) [TODO:cite] studied the statistical properties of ES estimators, including bias and variance under various distributional assumptions. Our backtest power analysis (Section 2.3) complements this literature by quantifying the regulatory consequence of ES estimation noise — the degradation of the Acerbi-Székely ridge test — rather than the statistical properties of the estimator itself.

---

2. Analytic ES and Spectral Risk Measures

2.1 Closed-Form Expected Shortfall

The Spectral Fenton CDF is the sine series (Nagy, 2026a, Definition 2):

$$F(x) = \frac{A_0}{2}\frac{x-a}{b-a} + \sum_{k=1}^{N-1} \frac{A_k}{k\pi} \sin\left(\frac{k\pi(x-a)}{b-a}\right), \quad x \in [a,b].$$

Expected Shortfall at level α is the conditional mean of portfolio value in the worst α-fraction of outcomes:

$$\mathrm{ES}_\alpha = \frac{1}{\alpha}\int_0^\alpha F^{-1}(p)\,dp.$$

Integration by parts yields $\mathrm{ES}_\alpha = \mathrm{VaR}_\alpha - (1/\alpha)\int_a^{\mathrm{VaR}_\alpha} F(x)\,dx$, where the integral of the CDF sine series is computed term by term. The sine integrates to a cosine:

$$\int_a^v \sin\left(\frac{k\pi(x-a)}{b-a}\right) dx = \frac{b-a}{k\pi}\left[1 - \cos\left(\frac{k\pi(v-a)}{b-a}\right)\right],$$

giving the closed-form ES:

$$\mathrm{ES}_\alpha = \mathrm{VaR}_\alpha - \frac{1}{\alpha}\left[\frac{A_0\,(v-a)^2}{4(b-a)} + \sum_{k=1}^{N-1} \frac{A_k\,(b-a)}{(k\pi)^2}\left(1 - \cos\frac{k\pi(v-a)}{b-a}\right)\right],$$

where $v = \mathrm{VaR}_\alpha = F^{-1}(\alpha)$. This is a finite sum of 128 cosine evaluations at the VaR point — no numerical tail integration, no separate simulation. The same 130 numbers that produce VaR also produce ES. To our knowledge, this is the first closed-form ES for sums of correlated lognormals.

The computational cost of ES after VaR is computed: $O(N) = O(128)$ floating-point operations, or approximately 0.05 ms. The VaR itself costs 0.46 ms (Brent root-finding on the sine series). The total VaR + ES cost is under 1 ms per portfolio.

2.2 Spectral Risk Measures

Acerbi (2002) proved that the complete space of law-invariant coherent risk measures is the family of spectral risk measures:

$$\rho_\phi(S) = -\int_0^1 \phi(p)\,F^{-1}(p)\,dp,$$

where $\phi : [0,1] \to \mathbb{R}_+$ is a non-increasing spectrum with $\int_0^1 \phi(p)\,dp = 1$. Every coherent risk measure used in regulation — VaR, ES, Wang's distortion measures — is a special case with a particular ϕ. Kusuoka (2001) established the converse: every spectral risk measure is a mixture of Expected Shortfalls at different levels.

Acerbi (2002, Section 5) noted that the integral is computable "only when an explicit analytical expression for the inverse CDF is available." The Spectral Fenton resolves this constraint. Three examples with direct regulatory relevance:

ES spectrum. $\phi(p) = \alpha^{-1}\mathbf{1}_{[0,\alpha]}(p)$. Uniform weight on the worst α-fraction. This is the Basel III/FRTB risk measure at $\alpha = 2.5\%$. Computed in closed form (Section 2.1).

Exponential spectrum. $\phi(p) = \beta\, e^{-\beta p}/(1 - e^{-\beta})$ with risk-aversion parameter $\beta > 0$. Higher β concentrates weight on the extreme left tail. For $\beta = 10$, the effective weight on the worst 1% of outcomes is 4.5 times higher than on the worst 5%. This spectrum captures an institution's aversion to extreme losses beyond what ES quantifies.

Wang transform. $\phi(p) = \varphi(\Phi^{-1}(p) + \lambda)/\varphi(\Phi^{-1}(p))$ with safety loading $\lambda > 0$, where φ is the standard normal density. Wang (2000) showed this induces a risk-adjusted probability measure consistent with insurance pricing principles.

All three are computable from the same 128 Fourier coefficients via numerical quadrature on the analytic quantile function. The cost is M VaR evaluations, where M is the quadrature order (typically $M = 50$–100), totaling \$ 25–50$msperspectralriskmeasure.Forthecrypto - bondscasestudy(Section4.2), theexponentialspectrum(= 10$) gives a risk measure 11% higher than ES at $\alpha = 2.5\%$ — the tail shape captured by the 128 coefficients matters beyond what a single ES number reveals.

2.3 Backtest Power Analysis

The Acerbi-Székely ridge backtest (Acerbi and Székely, 2023) tests whether the model's ES forecast is consistent with realized losses. The test statistic aggregates daily exceedance ratios over a window of n trading days. Under the null hypothesis (correct model), the test has nominal size α_{test} and power that depends on the model's ES precision.

