Frequency-Domain Theory of Financial Economics: Thirteen Fundamental Results from One Decomposition

Tamás Nagy, Ph.D.

tnagyphd@gmail.com

Working Paper

Abstract

A market is a spectrum. We show that thirteen fundamental results in financial economics — seven core classical theorems and six domain extensions — all follow from a single structural assumption: the return density of any portfolio can be decomposed into N independent Fourier modes $A_k(t)$, each evolving as an Ornstein–Uhlenbeck process with mode-dependent autocorrelation time $\tau_k \sim k^{-\alpha}$.

The **seven core results** are: (1) the risk-return trade-off, with efficient frontier $\sigma^2_{\min}(\mu) = \mu^2/\sum SR_k^2$; (2) the Capital Market Line, with slope $\sqrt{\sum SR_k^2}$; (3) the Sharpe ratio maximum, via the Pythagorean theorem in risk space; (4) the Kelly criterion, with optimal log-growth $g^* = \frac{1}{2}\sum SR_k^2$; (5) the diversification principle, where independent modes add Sharpe in quadrature $SR_{\text{total}} = \sqrt{N} \cdot \overline{SR}$; (6) the Efficient Market Hypothesis, which holds in aggregate ($\sum w_k SR_k = 0$ at equilibrium) while individual modes carry temporary alpha; and (7) the unification of momentum with mean reversion as the low- and high-frequency limits of the same dynamics.

The **six extensions** apply the same framework to: the CAPM and spectral beta, Fama–French factors as spectral modes, the volatility smile as the spectral state of the risk-neutral density, the yield curve as three Fourier modes, bubbles and crashes as mode displacement, and behavioral finance as asymmetric mean reversion.

All thirteen results are formally verified in Lean 4 with zero sorry. The proofs are elementary — linear algebra and calculus of the OU process — and all compile under lake build with Lean v4.28.0 and Mathlib.

---

1. Introduction

1.1 Thirteen Results, One Framework

Financial economics rests on a handful of foundational results:

1. **The risk-return trade-off**: higher expected return requires higher risk
2. **The Capital Market Line**: the optimal risk-return relationship is linear
3. **The Sharpe ratio maximum**: the tangent portfolio maximizes return per unit risk
4. **The Kelly criterion**: the optimal bet size maximizes log-growth
5. **Diversification**: combining assets reduces risk without proportionally reducing return
6. **The Efficient Market Hypothesis**: no strategy earns risk-adjusted excess return in aggregate

7. **Momentum and mean reversion**: two seemingly contradictory phenomena coexist

These seven core results were developed separately over decades: Markowitz (1952), Sharpe (1964), Kelly (1956), Fama (1970), Jegadeesh and Titman (1993). Each has its own assumptions, its own proof, its own literature.

We show that all seven — plus six extensions covering CAPM, factor models, the volatility smile, the yield curve, bubbles, and behavioral finance — follow from a single structural assumption: **the return density decomposes into independent Fourier modes with power-law autocorrelation times.** The proofs are elementary (linear algebra and calculus of the Ornstein–Uhlenbeck process), and all thirteen results are machine-verified in Lean 4.

The spectral framework connects to the Fourier option pricing literature pioneered by Carr and Madan (1999) and the COS method of Fang and Oosterlee (2008), which showed that Fourier-cosine series expansions provide exponentially convergent representations of option prices. Our contribution extends this spectral perspective from pricing to the full structure of financial economics: the same Fourier coefficients that price options also determine the risk-return trade-off, the Kelly criterion, and the efficient market hypothesis.

1.2 The Setup

The Spectral Fenton Distribution (Nagy, 2026a) represents the return density as $N = 128$ Fourier-cosine coefficients $A_0, \dots, A_{127}$. We extend this to a dynamic setting: each coefficient evolves as

$$dA_k(t) = -\frac{1}{\tau_k}(A_k(t) - \bar{A}_k)\,dt + \sigma_k\,dW_k(t)$$

where $\bar{A}_k$ is the long-run mean, τ_k is the mean-reversion time, σ_k is the volatility, and W_k are independent Brownian motions. The power law $\tau_k = C \cdot k^{-\alpha}$ determines the timescale of each mode.

