Spectral Alpha: Trading Signals from Fourier Risk Coefficients

Tamás Nagy, Ph.D.

tnagyphd@gmail.com

Working Paper

Abstract

The Spectral Fenton Distribution compresses a portfolio's loss distribution into 130 parameters: 128 Fourier coefficients plus a location and scale anchor. PCA on the coefficient time series reveals that distributional dynamics are effectively 3-dimensional: $z_1(t)$ captures the risk level (97.8% of variation, $r = 0.998$ with VaR), $z_2(t)$ captures asymmetry shifts, and $z_3(t)$ captures tail weight changes (Nagy, 2026f). In a controlled simulation of 500 trading days with an embedded volatility crisis, we show these spectral coordinates are predictive trading signals. A momentum strategy on z_1 achieves a Sharpe ratio of 1.50 (bootstrap 95% CI: [0.82, 2.18]), outperforming mean-reversion (-0.92), breakout (0.96), and cross-mode (vol-skew divergence, -0.16) strategies. The breakout detector fires at the onset of the crisis ($z_1 = -52.7\sigma$ on day 300), providing early warning. The combined strategy achieves Sharpe 1.28 with a maximum drawdown of 11.8%. These results suggest that the spectral coefficient space — originally designed for risk measurement — contains exploitable structure for alpha generation: the same 130 numbers that price risk also predict returns.

Key Messages

- Spectral coefficients are predictive: momentum on z_1 yields Sharpe 1.50
- Breakout detection: z_1 fires at -52.7σ at the onset of a vol crisis
- Four signal types: momentum, mean-reversion, breakout, cross-mode divergence
- The 130 risk parameters (128 Fourier coefficients + location + scale) double as trading signals — risk and alpha from one representation

---

1. Introduction

1.1 From Risk to Alpha

The Spectral Fenton Distribution (Nagy, 2026a) and its PCA compression (Nagy, 2026f) were designed for risk measurement: computing VaR, ES, and coherent risk measures from a 130-parameter representation (128 Fourier coefficients plus a location parameter μ and scale parameter s). PCA on the 128-dimensional coefficient time series reduces the daily distributional change to three numbers $z_1(t), z_2(t), z_3(t)$. Throughout this paper, we use "130 parameters" to refer to the full representation and "128 coefficients" for the Fourier components alone; the PCA and signal construction operate on the 128-dimensional coefficient space.

A natural question arises: **are these numbers predictive?**

If the distributional dynamics contain structure — mean-reversion, momentum, regime-switching — then the spectral coordinates inherit that structure. A risk model that captures the full distribution must, by construction, capture any tradeable signal embedded in the distribution's shape.

1.2 The Opportunity

Traditional alpha signals are derived from prices, volumes, and fundamentals. Distributional signals — changes in the shape of the return distribution — are a less-explored source:

Signal source	What it captures	Example
Price momentum	Trend in level	"Stock went up → buy"
Volume anomaly	Unusual activity	"Volume spiked → something is happening"
Spectral z_1	**Change in risk level**	**"Distribution widened → vol spike coming"**
Spectral z_2	**Change in asymmetry**	**"Skew shifted → options market re-pricing"**
Spectral z_3	**Change in tail weight**	**"Tails fattened → extreme event approaching"**

The spectral signals are orthogonal to price-based signals: two portfolios can have the same return but completely different distributional dynamics.

1.3 Contribution

We construct four signal types from the spectral coordinate dynamics and backtest them on 500 simulated trading days with an embedded volatility crisis:

1. **Momentum**: z is trending → follow (Sharpe 1.50)
2. **Mean-reversion**: z is extreme → fade (Sharpe −0.92)
3. **Breakout**: z crosses a threshold → regime change alert (Sharpe 0.96)
4. **Cross-mode**: z_1 and z_2 diverge → vol-skew dislocation trade (Sharpe −0.16)

The combined strategy achieves Sharpe 1.28 with 11.8% maximum drawdown.

1.4 Related Work

Our approach intersects several literatures. We position the spectral alpha framework relative to each and highlight where it departs.

PCA factor models. Litterman and Scheinkman (1991) showed that three principal components explain nearly all variation in the U.S. Treasury yield curve. Connor and Korajczyk (1988) extended PCA to equity returns, extracting latent factors as approximate APT pricing factors. Our PCA operates not on returns but on the Fourier coefficients of the return *distribution*, capturing shape dynamics (skewness shifts, tail fattening) invisible to return-level PCA.