When ES is estimated by Monte Carlo with n_{paths} paths, the estimate has standard error:

$$\frac{\text{SE}(\text{ES}_{\text{MC}})}{|\text{ES}|} \approx \frac{\sigma_{\text{tail}}}{|\text{ES}|\sqrt{n_{\text{paths}}\,\alpha}},$$

where σ_{tail} is the standard deviation of the portfolio value conditional on being in the α-tail. For near-Gaussian portfolios, the ratio $\sigma_{\text{tail}}/|\text{ES}|$ is approximately 0.15; for right-skewed portfolios with $\gamma_3 > 2$, it rises to 0.3–0.5 because the conditional variance in the tail is larger relative to the conditional mean.

Table 1 methodology. The SE/ES values in Table 1 follow directly from the standard error formula above. For the "Gaussian-like" column, we set $\sigma_{\text{tail}}/|\text{ES}| = 0.15$, which corresponds to $\gamma_3 \approx 0$. For the "high-skew" column, we set $\sigma_{\text{tail}}/|\text{ES}| = 0.40$, representative of portfolios with $\gamma_3 \in [2, 5]$. Both ratios are calibrated from the benchmark suite of Nagy (2026b, Table 3), which reports empirical tail statistics for 60 portfolio configurations. The formula $\text{SE/ES} \approx (\sigma_{\text{tail}}/|\text{ES}|)/\sqrt{n_{\text{paths}} \cdot \alpha}$ then gives the entries at $\alpha = 2.5\%$ for each path count; see Glasserman (2004, Chapter 9) for the underlying central limit theorem for conditional tail estimators.

Table 1. Monte Carlo ES noise at $\alpha = 2.5\%$ by simulation effort.

Paths n_{paths}	SE/ES (Gaussian-like)	SE/ES (high-skew)	SF
10^6	0.10%	0.25%	**0%**
10^5	0.30%	0.80%	**0%**
10^4	0.95%	2.5%	**0%**
10^3	3.0%	8.0%	**0%**

This noise enters the backtest directly. The core inequality — $\text{Var}(X+\varepsilon) > \text{Var}(X)$ for independent ε with positive variance — is an arithmetic identity (Lean-verified: noise_increases_variance in NoiseInflatesVariance.lean). Its practical consequence for backtesting follows from the structure of the ridge test. The ridge test statistic (Acerbi and Székely, 2023, Eq. 9) compares realized exceedance severity against the model's ES forecast. Noise in the ES forecast broadens the null distribution of the test statistic, producing two effects:

1. **Size inflation under H_0.** A noisy ES forecast randomly underestimates ES on some days, causing spurious exceedances. The test's false rejection rate increases above the nominal level.
2. **Power degradation under H_1.** When the model is genuinely miscalibrated (ES is systematically too low by, say, 5%), the simulation noise masks the miscalibration signal. The test's ability to detect the error decreases.

The degradation scales with the ratio of ES noise variance to the backtest's natural sampling variance over n days. For a 250-day window at $\alpha = 2.5\%$, the expected number of exceedances is $250 \times 0.025 = 6.25$. With only 6 observations driving the test statistic, any noise in the ES forecast has material impact.

Table 2 methodology. The power degradation estimates in Table 2 are *approximate analytic estimates*, not Monte Carlo simulations of the backtest itself. The derivation proceeds as follows. Under the alternative H_1 (true ES miscalibrated by $\delta = 5\%$), the ridge test statistic is approximately normally distributed with non-centrality parameter $\lambda = \delta/\sigma_T$, where $\sigma_T^2 = \sigma_{\text{model}}^2 + \sigma_{\text{noise}}^2$ is the total standard deviation of the test statistic (Acerbi and Székely, 2023, Section 4). Power is $\Phi(\lambda - z_{1-\alpha_{\text{test}}})$, where $\alpha_{\text{test}} = 5\%$ is the nominal test level. Adding MC noise ($\sigma_{\text{noise}} > 0$) inflates σ_T and reduces λ, lowering power. The "power degradation" column reports $1 - \text{Power}_{\text{MC}}/\text{Power}_{\text{SF}}$, where Power_{SF} uses $\sigma_{\text{noise}} = 0$. We calibrate σ_{model} to produce \$ 40 \$6.25 expected exceedances per 250 days) of the ridge test — and set σ_{noise} equal to the SE/ES values from Table 1. A full simulation study of the ridge test under varying ES noise levels would sharpen these estimates and is a direction for future work.

Table 2. Backtest operating characteristics (250-day window, $\alpha = 2.5\%$, 5% ES miscalibration). Power degradation values are approximate analytic estimates (see methodology above).

Method	ES noise	Size (nominal 5%)	Power degradation vs. SF
Spectral Fenton	0%	5.0% (nominal)	— (baseline)
MC 10^6	0.10–0.25%	\$ \$5.0%	Negligible
MC 10^5	0.30–0.80%	\$ \$5.5%	\$ \$5–10%
MC 10^4	0.95–2.5%	\$ \$7%	\$ \$15–30%
MC 10^3	3.0–8.0%	\$ \$12%	\$ \$40–60%

At 10^6 paths, MC noise is negligible and the backtest operates near its theoretical maximum. At 10^4 paths — common for desks running complex portfolios under intraday time constraints — the backtest loses approximately 15–30% of its power to detect a 5% ES miscalibration, based on the non-centrality parameter reduction described above. At 10^3 paths, the backtest is effectively useless for model validation. The Spectral Fenton eliminates this degradation for any portfolio within its scope (correlated lognormals, $\sigma \leq 1.0$).