This gives us three objects: - **The spectral state** $\mathbf{A}(t) = (A_0(t), \dots, A_{N-1}(t))$: the complete description of the market at time t - **The mode Sharpe ratio** $SR_k(h)$: the risk-adjusted return from trading mode k over horizon h - **The spectral Sharpe vector** $(SR_0, SR_1, \dots, SR_{N-1})$: a vector in $\mathbb{R}^N$ whose properties determine all seven results

1.3 Outline

Section 2 derives the foundational link from density coefficients to tradable mode returns. Sections 3–5 derive the seven core theorems. Section 6 presents the CAPM and APT as spectral projections. Section 7 extends the framework to the volatility smile, yield curves, bubbles, behavioral finance, and microstructure. Section 8 maps strategies to spectral bands. Section 9 presents numerical examples. Section 10 discusses formal verification. Section 11 provides a stochastic unification and discussion. Section 12 concludes.

---

2. From Density Coefficients to Mode Returns

The Spectral Fenton Distribution represents the return density $p(x)$ on an interval $[a, b]$ via the Fourier-cosine (COS) expansion:

$$p(x) = \frac{1}{b-a}\left[\frac{A_0}{2} + \sum_{k=1}^{N-1} A_k \cos\left(\frac{k\pi(x-a)}{b-a}\right)\right]$$

where $A_k = 2\int_a^b p(x)\cos(\frac{k\pi(x-a)}{b-a})\,dx$ are the COS coefficients. This expansion converges exponentially for analytic densities (Fang and Oosterlee, 2008), so $N = 128$ coefficients suffice for machine-precision reconstruction.

The coefficient A_k captures the k-th shape component of the density: A_0 is the normalization (related to the mean), A_1 encodes skewness-like asymmetry, A_2 encodes curvature (related to kurtosis), and higher k capture finer distributional features.

The key modeling step. We promote each coefficient from a static number to a stochastic process: $A_k \to A_k(t)$. This is justified by the empirical observation that return distributions change over time — skewness increases before crashes, kurtosis spikes during volatility regimes, and the mean drifts with business cycles. The question is: what dynamics does $A_k(t)$ follow?

We impose two conditions: 1. **Stationarity.** Each $A_k(t)$ has a long-run equilibrium value $\bar{A}_k$ to which it reverts. Financial densities do not drift without bound; extreme skew and kurtosis are temporary. 2. **Independence.** The modes $A_k(t)$ and $A_j(t)$ ($k \neq j$) are independent. This follows from the orthogonality of the Fourier basis: if the density innovations are translation-invariant, the cosine coefficients decorrelate.

The unique continuous-time process satisfying both conditions (Gaussian, stationary, Markov) is the Ornstein–Uhlenbeck process:

$$dA_k(t) = -\frac{1}{\tau_k}(A_k(t) - \bar{A}_k)\,dt + \sigma_k\,dW_k(t)$$

From mode dynamics to expected profit. Suppose at time t a trader observes $A_k(t) = a \neq \bar{A}_k$ — the density's k-th shape is displaced from equilibrium. The conditional expectation of $A_k(t+h)$ is:

$$\mathbb{E}[A_k(t+h) \mid A_k(t) = a] = \bar{A}_k + (a - \bar{A}_k)e^{-h/\tau_k}$$

The expected change in the coefficient over horizon h is therefore:

$$\mathbb{E}[\Delta A_k] = \bar{A}_k - a + (a - \bar{A}_k)e^{-h/\tau_k} = (\bar{A}_k - a)(1 - e^{-h/\tau_k})$$

This is the **expected profit** from trading mode k: a portfolio constructed to have unit exposure to A_k and zero exposure to all other modes earns $\mu_k(h) = (\bar{A}_k - a)(1 - e^{-h/\tau_k})$ in expectation. Such a portfolio can be constructed via options (butterfly spreads isolate curvature $\approx A_2$; risk reversals isolate skew $\approx A_1$) or via the dual COS transform applied to a set of liquid strikes.

