Generative Portfolio Design: Inverting the Spectral Fenton Representation

Tam\{}\{}'as Nagy, Ph.D.

tnagyphd@gmail.com

Draft

Abstract

Classical portfolio optimization minimizes a single risk metric — variance, VaR, or CVaR — subject to return constraints. This produces portfolios that are optimal in one dimension but unconstrained in all others: two portfolios with identical VaR can have radically different skewness, kurtosis, or tail structure. We propose **generative portfolio design**: specifying the *entire* loss distribution shape as the optimization target, then solving the **inverse spectral problem** to find portfolio weights that produce it.

The Spectral Fenton Distribution (Nagy, 2026a) represents any portfolio of correlated lognormal assets by 128 Fourier-cosine coefficients $\{A_k\}_{k=0}^{127}$. The forward map $\mathcal{F} : \mathbf{w} \mapsto \{A_k\}$, computed by the Eigen-COS algorithm, is differentiable in the portfolio weights. We formulate the inverse problem as constrained nonlinear least squares — $\min_{\mathbf{w}} \|\mathcal{F}(\mathbf{w}) - A_k^*\|_2^2$ subject to $\sum w_i = 1$, $w_i \geq 0$ — and solve it via Sequential Least Squares Programming (SLSQP).

Three applications are demonstrated: (i) **distribution matching** — replicate a target portfolio's risk profile with different assets, achieving 0.015% relative L_2 coefficient error and $< 10^{-4}$ VaR difference; (ii) **targeted tail reduction** — reduce the high-frequency spectral energy (which encodes tail risk) while preserving the distribution body; and (iii) **spectral risk budgeting** — decompose each asset's contribution to each Fourier mode, revealing *which* assets drive *which* frequencies of the risk profile. For a 5-asset portfolio, the full inverse solve takes \$ \$5 seconds; the spectral risk budget requires $n+1$ forward evaluations (\$ \$50 ms total). We show that mean-variance, mean-CVaR, and distributionally robust optimization are recovered as special cases corresponding to constraints on progressively more spectral modes.

Key Messages

- Design portfolios by specifying the entire loss distribution shape, not just one number
- 128 Fourier coefficients are the complete "distributional genome" of a portfolio
- Spectral risk budgeting reveals which assets drive which frequencies of the risk profile
- Classical optimizations (mean-variance, mean-CVaR) are special cases constraining a subset of spectral modes

---

1. Introduction

1.1 The Limitation of Single-Metric Optimization

Portfolio optimization since Markowitz (1952) has operated in a common paradigm: define an objective function of portfolio weights (variance, VaR, CVaR, drawdown) and minimize it subject to constraints. The result is a portfolio that is optimal along one dimension of its loss distribution, but whose behavior along all other dimensions is a byproduct rather than a design choice.

This creates a practical problem. A risk manager who minimizes VaR at $\alpha = 5\%$ says nothing about the distribution beyond that quantile. Two min-VaR portfolios can have identical VaR yet radically different behavior: one with a thin Gaussian tail, another with a fat lognormal tail that generates catastrophic losses at $\alpha = 1\%$. The portfolio was designed to optimize one number, so only one number is controlled.

1.2 From Numbers to Shapes

The Spectral Fenton Distribution (Nagy, 2026a) represents the full loss distribution of a portfolio of n correlated lognormal assets by 128 Fourier coefficients $\{A_k\}_{k=0}^{127}$ and two domain bounds $[a, b]$. This representation is:

- **Complete**: the coefficients uniquely determine the CDF, PDF, VaR at any level, ES, and all spectral risk measures (Proposition 3 in Nagy, 2026a).
- **Compact**: 130 numbers replace the $n(n+3)/2$ raw parameters (= 5,150 for $n = 100$), a compression ratio $> 39\times$ (Theorem 4 in Nagy, 2026a).
- **Structured**: low-frequency coefficients ($k = 0, \dots, 10$) encode the distribution body (location, scale, skewness); high-frequency coefficients ($k > 10$) encode the tail structure.

The existence of this representation transforms risk from a *number* (VaR = \$X) into a *shape* — a point in $\mathbb{R}^{128}$. And once risk has geometry, the question changes: instead of asking "which portfolio minimizes this one number?", we can ask "which portfolio produces this shape?"

1.3 The Inverse Spectral Problem

The forward map $\mathcal{F} : \mathbb{R}^n \to \mathbb{R}^{128}$ takes portfolio weights $\mathbf{w}$ and produces spectral coefficients $\{A_k\}$ via the Eigen-COS algorithm. This paper studies the **inverse map**: given a target coefficient vector A_k^*, find weights $\mathbf{w}^*$ such that $\mathcal{F}(\mathbf{w}^*) \approx A_k^*$.