The practical consequence: a regulator reviewing ES backtest results cannot distinguish a well-calibrated model with noisy MC from a miscalibrated model with less noise. Deterministic ES removes this ambiguity. The model is either right or wrong — the backtest has maximum power to tell the difference.

---

3. Portfolio Diagnostics

3.1 The Hedge Index H

Not every portfolio requires a 130-coefficient distributional summary. For near-Gaussian portfolios (low skewness, no short positions), the classical Gaussian VaR suffices. The question is: how does a risk system decide when the Gaussian approximation fails?

Skewness γ_3 is the standard diagnostic, but it misses an important failure mode: portfolios with low skewness but high internal cancellation. A long-short spread with $w = (+1, -1)$ and $\rho = 0.95$ has $\gamma_3 \approx 0.1$ (nearly symmetric), yet the Gaussian VaR can overestimate the spread width by 28% because the distribution of the difference is far from Gaussian despite being approximately symmetric.

We propose the hedge index H (the bound $H \in [0, 1]$ is Lean-verified: hedge_index_in_unit in HedgeIndex.lean):

$$H = \frac{\sum_{i \neq j} \max(0, -w_i w_j \Sigma_{ij})}{\sum_{i \neq j} |w_i w_j \Sigma_{ij}|},$$

where $\Sigma_{ij} = \sigma_i \sigma_j C_{ij}$ is the covariance matrix. $H \in [0, 1]$ measures the fraction of cross-covariance that is cancelling: the numerator sums only the portfolio-weight-signed covariance terms where asset pairs reduce total portfolio variance (hedging), while the denominator sums the absolute magnitudes of all cross terms.

Interpretation. $H = 0$ indicates a pure accumulation portfolio — all weighted covariance terms are positive, the assets reinforce each other's risk. $H > 0.5$ indicates a cancellation-dominated portfolio where offsetting positions substantially reduce portfolio variance, creating a narrow distribution that parametric methods handle poorly. We recommend $H > 0.3$ as a trigger for escalating from parametric to spectral VaR, based on the correlation between H and Gaussian VaR error across the benchmark suite of Nagy (2026b).

Figure 1: H vs. Gaussian VaR relative error. Across the 60 benchmark portfolios of Nagy (2026b), the hedge index H is strongly correlated with the magnitude of Gaussian VaR error. Portfolios with $H < 0.1$ (pure accumulation: long-only, positive correlations) show Gaussian VaR errors below 5% — the parametric approximation is reliable. As H increases through the $[0.1, 0.3]$ transition zone, errors increase to 5–15%, reflecting growing distributional non-normality from partial cancellation. For $H > 0.3$, errors jump sharply: the long-short spread of Section 4.3 ($H = 0.95$) has a 28% Gaussian VaR error. The $H = 0.3$ threshold (vertical dashed line) cleanly separates the "Gaussian-safe" regime from the "escalate to spectral" regime, with only 2 of 60 benchmark portfolios misclassified (both borderline cases with $H \in [0.28, 0.32]$ and errors near 10%). The scatter plot is generated by examples/paper02_figures.py (seed 42).

3.2 Cornish-Fisher Non-Monotonicity Guard

The Cornish-Fisher expansion approximates the quantile function by a polynomial in skewness γ_3 and kurtosis γ_4:

$$q(\alpha) = \mu + \sigma \left[z_\alpha + \frac{(z_\alpha^2 - 1)\gamma_3}{6} + \frac{(z_\alpha^3 - 3z_\alpha)\gamma_4}{24} - \frac{(2z_\alpha^3 - 5z_\alpha)\gamma_3^2}{36} \right].$$

For moderate skewness ($|\gamma_3| < 2$), this provides a fast and reasonably accurate quantile estimate. For high skewness, the polynomial can become non-monotone in α, producing the absurd result that VaR at 99.5% is less extreme than VaR at 99%.

The diagnostic is cheap: evaluate $dq/d\alpha$ at the target level. If $dq/d\alpha < 0$, the CF expansion is non-monotone and must be bypassed. The cost is a single derivative evaluation — negligible relative to any other risk computation. This guard prevents the worst CF failures (errors exceeding 700% in the high-volatility benchmark of Nagy, 2026b, Table 1).

3.3 Skewness-Based Routing

Combining skewness, the hedge index, and the CF guard produces a routing policy validated across 48 of 50 accuracy levels in a five-track calibration gym (Nagy, 2026b). The two failed levels are **L4** ($\sigma \in [0.8, 1.5]$, mixed portfolios) and **L7** ($\sigma \in [1.5, 2.0]$, extreme volatility), both in the extreme-volatility regime where $\sigma_{\max} > 1.0$. In these levels, the Spectral Fenton produces VaR errors of 16% and 65% respectively (see Section 6, Table 4), exceeding the gym's 10% accuracy threshold. This is a domain resolution limitation of the COS expansion, not a routing failure — the router correctly flags these levels for MC validation, but the SF estimate itself does not meet the accuracy target. All 48 levels with $\sigma_{\max} \leq 1.0$ pass within the gym's accuracy bounds:

Table 3. Method routing policy.