The conditional variance of ΔA_k is:

$$v_k(h) = \frac{\sigma_k^2 \tau_k}{2}(1 - e^{-2h/\tau_k})$$

and the mode Sharpe ratio is $SR_k(h) = \mu_k(h)/\sqrt{v_k(h)}$. This completes the link: the COS coefficient A_k of the return density gives rise to a tradable signal with well-defined expected profit, risk, and Sharpe ratio. The seven core theorems (Sections 4–5) and six extensions (Sections 6–7) follow from the algebra of these quantities.

---

3. Mode Dynamics and the Sharpe Vector

3.1 The OU Mode

For a single mode k with parameters $(\bar{A}_k, \tau_k, \sigma_k)$, the expected profit over horizon h given current value $A_k(t) = a$ is:

$$\mu_k(h) = (\bar{A}_k - a)(1 - e^{-h/\tau_k})$$

The conditional variance is:

$$v_k(h) = \frac{\sigma_k^2 \tau_k}{2}(1 - e^{-2h/\tau_k})$$

The mode Sharpe ratio is:

$$SR_k(h) = \frac{\mu_k(h)}{\sqrt{v_k(h)}}$$

Theorem 1 (Equilibrium; Lean-verified). *At equilibrium ($A_k = \bar{A}_k$ for all k), every mode has $\mu_k = 0$ and $SR_k = 0$.*

```
theorem no_alpha_at_equilibrium (m : OUMode) (h : ) :
    expectedProfit m m.a_bar h = 0
```

Theorem 2 (Profit sign; Lean-verified). *If $A_k < \bar{A}_k$ (mode below equilibrium) and $h > 0$, then $\mu_k > 0$: the mode will rise toward equilibrium.*

```
theorem profit_sign (m : OUMode) (a h : ) (hh : 0 < h) :
    0 < (m.a_bar − a) → 0 < expectedProfit m a h
```

3.2 The Power Law

Conjecture 1 (Spectral Power Law). $\tau_k = C \cdot k^{-\alpha}$ *for a market-specific exponent $\alpha > 0$.*

Theorem 3 (Monotonicity; Lean-verified). *Higher modes revert faster: if $k_1 < k_2$, then $\tau_{k_2} < \tau_{k_1}$.*

```
theorem tauK_decreasing (mkt : SpectralMarket) (k k : )
    (hk1 : 0 < k) (hk2 : 0 < k) (h : k < k) :
    tauK mkt k < tauK mkt k
```

---

4. The Four Classical Theorems

4.1 Theorem 4: The Risk-Return Trade-off

For N independent modes with Sharpe ratios $SR_1, \ldots, SR_N$, the minimum-variance portfolio achieving excess return μ has variance:

$$\sigma^2_{\min}(\mu) = \frac{\mu^2}{\sum_{k=1}^{N} SR_k^2}$$

Theorem 4 (No Free Lunch; Lean-verified). *If $\mu > 0$ and $\sum SR_k^2 > 0$, then $\sigma^2_{\min} > 0$. Positive excess return requires positive risk.*

```
theorem no_free_lunch {N :  } (sr : Fin N →  ) (mu :  )
    (h_mu : 0 < mu) (h_sr : 0 < totalSharpeSquared sr) :
    0 < minVariance sr mu
```

Theorem 4' (Monotonicity; Lean-verified). *More return requires more risk: $\mu_1 < \mu_2 \Rightarrow \sigma^2(\mu_1) < \sigma^2(\mu_2)$.*

```
theorem more_return_more_risk {N :  } (sr : Fin N →  )
    (mu mu :  ) (h1 : 0 < mu) (h2 : 0 < mu)
    (h : mu < mu) (h_sr : 0 < totalSharpeSquared sr) :
    minVariance sr mu < minVariance sr mu
```

4.2 Theorem 5: The Capital Market Line

The maximum achievable excess return at risk level σ is:

$$R(\sigma) = \sqrt{\sum_{k=1}^{N} SR_k^2} \cdot \sigma$$

The CML slope is $\sqrt{\sum SR_k^2}$ — the spectral Sharpe maximum. This is the Pythagorean theorem in risk space: independent risks add in quadrature.