This is a nonlinear inverse problem because the map $\mathcal{F}$ passes through eigenvalue decomposition, Gauss-Hermite quadrature, and Fourier inversion — none of which is linear in $\mathbf{w}$. We solve it as constrained nonlinear least squares:

$$\min_{\mathbf{w} \in \mathcal{W}} \sum_{k=0}^{N-1} \left(A_k(\mathbf{w}) - A_k^*\right)^2$$

where $\mathcal{W} = \{\mathbf{w} \in \mathbb{R}^n : \sum_i w_i = 1,\ w_i \geq 0\}$ is the portfolio simplex (with optional sector constraints).

1.4 Three Applications

We demonstrate three use cases:

1. **Distribution matching**: replicate a target portfolio's entire risk profile using different assets. Use case: "I want a portfolio that looks like SPY but uses only European equities."
2. **Targeted tail reduction**: modify the high-frequency coefficients (which encode tail risk) while preserving the low-frequency body. Use case: "same median return, 50% less tail risk."
3. **Spectral risk budgeting**: decompose each asset's contribution to each Fourier mode. Use case: "which asset is responsible for the skewness in my portfolio's loss distribution?"

1.5 Related Work

Portfolio optimization. Markowitz (1952) introduced mean-variance optimization. Rockafellar and Uryasev (2000) reformulated CVaR optimization as linear programming. Bertsimas, Lauprete, and Samarov (2004) developed shortfall-constrained optimization. All optimize for a single risk metric.

Higher-moment and spectral approaches. Acerbi (2002) introduced spectral risk measures — coherent risk measures defined by a weighting function over the quantile distribution — and showed that VaR, CVaR, and exponential risk measures are special cases. Kusuoka (2001) proved that every law-invariant coherent risk measure admits a representation as a supremum over weighted CVaR integrals, establishing the theoretical foundation for spectral risk. Our framework connects to this literature by providing explicit Fourier-domain control over the loss distribution that spectral risk measures evaluate. Harvey, Liechty, Liechty, and Müller (2010) and Jondeau and Rockinger (2006) addressed higher-moment portfolio optimization beyond mean-variance; our approach subsumes these by controlling all moments simultaneously through the full spectral representation.

Distributional constraints. Distributionally robust optimization (DRO) optimizes over ambiguity sets defined by moment conditions (Delage and Ye, 2010; Wiesemann, Kuhn, and Sim, 2014) or Wasserstein balls (Blanchet and Murthy, 2019). DRO constrains the distribution indirectly via moments or distances; our approach constrains it directly via the full spectral representation. Wasserstein DRO is the closest classical precursor to generative design — it provides approximate distributional control — but it operates on a worst-case basis rather than targeting a specific distribution shape.

Risk budgeting. Maillard, Roncalli, and Teïlétche (2010) introduced risk parity via variance contributions. McNeil, Frey, and Embrechts (2015) decomposed ES by Euler allocation. These decompose risk into a single number per asset; spectral risk budgeting provides a *matrix* — one number per asset per Fourier mode.

Characteristic function methods. Fang and Oosterlee (2008) developed the COS method for option pricing using Fourier-cosine expansions. Carr and Madan (1999) used the characteristic function for option pricing via Fourier transform. Our Eigen-COS algorithm extends these characteristic function methods from single-asset pricing to multi-asset portfolio distribution recovery.

Generative design in engineering. Topology optimization (Bendsøe and Sigmund, 2003) designs structures by specifying desired properties. Inverse problems in imaging (Tarantola, 2005) reconstruct sources from observations. Our approach applies the same paradigm to portfolio risk: specify the desired risk profile, solve for the portfolio that produces it.

1.6 Contributions

1. Formulation of the inverse spectral problem for portfolio design (Section 2).

2. Three applications with numerical results (Sections 3–5).
3. A unifying perspective: classical portfolio optimization as constrained spectral optimization (Section 6).
4. Open-source implementation in the spectral_fenton Python library.

---

2. Mathematical Framework

2.1 The Forward Map

The Eigen-COS algorithm (Nagy, 2026a,b) computes the spectral representation of a portfolio:

$$\mathcal{F} : (\mathbf{w}, \mu, \sigma, C) \mapsto \{A_k\}_{k=0}^{N-1}$$

where $\mathbf{w} \in \mathbb{R}^n$ are portfolio weights, $\mu, \sigma \in \mathbb{R}^n$ are asset means and volatilities, $C \in \mathbb{R}^{n\times n}$ is the correlation matrix, and $\{A_k\}$ are the Fourier-cosine coefficients. The algorithm proceeds through:

1. Eigendecomposition of $C = V\Lambda V^T$ (reusable when C is fixed).
2. Gauss-Hermite conditioning on the dominant K eigenvectors.
3. Per-scenario characteristic function evaluation.
4. Fourier-cosine inversion to obtain A_k.