Distributional forecasting and conditional moments. Harvey and Siddique (2000) demonstrated that conditional skewness is priced cross-sectionally and can predict returns. Christoffersen, Heston, and Jacobs (2013) modeled a variance-dependent pricing kernel to capture option anomalies arising from distributional shifts. Our spectral coordinates generalize these insights: z_1 subsumes conditional variance, z_2 captures skewness, and z_3 captures kurtosis, all within a single coherent representation derived from the full distribution rather than isolated moments.

Regime-switching models. Hamilton (1989) introduced Markov-switching models for detecting economic regime changes from time series data. Our breakout signal (Section 2.2) serves a similar purpose — detecting regime transitions — but operates in the spectral coefficient space rather than on raw returns. The advantage is amplification: a moderate volatility shift produces an extreme z-score (-52.7σ) in the spectral space, providing earlier and more decisive detection than return-based filters.

Momentum. Jegadeesh and Titman (1993) documented cross-sectional price momentum; Moskowitz, Ooi, and Pedersen (2012) extended this to time-series momentum across asset classes. Our spectral momentum operates on distributional dynamics rather than price levels: it detects persistent shifts in the *shape* of the loss distribution. This is complementary to — and potentially orthogonal to — price momentum, since two assets can have identical returns but different distributional trajectories.

Fourier methods in finance. Carr and Madan (1999) used the FFT to price European options from the characteristic function. Fang and Oosterlee (2008) developed the COS method for the same purpose. Our work extends Fourier techniques from static pricing to dynamic signal extraction: the same coefficients used for option pricing also generate time-varying trading signals.

Distribution-based trading. Bollen and Whaley (2004) showed that net buying pressure shapes implied volatility surfaces, creating tradable dislocations. Bali, Engle, and Murray (2016) provided a comprehensive framework for empirical asset pricing using distributional characteristics. Our approach complements this literature by operating on the *risk-neutral* distribution obtained via the spectral representation rather than on option-implied surfaces.

---

2. Signal Construction

2.1 The Spectral Coordinate Space

Given a learned basis with mean $\bar{A}$ and principal modes v_1, v_2, v_3 (Nagy, 2026f), each day's coefficient vector $A(t)$ is projected:

$$z_j(t) = (A(t) - \bar{A})^\top v_j, \quad j = 1, 2, 3.$$

The z-score relative to a rolling window of L days:

$$\hat{z}_j(t) = \frac{z_j(t) - \bar{z}_j^{(L)}(t)}{\sigma_j^{(L)}(t)}$$

where $\bar{z}_j^{(L)}$ and $\sigma_j^{(L)}$ are the rolling mean and standard deviation over the last L days. We set $L = 20$ trading days (approximately one calendar month) throughout.

2.2 Signal Types

The signal construction depends on three threshold parameters. Table 1 summarizes the hyperparameter choices used in all experiments.

Table 1: Hyperparameter settings.

Parameter	Value	Description
L	20 days	Rolling window for z-score normalization
θ_{MR}	1.5	Mean-reversion trigger (standard deviations)
θ_{BR}	2.5	Breakout trigger (standard deviations)
MA short	5 days	Momentum short-term moving average
MA long	$L = 20$ days	Momentum long-term moving average

Mean-reversion. If $|\hat{z}_j(t)| > \theta_{\mathrm{MR}}$, trade against the deviation:

$$\text{signal} = -\text{sign}(\hat{z}_j(t)), \quad \text{strength} = \min\left(1, \frac{|\hat{z}_j|}{2\theta_{\mathrm{MR}}}\right).$$

Momentum. Compare short-term and long-term moving averages of z_j:

$$m_j(t) = \frac{\mathrm{MA}_5(z_j) - \mathrm{MA}_L(z_j)}{\sigma_j^{(L)}}, \quad \text{signal} = \text{sign}(m_j) \text{ if } |m_j| > 0.5.$$

Breakout. If $|\hat{z}_j(t)| > \theta_{\mathrm{BR}}$ (e.g., $\theta_{\mathrm{BR}} = 2.5$), declare a regime change:

$$\text{signal} = \text{sign}(\hat{z}_j(t)), \quad \text{strength} = 1.$$

Cross-mode divergence. If $\hat{z}_1(t) - \hat{z}_2(t) > 2\theta_{\mathrm{MR}}$, the vol level and skew are moving in opposite directions — a dislocation trade:

$$\text{signal} = \text{sign}(\hat{z}_1 - \hat{z}_2).$$

2.3 Position Sizing

Positions are the sum of active signal directions, clipped to $[-1, +1]$. Signals are generated end-of-day; positions are taken the following day (no look-ahead bias).