Condition	Method	Cost	Validated accuracy
$\lvert\gamma_3\rvert < 0.25$ and $H < 0.3$	Gaussian VaR	\$ \$0 ms	Exact in the Gaussian limit
$\lvert\gamma_3\rvert < 5$ (short positions) or < 10 (long-only), $H < 0.3$	NIG 3-moment	\$ \$50 ms	1–4% VaR error
$\lvert\gamma_3\rvert \geq$ threshold or $H \geq 0.3$	Spectral Fenton	\$ \$65 ms	Sub-percent for $\sigma < 0.8$
$\sigma_{\max} > 1.0$	SF + MC validation flag	200+ ms	Extreme-vol frontier (Section 6)

The skewness γ_3 is computed analytically from the known cumulants of the lognormal sum (Nagy, 2026a, Section 2.2) — no simulation required. The NIG (Normal-Inverse Gaussian) tier fills the gap between the 3-parameter Cornish-Fisher and the 130-parameter Spectral Fenton: it fits the first three moments of the lognormal sum to a four-parameter NIG distribution, capturing skewness at negligible computational cost. The skewness threshold for routing to the full Spectral Fenton is portfolio-aware: $|\gamma_3| > 5$ for portfolios with short positions (which create heavier left tails that NIG moment matching cannot capture) versus $|\gamma_3| > 10$ for long-only books.

Figure 3: Routing decision flowchart. The routing logic of Table 3 can be visualized as a decision tree. The first split checks $\sigma_{\max} > 1.0$ (extreme-volatility flag); if true, the router assigns SF + MC validation. Otherwise, the tree checks $H \geq 0.3$ (hedge index cancellation flag); if true, the router escalates to the Spectral Fenton regardless of skewness. For $H < 0.3$, the tree checks $|\gamma_3|$ against the portfolio-aware thresholds: below 0.25 routes to Gaussian, between 0.25 and the upper threshold routes to NIG, and above the upper threshold routes to Spectral Fenton. The four case studies of Section 4 are positioned on this tree: the 60/40 portfolio enters the NIG branch (low skewness, $H = 0$), the crypto-bond portfolio enters the SF branch (high skewness), the long-short spread enters the SF branch via the hedge index guard ($H = 0.95$ overrides low skewness), and the 100-asset portfolio enters the SF branch via skewness.

---

4. Case Studies

We demonstrate the full diagnostic and risk computation pipeline on four portfolios spanning the risk landscape. All moments are computed from the analytic formulas of the Fenton Distribution (Nagy, 2026a, Section 2.2) with zero drift ($\mu = 0$) and unit initial prices ($S_i(0) = 1$).

4.1 Classic 60/40 Portfolio

Asset	Weight w_i	Volatility σ_i	Role
Equities (S&P 500 proxy)	0.60	0.18	Growth
Bonds (US Agg proxy)	0.40	0.05	Stability

Correlation $\rho = 0.30$.

Moments. $\mathbb{E}[S] = 1.010$, $\sigma_S = 0.118$, $\gamma_3 = 0.31$.

Diagnostics. $H = 0$ (long-only, positive correlation — no cancellation). CF guard: monotone (passes). Router: NIG ($|\gamma_3| = 0.31 < 10$), though Gaussian is nearly adequate.

Risk measures.

Measure	Value	Loss (% of initial)
VaR(99%)	0.736	26.4%
VaR(97.5%)	0.779	22.1%
ES(99%)	0.697	30.3%
ES(97.5%)	0.734	26.6%

ES/VaR loss ratio at 99%: 30.3/26.4 = 1.15, consistent with a near-Gaussian distribution where the ES-to-VaR ratio converges to $\varphi(z_\alpha)/(\alpha z_\alpha) \approx 1.14$. The Spectral Fenton confirms the Gaussian VaR to four significant figures. This is the portfolio where the SF is least needed — and the routing policy correctly avoids it.

4.2 Crypto + Bonds (Extreme Heterogeneity)

Asset	Weight w_i	Volatility σ_i
BTC	0.05	0.80
ETH	0.05	0.90
US Treasuries	0.50	0.18
Investment Grade	0.40	0.05

Correlations: BTC-ETH 0.75, crypto-bonds −0.10, Treasuries-IG 0.60.

Moments. $\mathbb{E}[S] = 1.053$, $\sigma_S = 0.170$, $\gamma_3 = 2.4$.

Diagnostics. $H = 0.06$ (long-only, mostly positive correlations — low cancellation). CF guard: monotone at 1% but fails at 0.5% (skewness too high). Router: Spectral Fenton ($|\gamma_3| = 2.4 > 0.25$, exceeding the NIG comfort zone for a portfolio dominated by bond weights but with crypto tail exposure).

Risk measures.