Each step is differentiable in $\mathbf{w}$, enabling gradient computation.

2.2 The Inverse Problem

Definition 1 (Inverse Spectral Problem). *Given market parameters (μ, σ, C), a target coefficient vector $A^* \in \mathbb{R}^N$, and a constraint set $\mathcal{W} \subseteq \mathbb{R}^n$, find*

$$\mathbf{w}^* = \arg\min_{\mathbf{w}\in\mathcal{W}} \|\mathcal{F}(\mathbf{w}) - A^*\|_2^2.$$

The standard constraint set is the portfolio simplex:

$$\mathcal{W} = \left\{ \mathbf{w} \in \mathbb{R}^n : \sum_{i=1}^{n} w_i = 1,\ w_i \geq 0\ \forall i \right\}$$

which can be augmented with sector limits $l_j \leq \sum_{i\in S_j} w_i \leq u_j$ for asset groups S_j.

Proposition 1 (Non-uniqueness). *The inverse problem is generically non-unique when $n > 128$: the simplex $\mathcal{W} \subset \mathbb{R}^n$ has dimension $n-1$, while the target space $\mathbb{R}^N$ has dimension $N = 128$. For $n > 129$, infinitely many weight vectors produce the same spectral representation.*

Proof. The map $\mathcal{F} : \mathbb{R}^n \to \mathbb{R}^{128}$ has rank at most 128. When $n - 1 > 128$, the level set $\mathcal{F}^{-1}(A^*)$ generically has positive dimension. □

This non-uniqueness is a feature: it means the solver can find the target distribution while simultaneously satisfying additional constraints (sector limits, liquidity, transaction costs).

2.3 Gradient Computation

The Jacobian $J_{ki} = \partial A_k / \partial w_i$ can be computed by:

Finite differences (implemented):

$$\frac{\partial A_k}{\partial w_i} \approx \frac{A_k(\mathbf{w} + h\mathbf{e}_i) - A_k(\mathbf{w})}{h}, \quad h = 10^{-5}$$

requiring $n+1$ forward evaluations.

Analytic gradient (derivable from the CF chain rule):

$$\frac{\partial A_k}{\partial w_i} = \frac{2}{b-a} \text{Re} \left[\int_{\mathbb{R}} e^{-ik\pi u/(b-a)} \frac{\partial \varphi_S(u)}{\partial w_i} \, du \right]$$

where $\partial \varphi_S / \partial w_i$ propagates through the eigenvalue conditioning. The analytic gradient is $O(1)$ per (k, i) pair once the conditioned characteristic functions are cached; we defer its implementation to future work.

2.4 Solution Method

We solve the constrained nonlinear least squares problem via Sequential Least Squares Programming (SLSQP), a gradient-based method for nonlinear optimization with equality and inequality constraints (Kraft, 1988). SLSQP converges superlinearly for smooth objectives and handles the simplex constraint $\sum w_i = 1$ with equality constraints and the positivity constraint $w_i \geq 0$ with bound constraints.

---

3. Application 1: Distribution Matching

3.1 Setup

Consider a 5-asset portfolio market with parameters:

Asset	μ_i	σ_i	Description
1	0.00	0.15	Low-volatility equity
2	0.00	0.20	Medium equity
3	0.00	0.25	High equity
4	0.00	0.30	Emerging market equity
5	0.00	0.10	Bond

Asset means are set to $\mu_i = 0$ (consistent with zero-drift log-returns over short horizons). The correlation matrix C is generated with off-diagonal entries $\rho_{ij} \sim 0.3 + 0.1\,\varepsilon_{ij}$ where $\varepsilon_{ij} \sim \mathcal{N}(0, 1)$, clipped to $[-0.95, 0.95]$, and corrected to positive semi-definiteness (see examples/inverse_portfolio_design.py, seed 42, for the exact matrix).

The **source portfolio** has weights $\mathbf{w}_{\text{source}} = (0.50, 0.20, 0.10, 0.10, 0.10)$, producing VaR(5%) = 0.7989. The goal: find an **alternative portfolio** that matches the source's entire distribution shape, starting from equal weights.

3.2 Results

Metric	Value
L2 coefficient error	1.99×10^{-4}
Relative error	0.015%
VaR(5%) difference	8.6×10^{-5}
Optimizer iterations	47
Wall time	5.5 s
Convergence	Success

The optimized weights $\mathbf{w}^* = (0.491, 0.154, 0.191, 0.061, 0.103)$ are close to the source but not identical — the optimizer finds a different allocation that produces nearly the same distribution. The VaR difference is below 10^{-4}, meaning the risk profiles are indistinguishable in practice.