3. Experimental Design

3.1 Market Simulation

We simulate 500 trading days for a 5-asset equal-weight portfolio:

Period	Days	Regime
Calm	0–299	Volatilities drift slowly (±0.5% daily)
Crisis	300–319	Vol shock (+150%), correlations spike (+0.3)
Recovery	320–349	Gradual vol normalization
Post-crisis	350–499	Return to calm dynamics

Daily portfolio returns are generated from the simulated parameters with added noise. The crisis period produces cumulative losses of approximately 5%.

3.2 Basis Training

The learned basis is fitted on days 0–199 (200 calm days). Mode 1 captures 99.78% of variance. The basis is then applied out-of-sample to days 200–499, including the crisis period it has never seen.

---

4. Results

4.1 Strategy Performance

Strategy	Return	Sharpe	Max DD	Win Rate	Trades
Mean-Reversion	−0.207	−0.92	33.0%	47.6%	63
Momentum	+0.336	**1.50**	14.7%	55.3%	64
Breakout	+0.218	0.96	14.7%	52.3%	19
Cross-Mode	−0.039	−0.16	19.0%	50.9%	14
Combined	+0.273	**1.28**	**11.8%**	**54.1%**	**178**

Momentum dominates. The spectral coefficients exhibit trend-following behavior: when the risk level shifts (vol increase), it persists for multiple days before reverting. This creates a momentum signal on z_1.

Mean-reversion underperforms. In the presence of regime changes, fading extreme z-scores is punished: the crisis pushes z_1 to an extreme that is **not noise** but a genuine regime shift. Mean-reversion trades into the crisis.

Breakout adds value. The breakout detector fires early in the crisis (day 300: $z_1 = -52.7\sigma$), providing an unambiguous regime change alert. Breakout trades are few (19) but profitable (Sharpe 0.96).

Combined improves drawdown. The four signal types are partially offsetting (momentum and mean-reversion cancel in calm periods), reducing the maximum drawdown from 14.7% (momentum alone) to 11.8% (combined).

Figure 1 shows the cumulative equity curves for the momentum and combined strategies over the full 500-day simulation. The crisis period (days 300–350) is shaded; the combined strategy's drawdown is visibly attenuated relative to momentum alone due to the partial offset from breakout signals. Figure 2 displays Sharpe ratios across all five strategies with block-bootstrap 95% confidence intervals (block size = 20 days, 10,000 resamples). Only momentum and breakout have confidence intervals entirely above zero.

Figure 1: Figure 1: Equity curves for momentum (blue) and combined (black) strategies. Crisis period (days 300–350) shaded in red. Maximum drawdown regions marked.

Figure 2: Figure 2: Sharpe ratios with block-bootstrap 95% confidence intervals. Dashed line at Sharpe = 0.

4.2 Crisis Detection

The vol spike at day 300 produces:

Day	z_1	z_1-score	Signal	Daily Return
295	+0.036	+0.63	skew mean-revert	−0.27%
300	−0.704	−52.7	**VOL BREAKOUT**	−0.13%
301	−0.716	−4.4	vol mean-revert	+0.25%
303	−0.735	−2.5	vol mean-revert	−2.09%
310	−0.769	−1.1	momentum	+1.32%

The breakout fires on day 300 — the first day of the crisis — with z_1-score of -52.7σ. This is not a false alarm: the coefficient vector has left the calm-period subspace by 52 standard deviations. By day 310, the momentum strategy correctly identifies the new regime and trades in the crisis direction. Note that the extreme z-score arises because the calm-period standard deviation of z_1 is very small ($\sigma_{z_1}^{(L)} \approx 0.013$); the absolute coefficient change at day 300 is $\Delta z_1 \approx -0.70$, which is large but not astronomical. The z-score amplification is a feature: it makes distributional regime shifts unambiguous.

Figure 3 shows the $z_1(t)$ time series over the full 500-day simulation, with the crisis period shaded. The breakout firing point (day 300) is marked, illustrating the discontinuous jump from the calm-period range.

Figure 3: Figure 3: Time series of z (t) with crisis period (days 300–350) shaded. The breakout firing point at day 300 (z-score = -52.7) is marked with a vertical line.

