Nonlinear Portfolio Risk in Closed Form

Exact Distribution, VaR, ES, and Cross-Correlation of Derivative Portfolios via the Hermite-COS Framework

The GH grid is the Latent of the market — any payoff is just a projection

Tamás Nagy, Ph.D.

tnagyphd@gmail.com

Draft

Abstract

We extend the Hermite-COS framework from linear portfolios (weighted sums of correlated lognormals) to portfolios containing options, structured products, and arbitrary derivative payoffs. The key observation is that the Gauss-Hermite quadrature grid — which evaluates the underlying Gaussian variables at Q^n deterministic nodes — computes the joint realization of ALL underlying assets simultaneously. Any payoff function $g(S_1, \dots, S_n)$ can be evaluated on these same nodes without recomputing the grid. The COS inversion then produces the full distribution, VaR, and ES of the derivative portfolio.

The deeper consequence is exact cross-correlation between arbitrary derivative payoffs. Given two options $V_1 = g_1(S)$ and $V_2 = g_2(S)$, their correlation is:

$$\rho(V_1, V_2) = \frac{\sum_\ell w_\ell g_1(S_\ell) g_2(S_\ell) - (\sum_\ell w_\ell g_1(S_\ell))(\sum_\ell w_\ell g_2(S_\ell))}{\sqrt{\mathrm{Var}_w(g_1) \cdot \mathrm{Var}_w(g_2)}}$$

This gives practitioners the exact correlation matrix of an option book — something previously available only through Monte Carlo simulation. The entire computation is deterministic, subsecond, and requires no additional infrastructure beyond the linear portfolio framework.

We demonstrate on a mixed portfolio of spot positions, call options, put options, and a collar, computing the joint distribution and the full correlation matrix of all positions.

Keywords: derivative portfolio risk, option correlation, nonlinear payoff, closed-form distribution, Gauss-Hermite quadrature.

---

1. Introduction

1.1 The Gap

The linear portfolio risk framework (Nagy, 2026c) computes the exact distribution of $S = \sum w_i e^{Y_i}$. But real portfolios contain options, futures, swaps, and structured products. The portfolio value is:

$$V = \sum_{j=1}^{m} g_j(S_1, \dots, S_n)$$

where g_j are nonlinear payoff functions:

Instrument	Payoff $g(S)$
Spot position	$w_i S_i$
European call	$\max(S_i - K, 0)$
European put	$\max(K - S_i, 0)$
Call spread	$\max(S_i - K_1, 0) - \max(S_i - K_2, 0)$
Collar	$S_i + \max(K_p - S_i, 0) - \max(S_i - K_c, 0)$
Straddle	$\lvert S_i - K \rvert$
Basket option	$\max(\sum w_i S_i - K, 0)$
Worst-of put	$\max(K - \min_i S_i, 0)$

Monte Carlo can compute the distribution of V by sampling $Y \sim N(\mu, \Sigma)$ and evaluating g. But this requires 10^5–10^6 paths, is noisy in the tails, and seed-dependent. The delta-gamma approximation is fast but inaccurate for large moves.

1.2 The Key Insight

The Gauss-Hermite quadrature grid from the Hermite-COS pipeline already evaluates the underlying Gaussian vector Y at Q^n deterministic nodes (z_ℓ, w_ℓ). At each node, ALL underlying asset prices are known:

$$S_{i,\ell} = e^{\mu_i + (\text{chol}(\Sigma) \cdot z_\ell)_i}, \quad i = 1, \dots, n$$

Any function of these prices can be evaluated at the same nodes:

$$V_\ell = g(S_{1,\ell}, \dots, S_{n,\ell})$$

The COS inversion or KDE then produces the full distribution of V. No new quadrature. No new eigendecomposition. Just one function evaluation per node.

1.3 From the Latent Perspective

The GH grid IS the Latent representation of the joint market — a finite, deterministic encoding of the multivariate Gaussian measure. Every derivative is a projection of this Latent:

$$\Lambda_g = \{w_\ell, g(S_\ell)\}_{\ell=1}^{Q^n}$$

The linear portfolio is the projection $g(S) = \sum w_i S_i$. A call option is the projection $g(S) = \max(S_1 - K, 0)$. A basket option is $g(S) = \max(\sum w_i S_i - K, 0)$. The Latent doesn't change — only the projection does.

This is a **grade-2 operation** in the Latent hierarchy: the representation (what we store) stays fixed; only the extraction (what we compute from it) changes.