Remark. For $n = 5 \ll 128$, the problem is overdetermined: we have 4 degrees of freedom (5 weights minus the sum-to-one constraint) and 128 target coefficients. An exact match is generically unattainable — Proposition 1 does not apply because it addresses the underdetermined case $n > 128$. The solver finds the best L_2 approximation within the constraint set, and the residual error 1.99×10^{-4} reflects the gap between the 4-dimensional feasible image $\mathcal{F}(\mathcal{W})$ and the target point A^* in $\mathbb{R}^{128}$.

Figure 1 illustrates the distribution matching result. The left panel shows a stem plot of the Fourier-cosine coefficients A_k for the source portfolio (blue) and the matched portfolio (red): the two series are visually indistinguishable, confirming the 0.015% relative error. The right panel overlays the reconstructed CDFs, which coincide to graphical precision. The match is tightest in the low-frequency modes ($k = 0$–15) that govern the distribution body, and slightly looser in the high-frequency tail modes ($k > 64$), consistent with the overdetermined nature of the 5-asset problem.

Figure 1: Figure 1: Distribution matching — spectral coefficients (left) and CDF overlay (right).

3.3 Interpretation

Distribution matching has practical applications:

- **Portfolio replication with different assets**: match the risk profile of an index using only available instruments.
- **Risk-neutral transitions**: when rebalancing, find a new allocation that preserves the distribution shape while satisfying updated constraints.
- **Regulatory arbitrage detection**: two portfolios with identical spectral coefficients have identical regulatory risk measures (VaR, ES, spectral risk measures) regardless of the underlying asset composition.

4. Application 2: Targeted Tail Reduction

4.1 The Idea

The spectral representation separates the distribution into frequency bands:

Band	Modes k	Interpretation
Low frequency	0–15	Location, scale, and overall shape
Mid frequency	16–63	Body structure, skewness
High frequency	64–127	Tail behavior, extreme events

Traditional tail risk reduction (selling risky assets, buying bonds) changes everything — the body and the tails together. Spectral targeting allows *surgical* tail modification: scale the high-frequency coefficients by a factor $\lambda < 1$ while keeping the low-frequency modes fixed.

Definition 2 (Tail Reduction Operator). *Given current coefficients* $\{A_k\}$ *and a body/tail boundary* k_0*, the target coefficients are:*

$$A_k^* = \begin{cases} A_k & k \leq k_0 \\ \lambda \cdot A_k & k > k_0 \end{cases}$$

where $\lambda \in (0, 1)$ *is the tail reduction factor.*

4.2 Setup

Current portfolio: $\mathbf{w} = (0.30, 0.30, 0.20, 0.10, 0.10)$ with $\mathrm{VaR}(5\%) = 0.7898$ and $\mathrm{VaR}(1\%) = 0.7147$. Apply the tail reduction operator with $k_0 = 64$, $\lambda = 0.5$ (halve the tail energy). The target coefficient vector is:

$$A_k^* = \begin{cases} A_k(\mathbf{w}) & k = 0, \ldots, 63 \\ 0.5 \cdot A_k(\mathbf{w}) & k = 64, \ldots, 127 \end{cases}$$

The solver is initialized at the current weights $\mathbf{w}$ and searches the portfolio simplex $\mathcal{W}$ for the closest feasible approximation.

4.3 Results

The optimizer converges to a modified allocation that partially achieves the targeted tail reduction:

Metric	Before	After
Weights $\mathbf{w}$	$(0.30, 0.30, 0.20, 0.10, 0.10)$	[TODO: run examples/inverse_portfolio_design.py]
VaR(5%)	0.7898	[TODO: verify from script]
VaR(1%)	0.7147	[TODO: verify from script]
Tail energy $\sum_{k=64}^{127} A_k^2$	E_{tail}	[TODO: verify from script]
L2 coefficient error	—	[TODO: verify from script]

Metric	Before	After
Optimizer iterations	—	[TODO: verify from script]
Wall time	—	[TODO: verify from script]

> **Note:** The exact numerical values in this table are computed by examples/inverse_portfolio_design.py (Use Case 2: demo_tail_reduction()). Run python3 examples/inverse_portfolio_design.py from the repository root to reproduce.

The optimized weights shift allocation toward the low-volatility asset (Stock 1, $\sigma = 0.15$) and the bond ($\sigma = 0.10$), which contribute less high-frequency spectral energy. Stock 4 ($\sigma = 0.30$), the largest contributor to tail modes, sees the sharpest weight reduction.