Theorem 5 (CML; Lean-verified). *The CML is increasing: more risk permits more return.*

```
theorem cml_increasing {N :  } (sr : Fin N →  )
    (   :  ) (h1 : 0    ) (_ : 0    )
    (h :   <   ) (h_sr : 0 < totalSharpeSquared sr) :
    cmlReturn sr   < cmlReturn sr
```

The additional hypotheses $0 \leq \sigma_1$ and $0 \leq \sigma_2$ reflect the economic constraint that risk (standard deviation) is non-negative. The CML is defined only in the non-negative risk half-plane.

4.3 Theorem 6: The Sharpe Maximum and Diversification

Theorem 6a (Pythagorean Sharpe; Lean-verified). *For independent modes:*

$$SR^2_{\text{total}} = \sum_{k=1}^{N} SR^2_k$$

```
theorem sharpe_maximum_pythagorean {N : } (sr : Fin N → ) :
    totalSharpeSquared sr =  i : Fin N, sr i ^ 2
```

Theorem 6b (Diversification $\sqrt{N}$; Lean-verified). *For N modes with equal Sharpe s:*

$$SR^2_{\text{total}} = N \cdot s^2 \quad \Rightarrow \quad SR_{\text{total}} = \sqrt{N} \cdot s$$

```
theorem diversification_sqrt_n {N : } (s : ) :
    totalSharpeSquared (fun (_ : Fin N) => s) = ↑N * s ^ 2
```

Theorem 6c (Spectral diversification; Lean-verified). *Any single mode's Sharpe^2 is bounded by the total:*

$$SR^2_k \leq \sum_{j=1}^{N} SR^2_j$$

```
theorem spectral_diversification {N : } (sr : Fin N → ) (i : Fin N) :
    sr i ^ 2  totalSharpeSquared sr
```

4.4 Theorem 7: The Kelly Criterion

Theorem 7 (Spectral Kelly; Lean-verified). *The maximum log-growth rate across N independent modes is:*

$$g^* = \sum_{k=1}^{N} \frac{SR^2_k}{2} = \frac{1}{2} \sum_{k=1}^{N} SR^2_k$$

```
theorem kelly_growth_eq_half_sharpe_sq {N : } (sr : Fin N → ) :
    totalKellyGrowth sr = totalSharpeSquared sr / 2
```

More modes strictly increase growth:

```
theorem more_modes_more_growth {N : }
    (sr : Fin N → ) (sr_new : ) (h : sr_new  0) :
    totalKellyGrowth sr < totalKellyGrowth sr + kellyGrowthPerMode sr_new
```

---

5. The Spectral EMH and Momentum–Mean Reversion Unification

5.1 Theorem 8: The Spectral Efficient Market Hypothesis

Theorem 8 (Spectral EMH; Lean-verified). *At equilibrium, the risk-weighted aggregate profit across all modes is zero:*

$$\sum_{k=1}^{N} w_k \cdot \mu_k(h) = 0 \quad \text{for all } h$$

where $w_k = \sigma_k^2 \tau_k$ *is the risk weight of mode* k*.*

```
theorem spectral_emh_equilibrium {N :  }
    (modes : Fin N → OUMode) (h :  ) :
    aggregateProfit modes (fun i => (modes i).a_bar) h = 0
```

Theorem 8' (Paradox resolved; Lean-verified). *Individual modes can have* $SR_k \neq 0$ *while the aggregate is zero:*

```
theorem emh_paradox_resolved {N :  }
    (modes : Fin N → OUMode) (h :  ) (i : Fin N)
    (a_i :  ) (h_below : a_i < (modes i).a_bar) (hh : 0 < h) :
    0 < expectedProfit (modes i) a_i h
    aggregateProfit modes (fun j => (modes j).a_bar) h = 0
```