---

2. The Extended Framework

2.1 Setup

The only change to the linear pipeline is at Step 3:

Step	Linear	Nonlinear
1. Quadrature	z_ℓ, w_ℓ from GH	**same**
2. Asset prices	$S_{i,\ell} = w_i e^{Y_{i,\ell}}$	$S_{i,\ell} = e^{Y_{i,\ell}}$ (raw prices)
3. Portfolio	$V_\ell = \sum w_i S_{i,\ell}$	$V_\ell = g(S_{1,\ell}, \dots, S_{n,\ell})$
4. Distribution	COS/KDE on $\{w_\ell, V_\ell\}$	**same**

2.2 Cross-Correlation of Derivatives

Given m derivative positions with payoffs $g_1, \dots, g_m$, the exact moments are:

$$E[g_j] = \sum_\ell w_\ell \, g_j(S_\ell)$$

$$E[g_j g_k] = \sum_\ell w_\ell \, g_j(S_\ell) \, g_k(S_\ell)$$

$$\text{Cov}(g_j, g_k) = E[g_j g_k] - E[g_j]E[g_k]$$

$$\rho(g_j, g_k) = \frac{\text{Cov}(g_j, g_k)}{\sqrt{\text{Var}(g_j)\text{Var}(g_k)}}$$

This produces the **exact correlation matrix** of the derivative book from a single GH grid evaluation. The cost is $O(m^2 \cdot Q^n)$ — quadratic in the number of positions, linear in the quadrature grid size.

2.3 Why This Is Not Trivial

The correlation between two options on correlated underlyings is NOT a simple function of the underlying correlation. Consider:

- BTC call (strike K_1, on S_1) and ETH put (strike K_2, on S_2) with $\rho(S_1, S_2) = 0.75$.
- When BTC rises above K_1: call payoff > 0. But BTC rising implies ETH likely rises too ($\rho = 0.75$), which means put payoff $= 0$.
- The OPTION correlation is strongly negative, even though the underlying correlation is positive.

Computing this analytically requires evaluating the bivariate normal integral with nonlinear boundaries — doable for two assets but intractable for $n > 3$ with arbitrary payoffs. The GH grid computes it exactly for any n and any payoff.

---

3. Demonstration: Mixed Derivative Portfolio

3.1 Portfolio

Position	Instrument	Underlying	Strike	Notional
1	BTC spot	BTC	—	\$400K
2	ETH call	ETH	1.1	\$200K
3	S&P put (hedge)	S&P	0.95	\$300K
4	BTC-ETH call spread	basket	1.0/1.3	\$100K

Underlyings: BTC ($\sigma = 0.80$), ETH ($\sigma = 0.90$), S&P ($\sigma = 0.20$). Correlations: BTC-ETH = 0.75, BTC-S&P = 0.10, ETH-S&P = 0.08.

3.2 What the Framework Produces

1. **Full distribution** of portfolio value — a single PDF/CDF capturing the joint effect of all positions.
2. **VaR and ES** at any confidence level — including the diversification benefit of the hedge put.
3. **Position-level correlation matrix:**

	BTC spot	ETH call	S&P put	Basket spread
BTC spot	1.00	ρ_{12}	ρ_{13}	ρ_{14}
ETH call		1.00	ρ_{23}	ρ_{24}
S&P put			1.00	ρ_{34}
Basket spread				1.00

Each ρ_{jk} is computed exactly from the GH grid.

4. **Marginal risk contribution:** how much each position contributes to portfolio VaR.

3.3 What Practitioners Gain

- **Hedge effectiveness:** the correlation between the S&P put and the rest of the portfolio shows exactly how much protection the hedge provides — not in a delta-approximation, but for the full nonlinear payoff.
- **Diversification across payoff types:** a call option and a put option on correlated underlyings have complex correlation that depends on moneyness, volatility, and the joint distribution. The framework computes this exactly.
- **Portfolio-level Greeks are implicit:** the sensitivity of portfolio VaR to any parameter change is computable by re-evaluating g on the same grid with perturbed parameters.

---

4. Theoretical Foundations

4.1 Convergence

Proposition 1 (Nonlinear Extension). If $g : \mathbb{R}_+^n \to \mathbb{R}$ is piecewise smooth with at most polynomial growth, and the underlying follows $Y \sim N(\mu, \Sigma)$, then the GH quadrature:

$$E[g(e^Y)] = \sum_\ell w_\ell \, g(e^{\mu + L z_\ell}) + O(Q^{-(2Q - d_g)})$$

where d_g is the effective polynomial degree of $g \circ \exp$ and $L = \text{chol}(\Sigma)$.

For option payoffs ($g = \max(\cdot, 0)$), the integrand is continuous but has a kink at the strike. The convergence rate drops from exponential (smooth g) to algebraic ($\sim Q^{-2}$), but remains fast in practice: $Q = 30$ gives 6+ digit accuracy for standard options.