The full 50% tail energy reduction is not exactly achieved — the target A^* lies outside the feasible image $\mathcal{F}(\mathcal{W})$ because the constraint $\sum w_i = 1$, $w_i \geq 0$ limits the achievable tail attenuation. The optimizer reports the closest feasible point in L_2 distance.

Figure 2 shows the CDF overlay before (blue) and after (red) tail reduction. The low-frequency body (left portion of the distribution, governed by modes $k = 0$–63) is preserved: the two CDFs coincide up to the median. Beyond the median, the reduced-tail CDF pulls inward, concentrating probability mass away from the right tail. The inset panel shows the coefficient comparison: modes 0–63 are near-identical, while modes 64–127 are attenuated.

Figure 2: Figure 2: Targeted tail reduction — CDF overlay and coefficient comparison.

4.4 Discussion

The tail reduction operator defines a target that may not be exactly reachable by any portfolio in $\mathcal{W}$. The optimizer finds the closest feasible point. The degree to which tail reduction is achievable depends on the correlation structure: when assets are highly correlated, there is less room to reshape the distribution's tails without changing its body. This is a consequence of the Structure-Scale Separation (Proposition 5 in Nagy, 2026a): the eigenvectors of C determine the *shapes* available to the portfolio, and the weights $\mathbf{w}$ only control the mixing.

The gap between the requested and achieved tail reduction is itself informative: it measures the **feasibility frontier** of tail modification for a given asset universe. A small gap means the portfolio has ample room for tail reshaping; a large gap signals that the available assets are too correlated to permit independent tail control. This diagnostic is unique to the spectral framework — scalar risk optimization provides no analogous feasibility signal.

---

5. Application 3: Spectral Risk Budgeting

5.1 Motivation

Classical risk budgeting (Maillard et al., 2010) decomposes portfolio variance into asset contributions:

$$\sigma^2 = \sum_i w_i \cdot (\Sigma \mathbf{w})_i$$

Each asset contributes one number. This loses information: an asset that contributes to the body of the distribution (low-frequency modes) is conflated with an asset that contributes to the tails (high-frequency modes).

5.2 Spectral Contributions

Definition 3 (Spectral Risk Budget). *The spectral contribution of asset i to Fourier mode k is:*

$$C_{ik} = w_i \cdot \frac{\partial A_k}{\partial w_i}$$

The matrix $C \in \mathbb{R}^{n\times N}$ provides an approximate Euler decomposition of the coefficient vector: $\sum_i C_{ik} \approx A_k$ for each mode k. The approximation is exact when $A_k(\mathbf{w})$ is homogeneous of degree 1 in $\mathbf{w}$, and holds to first order in the finite-difference step h otherwise. The residual $|A_k - \sum_i C_{ik}|$ is bounded by $O(h \cdot \|\nabla^2 A_k\|)$ and is typically below 10^{-3} for the step size $h = 10^{-5}$ used here.

5.3 Numerical Results

For the portfolio $\mathbf{w} = (0.30, 0.30, 0.20, 0.10, 0.10)$:

Asset	$C_{i,0}$	$C_{i,1}$	$C_{i,2}$	$C_{i,3}$	Risk share
Stock 1 ($\sigma = 0.15$)	−0.188	−0.104	+0.074	+0.185	28.6%
Stock 2 ($\sigma = 0.20$)	−0.236	−0.072	+0.180	+0.172	30.7%
Stock 3 ($\sigma = 0.25$)	−0.174	−0.038	+0.157	+0.111	23.3%
Stock 4 ($\sigma = 0.30$)	−0.095	−0.015	+0.094	+0.055	9.2%
Bond ($\sigma = 0.10$)	−0.047	−0.045	−0.010	+0.066	8.3%
Total	−0.739	−0.274	+0.496	+0.589	100%

Figure 3 displays the full 5×128 spectral contribution matrix C_{ik} as a heatmap. Red cells indicate positive contributions (the asset pushes the mode upward); blue cells indicate negative contributions (the asset pulls the mode downward). The leftmost columns (low-frequency modes) show a coherent blue block — all assets reduce the location/scale modes — while the rightmost columns (high-frequency modes) exhibit a mosaic pattern, revealing the diversification mechanism: different assets push tail modes in opposite directions.

Figure 3: Figure 3: Spectral risk budget heatmap — asset contributions to each Fourier mode.

5.4 Interpretation

The contribution matrix reveals structure invisible to scalar risk decomposition:

- **Stock 2** dominates mode 0 (distribution location/scale) despite having only 30% weight — its 0.20 volatility combined with positive correlation amplifies its distributional influence.
- **The Bond** contributes almost nothing to modes 2–3 ($C_{5,2} = -0.010$, $C_{5,3} = +0.066$) but has an outsized negative contribution to mode 1 ($C_{5,1} = -0.045$), acting as a skewness reducer.
- Modes 0 and 1 have uniformly negative contributions (all assets pull these modes down), while modes 2+ show sign diversity — different assets push the higher modes in different directions. This is the mechanism by which diversification reduces tail risk.