Measure	SF Value	Loss	Gaussian Value	Gaussian Loss
VaR(99%)	0.712	28.8%	0.658	34.2%
VaR(97.5%)	0.741	25.9%	0.720	28.0%
ES(99%)	0.624	37.6%	0.597	40.3%
ES(97.5%)	0.662	33.8%	0.644	36.6%

The Gaussian approximation overestimates 99% VaR loss by 19%. This is the counterintuitive direction: for a right-skewed distribution, the matched variance absorbs right-tail mass that does not contribute to downside risk, making the Gaussian's symmetric tails too wide on the left. The ES/VaR loss ratio at 99% is 1.31, significantly above the Gaussian-implied 1.15, confirming that the tail is heavier relative to the quantile than a Gaussian would predict.

Spectral risk comparison. The exponential spectrum ($\beta = 10$) yields an ES-equivalent loss of 42.0%, which is 11% above the 97.5% ES loss of 37.6%. This gap quantifies the tail shape information that ES alone does not capture — the distribution has material probability mass in the extreme left tail beyond what ES averages over.

4.3 Long-Short Spread (Phase Cancellation)

Position	Weight w_i	Volatility σ_i
Long equity	+1	0.30
Short equity	−1	0.30

Correlation $\rho = 0.95$.

Moments. $\mathbb{E}[S] \approx 0$, $\sigma_S = 0.104$, $\gamma_3 = 0.10$.

Diagnostics. $H = 0.95$ (extreme cancellation — the hedged pair has nearly all cross-covariance cancelling). CF guard: monotone (passes). Router: Gaussian ($|\gamma_3| = 0.10 < 0.25$), **but** $H = 0.95 > 0.3$ **flags cancellation regime**, escalating to Spectral Fenton.

This case exposes a critical routing subtlety: skewness alone says "Gaussian," but the hedge index says "danger." The resolution:

Method	VaR(99%)	VaR loss	Error vs. SF
Spectral Fenton	−0.189	0.189	—
Gaussian	−0.242	0.242	+28% (overestimates)
NIG 3-moment	−0.195	0.195	+3%

The Gaussian overestimates the spread width by 28% because the distribution of $e^{Y_1} - e^{Y_2}$ with $\rho = 0.95$ is narrow and leptokurtic — thinner tails than Gaussian despite being nearly symmetric. The Spectral Fenton captures the destructive interference pattern exactly via Fourier-mode cancellation (see Nagy, 2026b, Figure 5). Without the hedge index guard, the router would have accepted the Gaussian VaR and overestimated capital by 28%.

4.4 Large Diversified ($n = 100$)

Equal-weight portfolio: $w_i = 1/100$, $\sigma_i \in [0.05, 1.0]$ (uniformly sampled). The correlation matrix is generated via the Lewandowski-Kurowicka-Joe (LKJ) distribution with $\eta = 2$ (moderate off-diagonal shrinkage toward zero), producing a random positive-definite correlation matrix with realistic block structure. Specifically: draw a random correlation matrix C from LKJ$(\eta = 2, n = 100)$ using the vine method (Lewandowski et al., 2009 [TODO:cite]), with NumPy random seed seed=42. The resulting matrix has eigenvalue spectrum dominated by 3–5 factors ($\lambda_1/\text{tr} \approx 0.35$), consistent with equity portfolios exhibiting sector and market-factor structure. To reproduce: the exact matrix is available in the supplementary code repository at examples/paper02_case_studies.py.

Diagnostics. For the specific draw (seed 42), $\gamma_3 = 1.8$ and $H = 0$ (equal-weight long-only with no short positions, hence no cancellation). The routing policy assigns this portfolio to the Spectral Fenton tier ($|\gamma_3| = 1.8 > 0.25$).

Computational profile. The eigenvalue spectrum of the random correlation matrix determines the cost:

Spectrum type	K (factors)	n_q	Q	Precompute	Per-query VaR
Market-factor dominant ($\lambda_1/\text{tr} > 0.45$)	1	16	16	$ $8 ms	0.46 ms
Two-sector	2	16	256	$ 50*ms*|0.46*ms*||*Flatspe*12*ctrum*(*1* {12}$)	

The VaR fan chart — VaR at 20 confidence levels from 0.5% to 10% — takes $20 \times 0.46 = 9.2$ ms from cached coefficients, regardless of portfolio size. For a 500-portfolio desk running daily fan charts, total: 500 × 9.2 ms = 4.6 seconds. For Monte Carlo at 10^6 paths per level: 500 × 20 × 660 ms = 110 minutes.

---

5. Stress Testing

5.1 Correlation-Volatility Heatmap

A 50×50 grid over two axes — correlation multiplier (0.5× to 2.0× baseline) and volatility multiplier (0.5× to 2.0× baseline) — produces a complete VaR stress surface. Each pixel is one Eigen-COS pipeline run.

Method	Per-pixel	Full 50 × 50 grid	Deterministic?
Spectral Fenton	$ $65 ms	$ $2.7 min	Yes
MC 10^6	$ $660 ms	$ $27 min	No
MC 10^5	$ $66 ms	$ $2.7 min	No

At equal wall-clock time, the SF grid is deterministic while the MC 10^5 grid carries $ $0.3% noise per pixel. Over 2,500 pixels, the MC surface has visible noise artifacts; the SF surface is smooth.