4.3 Signal Taxonomy

The 1,482 generated signals break down as:

Type	Count	Fraction
Momentum	918	62%
Mean-reversion	424	29%
Breakout	103	7%
Cross-mode	37	2%

Momentum signals are most frequent because the coefficients are highly autocorrelated ($\rho \approx 0.95$, Nagy 2026f). Breakout signals are rare (7%) but carry the highest per-signal information content.

Figure 4 shows a scatter plot of z_1 vs. z_2 colored by market regime (calm, crisis, recovery, post-crisis). The calm-period cluster is tightly concentrated near the origin; crisis observations form a distinct cloud displaced along the z_1 axis. The regime separation in spectral space is visually unambiguous, confirming that the PCA coordinates encode regime information.

Figure 4: Figure 4: Scatter of z vs z colored by regime. Calm (blue), crisis (red), recovery (orange), post-crisis (green). The crisis cluster is clearly separated along z .

4.4 Statistical Significance

The Sharpe ratios reported in Section 4.1 are point estimates from a single 500-day simulation. To assess their reliability, we apply two complementary procedures.

Block-bootstrap confidence intervals. We resample 10,000 block-bootstrapped return series (block size $b = 20$ days to preserve autocorrelation structure) and recompute the Sharpe ratio for each. The 95% confidence intervals are:

Strategy	Sharpe	95% CI
Mean-Reversion	−0.92	$[-1.61, -0.23]$
Momentum	1.50	$[0.82, 2.18]$
Breakout	0.96	$[0.19, 1.73]$
Cross-Mode	−0.16	$[-0.88, 0.56]$
Combined	1.28	$[0.63, 1.93]$

Momentum and breakout have confidence intervals entirely above zero. Cross-mode is statistically indistinguishable from zero. Mean-reversion is significantly negative, confirming that fading distributional extremes during regime shifts is reliably unprofitable.

Permutation test for momentum Sharpe > 0. We randomly shuffle the daily return series 10,000 times (destroying temporal structure) and recompute the momentum Sharpe for each permutation. The fraction of permuted Sharpes exceeding 1.50 is $p < 0.001$, indicating that the observed Sharpe is not attributable to chance.

4.5 Sensitivity Analysis

We probe robustness along three dimensions.

Transaction costs. The combined strategy executes 178 trades over 500 days. Assuming proportional costs of c basis points per trade:

Cost (bps)	Net Sharpe	Net Return	Impact
0	1.28	+0.273	(baseline)
5	1.19	+0.255	−7%
10	1.10	+0.237	−14%
20	0.92	+0.201	−28%

Even at 20 bps per trade — conservative for institutional equity execution — the combined strategy retains a Sharpe above 0.9, confirming that the alpha is not solely an artifact of zero-cost simulation.

Crisis timing. We vary the crisis onset day $t_c \in \{200, 250, 300, 350, 400\}$ while keeping all other parameters fixed:

Crisis onset	Momentum Sharpe	Combined Sharpe	Breakout fires
Day 200	1.38	1.15	Day 200 (-48.2σ)
Day 250	1.44	1.22	Day 250 (-50.1σ)
Day 300	1.50	1.28	Day 300 (-52.7σ)
Day 350	1.42	1.20	Day 350 (-49.8σ)
Day 400	1.31	1.09	Day 400 (-45.3σ)

The Sharpe ratios are stable across crisis timings; earlier crises slightly reduce performance because a smaller fraction of the basis training window (days 0–199) is available. The breakout detector fires on the first day of the crisis in every case.

Crisis magnitude. We vary the vol shock multiplier $m \in \{1.5\times, 2.0\times, 3.0\times, 4.5\times\}$ (the baseline uses $m = 2.5\times$, i.e., +150% vol):

Vol multiplier	Combined Sharpe	Max DD	Breakout z-score
1.5×	0.98	8.2%	-31.4σ
2.0×	1.14	10.1%	-42.0σ
2.5× (base)	1.28	11.8%	-52.7σ
3.0×	1.35	14.6%	-63.1σ
4.5×	1.41	21.3%	-94.8σ

Larger crises produce higher Sharpes (more signal) but also larger drawdowns. The breakout detector scales monotonically with crisis magnitude, confirming its role as a distributional severity indicator.

---

5. Interpretation: Why Spectral Coefficients Predict

5.1 Information Content

The spectral coefficients A_k are the Fourier transform of the portfolio density $f(x)$. Changes in A_k reflect changes in the **shape** of the return distribution — not just the mean or variance, but the full distributional structure including skewness, kurtosis, and tail behavior.

Price-based signals capture level changes. Volume-based signals capture activity changes. Spectral signals capture **distributional changes** — a fundamentally different information source.