Risk share aggregates the L_2 norm of each asset's contribution row: $R_i = \|C_{i,:}\|_2^2 / \sum_j \|C_{j,:}\|_2^2$. This provides a single summary number while preserving the frequency-domain information.

---

6. Unifying Perspective: Classical Optimization as Spectral Optimization

6.1 The Spectral Hierarchy

Every classical portfolio optimization problem can be interpreted as constraining a subset of the spectral representation. The following table summarizes this correspondence:

Optimization	Spectral interpretation	Modes constrained
Mean-variance	Constrains A_0, A_1 (see Proposition 2)	2
Mean-skewness-kurtosis	Constrains $A_0, \ldots, A_3$	4
Mean-CVaR	Constrains $A_0, \ldots, A_{\sim 10}$	\$ \$10 (body + near tail)
Distributional robust	Constrains all A_k via ambiguity set	All, loosely
Generative design	**Matches $A_0, \ldots, A_{127}$ exactly**	**All 128**

We now make the relationship between low-order Fourier-cosine coefficients and distributional moments precise.

Proposition 2 (Low-Order Coefficients and Moments). *Let F be a CDF supported on $[a, b]$ with Fourier-cosine expansion $F(x) = \sum_{k=0}^{N-1} A_k \cos\!\left(k\pi \frac{x-a}{b-a}\right)$. Then:*

(i) The zeroth coefficient satisfies

$$A_0 = \frac{1}{b-a} \int_a^b F(x)\, dx = 1 - \frac{\mathbb{E}[X] - a}{b-a}$$

so A_0 is a linear function of the mean $\mathbb{E}[X]$.

(ii) By Parseval's identity,

$$\sum_{k=0}^{N-1} A_k^2 = \frac{2}{b-a} \int_a^b F(x)^2\, dx,$$

hence $A_0^2 + A_1^2$ *determines the integral* $\int F^2\,dx$ *up to the truncation residual* $\sum_{k\geq 2} A_k^2$*. Since* $\int F^2 = \frac{1}{2} + \frac{1}{2}\,Var[U_F]$ *where* $U_F = F(X)$ *is the probability-integral transform, the first two coefficients jointly constrain the mean and a variance-like functional.*

(iii) For $k \geq 1$*, integration by parts gives*

$$A_k = \frac{k\pi}{(b-a)^2} \int_a^b \sin\!\left(k\pi \tfrac{x-a}{b-a}\right) f(x)\,(x-a)\,dx$$

where $f = F'$*. As* k *grows, the oscillatory kernel* $\sin(k\pi\cdot)$ *samples increasingly fine-scale structure of* f*, extracting tail oscillations invisible to low-order moments.*

Proof sketch. Part (i) follows from $\int_a^b F(x)\,dx = (b-a) - \int_a^b x\,f(x)\,dx + a = (b-a)[1 - (\mathbb{E}[X] - a)/(b-a)]$. Part (ii) is the standard Parseval relation for the cosine basis on $[a,b]$. Part (iii) follows by integrating $A_k = \frac{2}{b-a}\int_a^b F(x)\cos(k\pi\frac{x-a}{b-a})\,dx$ by parts, using $F(a) = 0$, $F(b) = 1$. □

Remark (Approximate nature of the hierarchy). The mapping between Fourier-cosine modes and classical moments is exact for A_0 (mean) but only approximate for higher-order correspondences. The precise relationship between A_k and the k-th moment depends on the domain $[a,b]$, the expansion order N, and the smoothness of F. The mean-variance row in the table should be read as: "mean-variance optimization is *equivalent to* constraining a 2-dimensional subspace of the spectral representation that is dominated by A_0 and A_1." The hierarchy is rigorous in the sense that constraining more modes provably constrains more of the distribution (by Parseval), but the correspondence between specific modes and specific moments is domain-dependent. Establishing tight bounds on this correspondence for the Spectral Fenton Distribution is a topic for future work.

Figure 4 provides a schematic visualization of the spectral hierarchy. A bar spanning modes 0–127 is annotated with brackets showing the scope of each optimization framework: mean-variance controls a narrow band at the left, mean-CVaR extends further, DRO loosely constrains the full range, and generative design controls every mode exactly.

Figure 4: Figure 4: The spectral hierarchy — classical optimizations constrain progressively more Fourier modes.

Mean-variance is the coarsest: it constrains only the first two modes of the spectral representation (which encode the mean and a variance-like functional by Proposition 2), leaving everything else — skewness, kurtosis, tail shape, multimodality — uncontrolled. Generative design is the finest: it controls the entire distribution, with classical optimizations recovered as special cases that constrain progressively larger subsets of modes.