4.2 The Latent Composition Theorem

Theorem 1 (Latent Composition). Let $\Lambda = \{w_\ell, z_\ell\}_{\ell=1}^{Q^n}$ be the GH Latent of the Gaussian measure $N(\mu, \Sigma)$. For any measurable $g : \mathbb{R}^n \to \mathbb{R}$ with $E[|g(e^Y)|] < \infty$:

$$\Lambda_g = \{w_\ell, g(e^{\mu + L z_\ell})\}_{\ell=1}^{Q^n}$$

is a Latent representation of the distribution of $g(e^Y)$.

The key property: the INDEX SET $\{w_\ell, z_\ell\}$ is universal — it depends only on (μ, Σ), not on g. Evaluating 100 different payoffs on the same grid costs 100 function evaluations, not 100 separate Monte Carlo simulations.

4.3 Cross-Correlation as Inner Product

The cross-correlation structure has a clean algebraic interpretation. Define the weighted inner product:

$$\langle g_1, g_2 \rangle_\Lambda = \sum_\ell w_\ell \, g_1(S_\ell) \, g_2(S_\ell)$$

Then $E[g_j g_k] = \langle g_j, g_k \rangle_\Lambda$ and the correlation matrix is the normalized Gram matrix:

$$\rho_{jk} = \frac{\langle g_j, g_k \rangle_\Lambda - \langle g_j, 1 \rangle_\Lambda \langle g_k, 1 \rangle_\Lambda}{\sqrt{(\langle g_j, g_j \rangle_\Lambda - \langle g_j, 1 \rangle_\Lambda^2)(\langle g_k, g_k \rangle_\Lambda - \langle g_k, 1 \rangle_\Lambda^2)}}$$

This is a **reproducing kernel inner product** on the space of payoff functions, with the GH grid providing the reproducing kernel. The correlation matrix of any option book is a Gram matrix in this inner product space.

---

5. Comparison with Existing Methods

Method	Handles nonlinear?	Cross-correlations?	Deterministic?	Speed
Delta-Normal	No (linearizes)	Approximate	Yes	ms
Delta-Gamma	Partially	Approximate	Yes	ms
Monte Carlo	Yes	Yes (noisy)	No	seconds–minutes
Historical simulation	Yes	Data-limited	N/A	seconds
Hermite-COS (this paper)	**Yes (exact)**	**Yes (exact)**	**Yes**	**$<$ 1 second**

The delta-gamma approximation captures second-order effects but fails for deep out-of-the-money options, barrier products, and portfolios with discontinuous payoffs. Monte Carlo handles everything but is noisy and expensive. The Hermite-COS framework combines the exactness of full revaluation with the speed and determinism of analytical methods.

---

6. Limitations

1. **GBM assumption:** underlyings follow correlated geometric Brownian motion. Extensions to stochastic volatility or jump-diffusion require modifying the inner quadrature.
2. **Dimensionality:** direct tensor-product GH scales as Q^n. For $n > 8$, eigenvalue conditioning reduces the effective dimension to $K \leq 3$.
3. **Path-dependent payoffs:** Asian options, barrier options, and lookbacks depend on the price path, not just the terminal distribution. These require extending the framework to multi-step quadrature (replacing one GH grid with a sequence).
4. **American exercise:** early exercise requires dynamic programming, which the single-step framework does not address. Longstaff-Schwartz regression could be adapted to use GH nodes instead of MC paths.

---

7. Conclusion

The Hermite-COS framework extends naturally from linear to nonlinear portfolios. The GH quadrature grid is a universal Latent representation of the joint market distribution — any payoff function is just a projection. This gives practitioners, for the first time:

1. The exact distribution of a derivative portfolio without Monte Carlo.
2. The exact cross-correlation matrix of an option book.
3. Sub-second, deterministic computation.

The conceptual insight is that the Latent of the market (the GH grid) is **deeper** than any individual portfolio or derivative. It encodes the full joint distribution; individual positions are views into this shared representation. Risk management is Latent projection.

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

- Nagy, T. (2026). Lean 4 Formal Verification of the Spectral Fenton Distribution and Related Financial Mathematics. *Working paper.*
- Nagy, T. (2026). The Fenton Distribution Solved. *Working paper.*
- Nagy, T. (2026). Nonlinear Portfolio Risk in Closed Form. *Working paper.*
- 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
- Longstaff, F.A., & Schwartz, E.S (2001). Valuing American options by simulation: A simple least-squares approach. *Review of Financial Studies*, 14(1), 113-147. DOI: 10.1093/rfs/14.1.113