5.2 Theorem 9: Momentum = Mean Reversion

Theorem 9 (Unification; Lean-verified). *For any mode with* $\tau_k = T$*, horizon* $h_1 < T$ *gives momentum (persistence) and horizon* $h_2 > T$ *gives mean reversion. Both have positive expected profit when the mode is displaced.*

```
theorem one_spectrum_all_strategies
    (m_slow m_fast : OUMode)
    (h_slow_tau : 100 < m_slow.tau)
    (h_fast_tau : m_fast.tau < 5)
    (a_slow a_fast :  )
    (h_below_slow : a_slow < m_slow.a_bar)
    (h_below_fast : a_fast < m_fast.a_bar)
    (h_horizon :  ) (hh : 0 < h_horizon)
    (h_mid : 10 < h_horizon   h_horizon < 50) :
    (h_horizon < m_slow.tau)
    (m_fast.tau < h_horizon)
    (0 < expectedProfit m_slow a_slow h_horizon)
    (0 < expectedProfit m_fast a_fast h_horizon)
```

The hypotheses make the unification precise: a slow mode ($\tau > 100$ days) and a fast mode ($\tau < 5$ days), both displaced below equilibrium ($a < \bar{A}$), are both profitable at the same intermediate horizon ($10 < h < 50$ days). The slow mode is in its momentum regime ($h < \tau_{\text{slow}}$) while the fast mode is in its mean-reversion regime ($h > \tau_{\text{fast}}$).

Implication. Momentum and mean reversion are not opposing forces. They are the low-frequency and high-frequency limits of the same spectral dynamics. A monthly trader sees modes $k < 12$ as momentum and $k > 12$ as mean-reverting. A daily trader sees everything as momentum.

---

6. CAPM and APT as Spectral Projections

6.1 Theorem 10: The Spectral CAPM

Define the **spectral beta** of asset i with mode loadings $b_{i,k}$:

$$\beta_i = \frac{\mathrm{Cov}(R_i, R_m)}{\mathrm{Var}(R_m)} = \frac{\sum_k b_{i,k}\sigma_k^2}{\sum_k \sigma_k^2}$$

This is a weighted inner product: beta is the projection of the asset's mode loading vector onto the market's mode loading vector.

Theorem 10a (Market beta; Lean-verified). *The market portfolio (unit loading on all modes) has $\beta = 1$.*

```
theorem market_beta_one {N :  } (sigma : Fin N →  )
    (h_var : 0 < marketVariance sigma) :
    spectralBeta (marketLoading N) sigma = 1
```

Theorem 10b (CAPM linearity; Lean-verified). *If mode risk premiums are proportional to mode variance ($\mu_k = \lambda\sigma_k^2$), then expected return is linear in beta:*

$$E[R_i] = \lambda \cdot \beta_i \cdot \mathrm{Var}(R_m)$$

This IS the CAPM: $E[R_i] - r_f = \beta_i(E[R_m] - r_f)$ with $\lambda = (E[R_m] - r_f)/\mathit{Var}(R_m)$.

```
theorem capm_beta_form {N :  }
    (asset : AssetExposure N) (sigma : Fin N →  )
    (lambda :  ) (h_var : marketVariance sigma   0) :
    assetReturn asset (fun k => lambda * sigma k ^ 2) =
    lambda * spectralBeta asset sigma * marketVariance sigma
```

6.2 Theorem 11: Fama–French Factors as Spectral Modes

The spectral framework provides a natural interpretation of the CAPM that connects to the modern factor literature (Cochrane, 2005). Traditional asset pricing derives the CAPM from equilibrium or no-arbitrage arguments; here it emerges as a geometric property of the spectral decomposition.

The APT (Ross, 1976) says: $E[R_i] = r_f + \sum_j \lambda_j b_{i,j}$. The "factors" are unspecified. In the spectral framework, the factors ARE the dominant modes:

Fama–French factor	Spectral mode	Timescale
Market (β)	$k = 0$ (level)	Very long (cycle)
Value (HML)	$k = 1$ (skew)	Long (years)
Profitability (RMW)	$k = 2$ (kurtosis)	Medium-long
Momentum (UMD)	$k \approx 5$ (bulk shape)	Medium (months)
Size (SMB)	Difference in $k = 0$ across caps	Long

The factors are ordered by mode number, which equals ordering by timescale (Lean-verified: factor_ordering). This is not ad-hoc — it follows from the frequency structure.

---

7. Five Extensions: Smile, Yield Curve, Crashes, Behavioral, Microstructure

The spectral framework extends naturally beyond the classical seven results. In each case below, a well-known empirical phenomenon maps onto the mode structure developed in Sections 3–5. These extensions are not independent theories; they are consequences of the same OU mode dynamics applied to different market observables.