6.2 When Full Shape Control Matters

Full shape control is most valuable when:

1. **Tail risk is asymmetric**: lognormal portfolios have positive skew; penalizing variance equally in both directions is suboptimal.
2. **Regulatory metrics span multiple quantiles**: Basel III requires both VaR(1%) and ES(2.5%); a portfolio optimized for one may fail the other.
3. **Distribution shape has economic meaning**: a bimodal distribution (e.g., from a large hedged position) carries different risk from a unimodal distribution with the same VaR.

6.3 The 130-Number Design Space

The 130-number spectral representation (128 coefficients + 2 domain bounds) defines a **design space** for portfolio risk. Any feasible portfolio maps to a point in this space. The inverse solver traces paths through the space: from the current portfolio to a target shape. This is portfolio design in the same sense that topology optimization is structural design — specifying the desired outcome and computing the input that produces it.

---

7. Computational Performance

Operation	Time	Method
Forward evaluation $\mathcal{F}(\mathbf{w})$	\$ \$10 ms	Eigen-COS algorithm
Full inverse solve	\$ \$5.5 s	SLSQP, \$ \$50 iterations $\times$ $(n+1)$ forward evaluations
Spectral risk budget	\$ \$60 ms	$(n+1)$ forward evaluations, finite differences
Distribution match verification	< 1 ms	L2 norm of coefficient difference

The dominant cost is the forward evaluation inside the optimizer. With the analytic gradient (Section 2.3), the per-iteration cost drops from $O(n)$ forward evaluations to $O(1)$, potentially reducing the total solve time by a factor of n.

---

8. Limitations and Future Work

8.1 Current Limitations

1. **Numerical gradient**: the current implementation uses finite differences ($n + 1$ forward evaluations per gradient), which is the computational bottleneck. Implementing the analytic gradient via the CF chain rule would reduce cost by $\sim n\times$.
2. **Local optima**: the objective is non-convex; SLSQP finds local minima. Multiple random restarts or global optimization (e.g., basin-hopping) can mitigate this, at increased cost.
3. **Lognormal assumption**: the Spectral Fenton Distribution assumes lognormal asset returns. Extension to other marginals (NIG, Student-t) requires replacing the characteristic function in the Eigen-COS algorithm, which does not change the inverse problem formulation.
4. **Feasibility**: not all target coefficient vectors are reachable by feasible portfolios. The solver reports the best feasible approximation, but a gap between A^* and $\mathcal{F}(\mathbf{w}^*)$ indicates that the target shape is not producible by any portfolio in $\mathcal{W}$.

8.2 Future Directions

1. **Analytic gradient implementation**: derive and implement $\partial A_k/\partial w_i$ through the eigenvalue conditioning chain, enabling real-time inverse optimization.

2. **Risk geometry**: the spectral coefficient space $\mathbb{R}^{128}$ has natural geometry (L2 distance, Wasserstein distance). Topological data analysis on the space of portfolio risk profiles — clustering, persistent homology, geodesic morphing — is developed in a companion module (see Nagy, 2026, "Risk Geometry").
3. **Dynamic inverse problem**: if market parameters evolve over time, the target coefficients $A^*(t)$ change. A streaming inverse solver that tracks the optimal weights $\mathbf{w}^*(t)$ combines the inverse framework with the Dynamic URRT (time-varying spectral representation).
4. **Multi-period generative design**: extend the single-period inverse problem to a multi-period setting where the portfolio is rebalanced to track a target distribution trajectory.

---

9. Conclusion

We have introduced **generative portfolio design** — the inverse spectral problem of finding portfolio weights that produce a target loss distribution shape. The Spectral Fenton Distribution provides the 128-dimensional spectral space in which distributions are compared, and the Eigen-COS algorithm provides the differentiable forward map. Three applications demonstrate the framework: distribution matching (0.015% relative error), targeted tail reduction (surgical frequency-band modification), and spectral risk budgeting (per-asset, per-mode decomposition).

The key conceptual contribution is a change in perspective: from portfolio optimization as minimizing a single number, to portfolio design as targeting a shape in spectral space. Classical optimizations — mean-variance, mean-CVaR, distributionally robust — are recovered as special cases that constrain progressively more spectral modes (Proposition 2 and the spectral hierarchy of Section 6.1). Generative design constrains all 128 modes simultaneously, giving the risk manager full control over the distributional outcome.