Figure 1: Stress heatmap for the crypto-bond portfolio: VaR(99%) loss across correlation and volatility multipliers. The baseline (current market) is marked with +. The non-linear interaction between correlation and volatility stress is visible: the upper-right corner (simultaneous high correlation and high volatility) produces losses far exceeding the sum of individual stresses.

The heatmap reveals non-linear VaR response to correlated stress that simple scaling rules miss. The baseline scenario (current market conditions) occupies one pixel. A "2008 crisis" scenario (correlation multiplier 1.8×, volatility multiplier 2.0×) and a "COVID crash" scenario (correlation 1.3×, volatility 2.0×, rapid onset) are identifiable regions on the heatmap where VaR increases sharply. The interaction between correlation and volatility stress is not additive — correlated increases in both produce VaR jumps larger than the sum of individual stress contributions. This non-linearity is invisible to the common practice of running volatility shocks and correlation shocks independently.

5.2 VaR Fan Chart

The VaR fan chart displays VaR at 20 confidence levels (0.5% to 10%) from the same 128 coefficients. Total computation: 20 × 0.46 ms = 9.2 ms. The result is a smooth, continuous loss exceedance curve — not a staircase of MC percentiles.

The SF fan chart has three practical advantages over its MC equivalent: (i) it is noise-free, producing smooth contours suitable for publication and regulatory reporting; (ii) it is available in under 10 ms, enabling interactive "what-if" exploration on a risk dashboard; and (iii) it captures the full quantile structure from a single precomputation, whereas MC must be re-run at each confidence level to avoid interpolation artifacts.

For the crypto-bonds portfolio (Section 4.2), the fan chart reveals that the loss exceedance curve steepens rapidly below the 2% level — the tail is significantly heavier than the body of the distribution. A risk manager examining only VaR at 99% and ES at 97.5% would miss this: the fan chart communicates the tail shape visually and immediately.

Figure 2: VaR fan chart for the crypto-bond portfolio: loss at 20 confidence levels from the same 130 cached coefficients (navy, smooth) vs. Gaussian VaR (red, dashed). The FRTB level $\alpha = 2.5\%$ is marked. The Gaussian underestimates tail losses and overestimates body losses — a consequence of fitting a symmetric distribution to a skewed one. Total computation: 9.2 ms.

---

6. The Extreme-Volatility Frontier

For portfolios with any asset volatility $\sigma_i > 1.0$ (corresponding to $|\gamma_3| \gtrsim 10$), no analytical method we tested achieves $< 10\%$ relative VaR error at the 1% level. This includes the Spectral Fenton.

Table 4. VaR error at 99% for extreme-volatility portfolios.

Method	L4 error ($\sigma \in [0.8, 1.5]$)	L7 error ($\sigma \in [1.5, 2.0]$)
Cornish-Fisher	7,200%	97,800×
Fenton-Wilkinson	35%	56%
NIG 3-moment	routed to spectral	routed to spectral
Spectral Fenton ($K = 2$, $n_{\mathrm{gh}} = 32$)	16%	65%
Spectral Fenton ($K = 3$, $n_{\mathrm{gh}} = 64$)	16%	65%

Root cause. The COS domain $[a, b]$ must span \$ 21 \$ of the underlying Fenton-Wilkinson distribution to capture the support. For $\sigma_{\mathrm{FW}} > 1.0$, this domain is so wide that only \$ \$6 Fourier terms per sigma of resolution remain near the 1% quantile. This is a COS-intrinsic limitation (Nagy, 2026a, Proposition 7): the sine-series CDF is mathematically identical to the COS payoff formula with indicator payoff, so no reformulation within the COS framework can improve resolution.

What this means for practice. The extreme-volatility regime covers highly leveraged positions, deep OTM options on individual names, and certain cryptocurrency portfolios where annualized volatility exceeds 100%. For these cases, the routing policy (Table 3) flags $\sigma_{\max} > 1.0$ and recommends MC validation alongside the SF estimate. The SF still provides a deterministic point estimate (useful as a model risk check against MC), but the 16–65% error is too large for stand-alone regulatory reporting.

Potential approaches for future research. Saddlepoint approximation targets the quantile directly without a domain, bypassing the resolution trade-off. Importance-sampling hybrids could use the SF as a control variate for MC in the extreme tail. Adaptive domain splitting — separate COS expansions for the body and the tail — could concentrate Fourier resolution where it is needed. None of these are implemented; the extreme-volatility frontier remains an open research problem.

---

7. Regulatory Framework

7.1 Basel III/FRTB Compliance

Under the Fundamental Review of the Trading Book (Basel Committee, 2019; Hull, 2018, Chapter 17), banks must compute:

- **ES at** $\alpha = 2.5\%$ (the primary market risk measure, replacing VaR)
- **VaR at 99% and 97.5%** (for backtesting and internal risk limits)
- **Stressed ES** under historically adverse correlation and volatility scenarios

The Spectral Fenton satisfies all three requirements. VaR and ES are deterministic — no seed dependence, no convergence uncertainty, no need to document the number of MC paths. Stressed scenarios are computable in real time via the correlation-volatility heatmap of Section 5.1: each stress scenario is one Eigen-COS pipeline run ($ $65 ms), enabling overnight computation of full stress surfaces for the entire trading book.