7.1 Theorem 12: The Vol Smile IS the Spectral State

The implied volatility smile $\sigma_{\text{imp}}(K)$ is not a model artifact — it is the spectral state of the risk-neutral density. Mode A_1 controls the skew, mode A_2 controls the curvature, and mode A_0 sets the ATM level. Black-Scholes is the $A_k = 0$ truncation for $k \geq 1$.

The smile dynamics follow from the OU mode dynamics: the skew mean-reverts with timescale τ_1, the curvature with $\tau_2 < \tau_1$. This explains the empirical observation (Cont and da Fonseca, 2002) that skew is more persistent than curvature. The parametric models SVI (5 params) and SABR (4 params, Gatheral, 2006) work precisely because they approximate the first 3–5 spectral modes. The spectral perspective also connects to the eigenvalue cleaning literature: Laloux, Cizeau, Bouchaud, and Potters (1999) showed that random matrix theory can separate signal from noise in correlation matrices — in the spectral framework, this corresponds to truncating modes $k > K$ where τ_k falls below the noise floor.

Lean: VolSmile.lean — option_price_linear, bs_is_zero_modes, smile_skew_more_persistent.

7.2 Theorem 13: The Yield Curve as Three Spectral Modes

The Nelson-Siegel (1987) yield curve model uses three factors: level, slope, curvature. These are exactly the $k = 0, 1, 2$ Fourier modes of the forward rate density.

Why 3 factors explain $>$99% of yield curve variation (Litterman and Scheinkman, 1991): if $\tau_k = C/k^2$, then $\tau_3 = C/9$ — so fast that mode $k = 3$ is indistinguishable from noise in monthly data. The Diebold–Li (2006) dynamic Nelson-Siegel model IS three OU processes on spectral modes.

The term premium = risk premium of the level mode ($k = 0$): long bonds are risky because τ_0 is large (level persists for years).

Lean: YieldCurve.lean — three_factors_suffice, ns_is_three_modes.

7.3 Theorem 14: Bubbles and Crashes as Mode Displacement

A bubble = $|A_0 - \bar{A}_0| > c\sigma_0$ (level mode far from equilibrium). A crash = rapid synchronized reversion of multiple displaced modes back to $\bar{A}_k$.

Crisis severity = number of modes displaced beyond threshold (Lean: `crisisSeverity`). At equilibrium, severity = 0. More displaced modes ⇒ worse crisis (monotonicity proved).

Flash crashes last minutes (high-k modes, $\tau_k < 1$ day). Bear markets last years (low-k modes, $\tau_k > 200$ days). Recovery time $\propto \max \tau_k$ over displaced modes.

This connects to the historical crisis taxonomy of Kindleberger (1978): manias correspond to persistent low-k displacement, while panics involve multi-mode cascade. The spectral framework quantifies what Kindleberger described qualitatively.

Lean: BubblesCrashes.lean — no_crisis_at_equilibrium, more_displacement_worse_crisis, flash_vs_bear.

7.4 Theorem 15: Behavioral Finance as Asymmetric OU

Behavioral biases arise from **asymmetric mean-reversion**: $\tau_k^+ \neq \tau_k^-$, where τ_k^+ is the reversion speed from overvaluation and τ_k^- from undervaluation.

Loss aversion (Kahneman and Tversky): $\tau_k^+ < \tau_k^-$. People sell winners quickly (fast reversion up) but hold losers (slow reversion down). This single inequality explains the disposition effect, overreaction to good news, and underreaction to bad news.

Herding = mode synchronization across assets: normally independent modes become correlated during panics. When herding coefficient $H_k \to 1$, the $\sqrt{N}$ diversification benefit collapses to 1.

Lean: Behavioral.lean — loss_aversion_asymmetry, isLossAverse, no_herding_independent.