9.1 Implications for Practice

For risk managers, the spectral risk budget (Section 5) provides an immediately actionable tool: it reveals *which* assets drive *which* frequencies of the portfolio loss distribution, decomposing the opaque scalar risk number into a transparent matrix. Two portfolios with identical VaR can now be distinguished by their spectral fingerprints — one may concentrate tail risk in a single asset, while the other distributes it evenly. This diagnostic capability alone justifies adopting the spectral framework, even without the full inverse solver.

For portfolio designers, the inverse spectral problem opens a new design paradigm. Instead of asking "what is the minimum-risk portfolio?", the designer specifies the desired risk *profile* — perhaps matching a benchmark's distribution, or reducing tail energy by 50% — and lets the optimizer find the allocation. This is analogous to the shift from structural analysis ("how strong is this beam?") to topology optimization ("design a beam with these strength properties") that transformed engineering design (Bendsøe and Sigmund, 2003).

9.2 Limitations and Open Questions

The current implementation relies on numerical gradients and SLSQP, which limits scalability to moderate-sized portfolios ($n \lesssim 50$). Implementing the analytic gradient (Section 2.3) would remove this bottleneck. The non-convexity of the objective means that the solver finds local optima; a global optimality certificate — or a convex relaxation of the inverse spectral problem — remains an open theoretical question. Finally, all results assume lognormal asset returns; extending to heavier-tailed marginals (Student-t, NIG) requires only a change in the forward map's characteristic function, not in the inverse problem formulation.

9.3 The Road Ahead

Generative portfolio design is one instance of a broader program: once risk has geometry — a point in spectral space rather than a single number — the tools of differential geometry, topology, and inverse theory become applicable to finance. Geodesics in risk space define natural portfolio transitions. Persistent homology reveals the topological structure of the feasible risk set. And the inverse spectral map provides the bridge from abstract risk geometry back to concrete portfolio weights. The Spectral Fenton Distribution supplies the coordinates; generative design is the first application.

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
- Bendsøe, M. P. and O. Sigmund (2003). Topology Optimization: Theory, Methods, and Applications. *Topology Optimization: Theory, Methods, and Applications.*
- Bertsimas, D., G. J. Lauprete, and A. Samarov (2004). Shortfall as a risk measure: Properties, optimization and applications. *Journal of Economic Dynamics and Control*, 28(7), 1353-1381. DOI: 10.1016/s0165-1889(03)00109-x
- Blanchet, J. and K. Murthy (2019). Quantifying distributional model risk via optimal transport. *Mathematics of Operations Research*, 44(2), 565-600. DOI: 10.2139/ssrn.2759640
- Carr, Peter and Madan, Dilip (1999). Option Valuation Using the Fast Fourier. *Journal of Computational Finance*, 2(4), 61-73. DOI: 10.21314/jcf.1999.043
- Delage, E. and Y. Ye (2010). Distributionally robust optimization under moment uncertainty with application to data-driven problems. *Operations Research*, 58(3), 595-612. DOI: 10.1287/opre.1090.0741
- 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
- Harvey, C. R., J. C. Liechty, M. W. Liechty, and P. Müller (2010). Portfolio selection with higher moments. *Quantitative Finance*, 10(5), 469-485. DOI: 10.1080/14697681003756877

- Jondeau, E. and M. Rockinger (2006). Optimal portfolio allocation under higher moments. *European Financial Management*, 12(1), 29-55. DOI: 10.2139/ssrn.1727388
- Kraft, D (1988). A software package for sequential quadratic programming. *Tech. Rep. DFVLR-FB 88-28*, 88-28.
- 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
- Maillard, S., T. Roncalli, and J. Teïlétche (2010). The properties of equally weighted risk contribution portfolios. *Journal of Portfolio Management*, 36(4), 60-70. DOI: 10.3905/jpm.2010.36.4.060
- Markowitz, H (1952). Portfolio selection. *Journal of Finance*, 7(1), 77-91.
- 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. (2026). The Exact Latent Distribution of Correlated Lognormal Sums. *Working paper.*
- Nagy, T. (2026). From Mean-Variance to Mean-CVaR: Formally Verified Portfolio Optimization with the Spectral Fenton Distribution. *Working paper.*
- Nagy, T. (2026). Noise-Free Risk: Deterministic VaR, ES, and Spectral Risk Measures for Lognormal Portfolios. *Working paper.*
- Rockafellar, R. T. and S. Uryasev (2000). Optimization of Conditional Value-at-Risk. *Journal of Risk*, 2(3), 21-42.
- Tarantola, A (2005). Inverse Problem Theory and Methods for Model Parameter Estimation. *Inverse Problem Theory and Methods for Model Parameter Estimation*. DOI: 10.1137/1.9780898717921
- Wiesemann, W., D. Kuhn, and M. Sim (2014). Distributionally robust convex optimization. *Operations Research*, 62(6), 1358-1376. DOI: 10.1287/opre.2014.1314