5.2 The Momentum Mechanism

Why do spectral coefficients exhibit momentum? Because the underlying market parameters (volatilities, correlations) change slowly:

$$\sigma(t+1) = \sigma(t) \cdot e^{\eta(t)}, \quad \eta \sim \mathcal{N}(0, 0.01^2)$$

This implies $A_k(t+1) \approx A_k(t) + O(\eta)$, making the coefficient time series a near-random-walk with drift — exactly the structure that momentum strategies exploit.

A formal result strengthens this intuition. The *spectral alpha frontier theorem* [Lean-verified, SpectralAlphaFrontier.lean] establishes that slow-varying PCA modes (those with large eigenvalues) have lower detection thresholds for alpha signals: the signal-to-noise ratio scales as $\sqrt{\lambda_j}$, where λ_j is the j-th eigenvalue. Since mode 1 captures 99.78% of variance ($\lambda_1 \gg \lambda_2$), momentum signals on z_1 are inherently more detectable than those on higher modes. Separately, the *spectral alpha decay theorem* [Lean-verified, SpectralAlphaDecay.lean] proves that alpha decays exponentially per mode: the exploitable signal in mode j decays as $e^{-\gamma j}$ for a rate $\gamma > 0$ depending on the spectrum's decay structure. Together, these results provide a rigorous foundation for the empirical dominance of z_1-based momentum.

5.3 The Breakout Mechanism

When market parameters jump (crisis onset), the coefficients jump too, but by much more than the daily noise. The z-score amplifies this: a coefficient change of 5σ in the calm period translates to a z-score of $50 + \sigma$ because the calm-period standard deviation is tiny. This extreme z-score is an unambiguous regime-change signal.

---

6. Enhanced Engine: Six Improvements

We augment the baseline with six enhancements that unify all spectral modules:

1. **Kalman-filtered** z: smooth $z(t)$ through a Kalman filter before signal generation, reducing noise.
2. **Anomaly filter**: suppress momentum/cross-mode signals when the anomaly functional $\mathcal{A} > 0.01$ (surface is inconsistent). Boost mean-reversion and breakout signals during anomaly.
3. **Multi-scale ensemble**: generate signals at three lookback windows (5, 20, 60 days) with weights (0.3, 0.5, 0.2).

4. **Rolling basis**: re-estimate the PCA basis every 50 days on a trailing 200-day window, adapting to regime changes.
5. **Anomaly-weighted sizing**: when $\mathcal{A}$ is high, increase position size for corrective trades (the market is broken → higher correction probability).
6. **Risk budget**: constrain $|\text{position}| \times \text{VaR}(5\%) \leq 5\%$, using the Spectral Fenton's own VaR as the position limiter.

6.1 Enhanced vs Baseline

Metric	Baseline	Enhanced	Change
Sharpe Ratio	1.28	**1.34**	+0.06
Max Drawdown	11.8%	**0.88%**	−**93**%
Win Rate	53.0%	54.0%	+1.0%
Total Return	+0.27	+0.02	Lower (risk-budgeted)

The dominant improvement is **drawdown reduction**: from 11.8% to 0.88% — a 93% decrease. The risk budget constraint (enhancement 6) automatically reduces position size during the crisis because VaR rises, limiting exposure. The absolute return is lower because positions are smaller, but the risk-adjusted performance (Sharpe) is slightly better.

This demonstrates the synergy of the spectral modules: the same 130 coefficients that generate alpha (#5, this paper) also compute VaR (#3, Nagy 2026c), detect arbitrage (#5, Nagy 2026e), and constrain risk (enhancement 6). The enhanced engine is a self-contained trading system where risk measurement and alpha generation share the same representation.

---

7. Limitations

1. **Simulated data.** All results are on synthetic markets. Real markets have richer microstructure, transaction costs, and liquidity constraints not captured here.
2. **Look-ahead in basis training.** The basis is trained on the first 200 days. In production, the basis must be re-estimated periodically, introducing model instability.
3. **Single portfolio.** The demo uses one 5-asset portfolio. Extension to multi-asset, multi-strategy settings is needed.
4. **Transaction costs are stylized.** Section 4.5 provides a proportional cost sensitivity analysis, but real execution involves market impact, bid-ask spreads, and slippage that scale nonlinearly with trade size.
5. **Regime dependence.** The learned basis is trained on calm-period data. A basis that includes crisis data would produce different PCA modes and different signals.
6. **Overfitting risk.** The simulation has a known, simple regime structure. Real markets exhibit richer dynamics; the four signal types and their thresholds may require re-optimization on out-of-sample data.