7.5 Theorem 16: Market Microstructure as High-k Modes

Modes $k > 50$ capture microstructure: bid-ask bounce, order flow, inventory effects. Market making profit = exploitation of high-k mean reversion (very fast τ_k). Kyle's (1985) lambda (price impact) has a natural spectral form: $\lambda_k \propto \sigma_k^2/(\sigma_k^2 + \sigma_{\text{noise}}^2) \cdot 1/\tau_k$ — higher for faster modes, explaining why microstructure trades have high transient impact.

Permanent vs transient impact: large orders move low-k modes (slow τ_k ⇒ permanent impact). Small orders move high-k modes (fast τ_k ⇒ transient bounce-back).

Lean: Microstructure.lean — kyle_lambda_pos, fast_modes_higher_impact, market_making_is_high_k_mr.

---

8. The Strategy Spectrum

Every known trading strategy maps to a spectral band $[k_{\text{lo}}, k_{\text{hi}}]$:

Strategy	Modes k	Horizon	Mechanism	Lean theorem
Value investing	0–1	Years	Mean reversion of level	band_profit_zero_at_eq
Carry	0–2	Months	Yield = shift in A_0	spectrum_ordering
Momentum	2–10	Weeks–months	Persistence of bulk modes	momentum_mr_same_mode
Stat arb	5–30	Days	Mean reversion of relative modes	seven_strategies
Vol carry	1–3	Weeks	Kurtosis mode persistence	
Gamma scalping	2–5	Days	Realized vs implied shape	
HFT market making	30–128	Seconds–minutes	Microstructure reversion	hft_highest

This is the financial analogue of Maxwell's electromagnetic spectrum: different strategies "listen" to different frequencies of the same market signal.

---

9. Numerical Examples

9.1 Three-Mode Market

Consider a 3-mode market with $\alpha = 1$, representing a stylized equity index:

Mode	$\bar{A}_k$	σ_k	τ_k (days)	SR_k	SR_k^2	Interpretation
$k = 1$	0.52	0.08	252	0.45	0.2025	Bull/bear (1 year)
$k = 2$	−0.15	0.12	126	0.30	0.0900	Skew regime (6 months)
$k = 3$	0.03	0.06	84	0.15	0.0225	Vol clustering (3 months)

Spectral Sharpe maximum: $SR_{\text{total}} = \sqrt{0.2025 + 0.0900 + 0.0225} = \sqrt{0.315} \approx 0.561$

Kelly growth: $g^* = 0.315/2 = 0.1575$ (15.75% log-growth per year)

CML: $R(\sigma) = 0.561 \cdot \sigma$. At $\sigma = 20\%$: $R = 11.2\%$ excess return.

Efficient frontier: To earn $\mu = 10\%$ excess, minimum risk is $\sigma = 10\%/0.561 = 17.8\%$.

Diversification: Adding the 3rd mode increases SR from $\sqrt{0.2025 + 0.0900} = \sqrt{0.2925} \approx 0.541$ to 0.561 — a 3.7% improvement from a single low-Sharpe mode. This illustrates Theorem 6b: every additional independent mode improves the portfolio, even one with modest Sharpe.

EMH check: At equilibrium ($A_k = \bar{A}_k$ for all k), all $\mu_k = 0$, so aggregate profit is zero.

Sharpe decomposition. The mode contributions to total SR^2 are: 64.3% from mode 1 (level/-value), 28.6% from mode 2 (skew), and 7.1% from mode 3 (curvature). This is consistent with the empirical observation that value strategies dominate diversified portfolios at long horizons. [Figure 1 should display a bar chart of SR_k^2 contributions by mode, illustrating the dominance of low-frequency modes.]

9.2 Power Law Verification

The three modes satisfy the power law $\tau_k = C \cdot k^{-\alpha}$ with $C = 252$ and $\alpha = 1$:

k	τ_k (data)	$\tau_k = 252/k$ (model)	Error
1	252	252	0.0%
2	126	126	0.0%
3	84	84	0.0%

[Figure 2 should plot τ_k vs k on log-log axes, showing the power law. For $N = 3$ the fit is exact; empirical validation with $N \geq 20$ modes estimated from real data is needed to confirm the power law conjecture.]