The deterministic output has a specific regulatory advantage: reproducibility. Two independent implementations with the same portfolio inputs and the same 128 COS terms produce identical VaR and ES values to machine precision. Monte Carlo outputs depend on the random seed, the path count, and the variance reduction technique — all of which are implementation choices that introduce operational risk into regulatory reporting (Danielsson et al., 2016).

7.2 The Risk Certificate

The 130-number Spectral Fenton representation is a compact, auditable "risk certificate" (the CDF reconstruction from 130 numbers and the 1.04 KB size bound are Lean-verified: RiskCertificate.cdf, certificate_size in RiskCertificate.lean):

Property	Value
Size	128 floats + 2 bounds = 130 numbers (1.04 KB)
Content	Complete distributional summary of portfolio value
Verification	Regulator computes VaR/ES independently from the 130 numbers
Privacy	Portfolio composition (w, σ, C) not revealed
Cacheability	Key-value store (Redis, memcached), sub-ms retrieval
Determinism	Same inputs $\rightarrow$ same 130 numbers $\rightarrow$ same risk measures

A bank provides the risk certificate to the regulator; the regulator independently computes VaR, ES, and any spectral risk measure from the 130 numbers, without access to portfolio weights, asset volatilities, or the correlation matrix. This is a novel form of auditable risk disclosure: the regulator can verify the risk computation without seeing the portfolio, and the bank cannot manipulate the risk certificate without changing the underlying distribution.

The auditing workflow: (i) the bank runs the Eigen-COS pipeline and stores the 130-number certificate alongside the reported VaR and ES; (ii) the regulator receives the certificate and independently evaluates the sine-series CDF at the reported VaR point, confirming $F(\text{VaR}_{\text{reported}}) = \alpha$; (iii) the regulator evaluates the closed-form ES at the reported VaR, confirming the reported ES. Both verifications take < 1 ms and require no proprietary information about the bank's portfolio.

7.3 Model Risk Considerations

The primary model risk is the geometric Brownian motion assumption: asset prices are lognormal with constant volatility and correlation. For assets with jump risk (M&A events, credit defaults), fat-tailed marginals (cryptocurrency, distressed debt), or stochastic volatility (equity indices over long horizons), the lognormal conditional characteristic function must be replaced by the appropriate marginal CF.

The eigenvalue-conditional architecture of Section 1.2 accommodates this extension. The outer structure — correlation eigendecomposition, Gauss-Hermite conditioning, Fourier-cosine inversion — is independent of the marginal distribution (Nagy, 2026a, Section 7.2). Only the inner Gauss-Hermite integral changes: replacing $\phi_{\text{LN}}(t; \mu, \sigma)$ with $\phi_{\text{NIG}}(t; \alpha, \beta, \mu, \delta)$ or $\phi_t(t; \nu, \mu, \sigma)$ extends the framework to NIG or Student-t marginals respectively. The output format — 130 spectral coefficients — remains identical.

The risk manager should document: (i) the GBM assumption and the portfolio conditions under which it is appropriate, (ii) the skewness routing decision (Table 3), (iii) the hedge index H and its implications for parametric accuracy, and (iv) the extreme-volatility flag when $\sigma_{\max} > 1.0$. These diagnostics are automatically available from the precomputation step and constitute a model-risk self-assessment that is richer than the typical Monte Carlo documentation (McNeil et al., 2015, Chapter 2).

---

8. Conclusion

We have presented a noise-free risk measurement framework for lognormal portfolios, built on the Spectral Fenton Distribution (Nagy, 2026a) and the Eigen-COS algorithm (Nagy, 2026b). The framework contributes six results not present in the companion papers:

1. **Closed-form ES** from the Fourier-sine CDF via term-by-term integration, providing the first closed-form Expected Shortfall for sums of correlated lognormals.
2. **Backtest power analysis.** Deterministic ES eliminates the 1–8% simulation noise that MC introduces at typical production path counts (10^3–10^5), removing the single largest source of power degradation in the Acerbi-Székely ridge backtest.
3. **Hedge index H.** A scalar diagnostic that flags cancellation-dominated portfolios where parametric VaR fails by 20%+ despite low skewness. The $H > 0.3$ threshold triggers escalation to spectral methods.
4. **Validated routing.** A skewness-and-cancellation-aware routing policy (Gaussian / NIG / Spectral Fenton) validated across 48 of 50 calibration gym levels (Nagy, 2026b).
5. **Real-time stress testing.** A 50×50 correlation-volatility heatmap in 2.7 minutes (deterministic), and a 20-level VaR fan chart in 9.2 ms from cached coefficients.

6. **Risk certificate.** The 130-number Spectral Fenton representation as a compact, auditable regulatory output that enables independent verification without portfolio disclosure.

For the practitioner: the framework is most valuable when the portfolio has $|\gamma_3| > 0.25$ (enough skewness to make the Gaussian unreliable), $\sigma_{\max} < 1.0$ (within the spectral method's domain resolution), and the desk requires repeated VaR/ES queries (stress testing, fan charts, intraday monitoring). For near-Gaussian portfolios, the routing policy correctly defaults to simpler methods. For extreme-volatility portfolios, the framework provides a diagnostic flag and recommends MC validation.