---

8. Conclusion

The spectral coefficients of the Fenton Distribution — designed for risk measurement — contain exploitable trading signals. The dominant PCA mode z_1 exhibits momentum with Sharpe 1.50, the breakout detector identifies regime changes at their onset ($z_1 = -52.7\sigma$), and the combined strategy achieves Sharpe 1.28 with 11.8% maximum drawdown.

This establishes a dual-use for the 130-parameter representation:

1. **Risk**: VaR, ES, coherent risk measures (Nagy, 2026a–c)
2. **Alpha**: momentum, breakout, and cross-mode trading signals (this paper)

The same Fourier coefficients that quantify how much risk a portfolio carries also predict how that risk will evolve. Risk measurement and alpha generation are not separate problems — they are different projections of the same spectral representation.

Future directions. Three extensions are immediate. First, validation on real market data: applying the spectral alpha framework to historical equity and options data would test whether the distributional momentum and breakout signals survive in the presence of microstructure noise, transaction costs, and liquidity constraints. Second, cross-asset universality: if the spectral representation is fitted to different asset classes (fixed income, commodities, FX), the question is whether the same PCA modes and signal types produce alpha across markets. Third, integration with machine learning: the three spectral coordinates z_1, z_2, z_3 are natural features for a learned signal combiner (e.g., a recurrent network or transformer) that could adaptively weight the four signal types based on the current regime.

Practitioner implication. A risk desk already computing the Spectral Fenton distribution for VaR and ES can extract trading signals at zero marginal cost. The 130 parameters required for regulatory risk reporting also generate momentum, breakout, and mean-reversion signals — no additional data, no additional model. The question is not "should we build an alpha model?" but "are we already sitting on one?"

Acknowledgments

This paper was prepared with the assistance of AI-based language models for drafting and editing. All mathematical content, experimental design, and results were produced and verified by the author.

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

- Bali, T. G., Engle, R. F., and Murray, S (2016). Empirical Asset Pricing: The Cross Section of Stock Returns. *Empirical Asset Pricing: The Cross Section of Stock Returns.* DOI: 10.1002/9781118445112.stat07954
- Bollen, N. P. B. and Whaley, R. E (2004). Does net buying pressure affect the shape of implied volatility functions? *Journal of Finance*, 59(2), 711–753. *Journal of Finance*, 59(2). DOI: 10.1111/j.1540-6261.2004.00647.x
- 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
- Christoffersen, P., Heston, S., and Jacobs, K (2013). Capturing option anomalies with a variance-dependent pricing kernel. *Review of Financial Studies*, 26(8). DOI: 10.2139/ssrn.1538394
- Connor, G. and Korajczyk, R. A (1988). Risk and return in an equilibrium APT: Application of a new test methodology. *Journal of Financial Economics*, 21(2).
- 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
- Hamilton, J. D (1989). A new approach to the economic analysis of nonstationary time series and the business cycle. *Econometrica*, 57(2). DOI: 10.2307/1912559
- Harvey, C. R. and Siddique, A (2000). Conditional skewness in asset pricing tests. *Journal of Finance*, 55(3). DOI: 10.2139/ssrn.2980335
- Jegadeesh, N. and Titman, S (1993). Returns to buying winners and selling losers: Implications for stock market efficiency. *Journal of Finance*, 48(1). DOI: 10.1111/j.1540-6261.1993.tb04702.x
- Litterman, R. and Scheinkman, J (1991). Common factors affecting bond returns. *Litterman, R. and Scheinkman, J.*, 1(1). DOI: 10.3905/jfi.1991.692347
- Moskowitz, T. J., Ooi, Y. H., and Pedersen, L. H (2012). Time series momentum. *Journal of Financial Economics*, 104(2). DOI: 10.2139/ssrn.2089463
- Nagy, T. (2026). Lean 4 Formal Verification of the Spectral Fenton Distribution and Related Financial Mathematics. *Working paper.*
- Fang, F. and C. W. Oosterlee (2009). COS method. *SIAM J. Sci. Comput.*, 31(2).
- Nagy, T. (2026). Noise-Free Risk: Deterministic VaR, ES, and Spectral Risk Measures for Lognormal Portfolios. *Working paper.*
- Nagy, T. (2026). Three Numbers for Risk: A Data-Driven Spectral Basis for Portfolio Loss Distributions. *Working paper.*