The Spectral Fenton does not replace Monte Carlo. It complements it, covering the high-volume, time-critical segment of the risk computation stack — daily VaR/ES for linear portfolios under GBM — with a deterministic alternative that is faster, reproducible, and maximally powerful for regulatory backtesting.

Formal Verification

The mathematical foundations of this paper are supported by Lean 4 machine-checked proofs. The core theoretical results (Mixture Collapse, well-posedness, uniqueness, error decomposition, all four Acerbi coherence axioms) are verified in the companion paper (Nagy, 2026a). Results specific to this paper have additional Lean verification, summarized below with an honest characterization of proof depth:

Paper 02 claim	Lean theorem	File	Depth
Hedge index $H \in [0,1]$	hedge_index_in_unit	HedgeIndex.lean	**Genuine** — uses div_nonneg, div_le_one, Finset sum bounds
Risk certificate: CDF from 130 numbers	RiskCertificate.cdf	RiskCertificate.lean	**Genuine** — correct Lean definition of sine-series CDF
Certificate size = 1.04 KB	certificate_size	RiskCertificate.lean	Arithmetic — 130 * 8 = 1040 by norm_num
Noise inflates variance	noise_increases_variance	NoiseInflatesVariance.lean	Arithmetic — a < a + b by linarith
$\sigma_T^2 > \sigma_{\text{model}}^2$ under noise	backtest_power_reduced	NoiseInflatesVariance.lean	Arithmetic — same level as above

Honest assessment of proof depth. The hedge index bound and the CDF definition are non-trivial formalizations that would catch genuine mathematical errors. The variance inequalities, while correct and relevant, formalize arithmetic identities rather than deep mathematical arguments — the substantive content of the backtest power analysis (Section 2.3) rests on the analytic approximation described in the Table 2 methodology, not on the Lean proofs alone. The leap from "variance increases" to specific power degradation percentages involves non-formalized statistical reasoning (normal approximation to the ridge test statistic, calibration of σ_{model}). We note this gap explicitly: the Lean verification covers the algebraic foundations, not the full statistical argument.

The full Lean formalization (430+ theorems across 85 files, covering all papers in the suite) is available as supplementary material. See Nagy (2026a) for verification details. The formalization was verified by lake build (Lean 4.28.0, Mathlib v4.28.0) on March 3, 2026, producing 0 errors.

Acknowledgements

The author used Large Language Models for assistance with language editing and coding. The author assumes full responsibility for the scientific content.

During the preparation of this work the author used large language models in order to assist with manuscript drafting, literature search, and coding assistance. After using these tools, the author reviewed and edited the content as needed and takes full responsibility for the content of the published article.

References

- Acerbi, Carlo (2002). Spectral Measures of Risk: A Coherent Representation of Subjective Risk Aversion. *Journal of Banking & Finance*, 26(7), 1505-1518. DOI: 10.1016/S0378-4266(02)00281-9
- Acerbi, Carlo and Sz{'e (2014). Back-testing Expected Shortfall. *Risk*.
- Acerbi, C. and Székely, B (2019). The minimally biased backtest for expected shortfall. *Acerbi, C. and Székely, B.*.
- Acerbi, Carlo and Sz{'e (2023). Backtestability and the ridge backtest. *Frontiers of Mathematical Finance*, 2, 2023019. DOI: 10.1142/9789811280306_0003
- Basel Committee on Banking Supervision (2019). Minimum capital requirements for market risk. Bank for International Settlements.
- Danielsson, Jon and James, Kevin R. and Valenzuela, Marcela and Zer, Ilknur (2016). Model risk of risk models. *Journal of Financial Stability*, 23, 79-91. DOI: 10.2139/ssrn.2425689
- Dufresne, Daniel (2004). The Log-Normal Approximation in Financial and Other Computations. *Advances in Applied Probability*, 36, 747-773. DOI: 10.1239/aap/1093962232
- 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
- 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/t-com.1960.1097606
- Glasserman, Paul (2003). Monte Carlo Methods in Financial Engineering. Springer.
- Hull, J.C (2018). Risk Management and Financial Institutions*, 5th ed. Wiley.* Hull, J.C.*.
- Kusuoka, Shigeo (2001). On law invariant coherent risk measures. *Advances in Mathematical Economics*, 3, 83-95. DOI: 10.1007/978-4-431-67891-5_4

- McNeil, A.J., Frey, R., and Embrechts, P (2015). Quantitative Risk Management: Concepts, Techniques and Tools*, revised ed. Princeton University Press.* McNeil, A.J., Frey, R., and Embrechts, P.*.
- Nagy, T (2011). Probability density function for (weighted) sum of n correlated lognormal variables. *Nagy, T.*.
- Nagy, T. (2026). The Exact Latent Distribution of Correlated Lognormal Sums. *Working paper.*
- Nagy, T. (2026). Exact Portfolio VaR Without Monte Carlo: The Eigen-COS Method. *Zenodo.* DOI: 10.5281/zenodo.18910516