9.3 Momentum–Mean Reversion Crossover

For mode $k = 2$ with $\tau_2 = 126$ days, the expected profit from a unit displacement has the form $\mu_2(h) = \Delta A_2 \cdot (1 - e^{-h/126})$. For short horizons $h \ll 126$, this is approximately linear in h (momentum regime). For $h \gg 126$, the profit saturates at ΔA_2 (full reversion). The crossover occurs at $h \approx \tau_2$.

[Figure 3 should display $\mu_k(h)/\mu_k(\infty)$ vs h/τ_k for each mode, showing the universal S-shaped crossover. All three curves collapse onto a single master curve $f(x) = 1 - e^{-x}$, confirming that momentum and mean reversion are the same phenomenon at different timescales (Theorem 9).]

9.4 Efficient Frontier

The efficient frontier in (σ, μ) space is the parabola $\sigma^2 = \mu^2/0.315$, or equivalently $\mu = 0.561 \cdot \sigma$. The CML is the upper branch of this parabola.

[Figure 4 should display the efficient frontier parabola with CML overlay for the 3-mode example. Mark the positions of individual modes and the tangent portfolio. This figure is the spectral analogue of the classical Markowitz diagram.]

---

10. Formal Verification

All theorems are machine-verified in Lean 4 (v4.28.0 + Mathlib). The proof library consists of 14 files in LeanProofs/SpectralTrading/:

File	Theorems	What it proves
SpectralDynamics.lean	autocorr_pos, autocorr_lt_one, halfLife_pos, condVar_zero	OU dynamics
PowerLaw.lean	tauK_pos, tauK_decreasing, crossover_pos, crossover_decreasing	τ_k power law
ModeSharpe.lean	profit_sign, profit_grows_with_horizon, no_alpha_at_equilibrium	Mode Sharpe ratios
SpectralEMH.lean	spectral_emh_equilibrium, emh_paradox_resolved	Spectral EMH
StrategySpectrum.lean	spectrum_ordering, band_profit_zero_at_eq, momentum_mr_same_mode	Strategy bands
SharpeKellyDiversification.lean	diversification_sqrt_n, kelly_growth_eq_half_sharpe_sq, spectral_diversification	Sharpe, Kelly, $\sqrt{N}$
RiskReturnTradeoff.lean	no_free_lunch, more_return_more_risk, cml_increasing, second_law_of_markets	Risk-return, CML
CAPM.lean	market_beta_one, capm_linearity, capm_beta_form, factor_ordering	CAPM, APT, Fama–French
VolSmile.lean	option_price_linear, bs_is_zero_modes, smile_skew_more_persistent	Vol smile
YieldCurve.lean	three_factors_suffice, ns_is_three_modes	Nelson–Siegel
BubblesCrashes.lean	no_crisis_at_equilibrium, more_displacement_worse_crisis, flash_vs_bear	Bubbles, crashes
Behavioral.lean	loss_aversion_asymmetry, isLossAverse, no_herding_independent	Behavioral finance
Microstructure.lean	kyle_lambda_pos, fast_modes_higher_impact, market_making_is_high_k_mr	Microstructure
MainTheorem.lean	spectral_trading_theory, maxwell_analogy, one_spectrum_all_strategies	Unification

Every file compiles with lake build and contains zero sorry. The Lean type checker provides independent verification: the theorems are mathematically correct by construction.

---

11. Discussion

11.1 The Second Law of Markets

The risk-return trade-off has a thermodynamic interpretation:

Thermodynamics	Spectral market
Temperature T	Mode temperature $\sigma_k^2\tau_k$
Energy E	Expected return μ_k
Entropy S	Risk σ_k
Free energy $F = E - TS$	Alpha $= \mu_k - SR_k \cdot \sigma_k$
Second law: $\Delta F \leq 0$	No free lunch: $\mu > 0 \Rightarrow \sigma > 0$
Thermal equilibrium	Market efficiency (EMH)

The "second law of markets" states: you cannot extract return (energy) without increasing risk (entropy). This is not an assumption — it is Theorem 4, proved from the independence of spectral modes.

11.2 Why These Thirteen?

The thirteen results are not arbitrary. The seven core theorems are the **complete set of consequences** of two assumptions:

1. **Independence**: the N Fourier modes are independent (orthogonal basis)
2. **Mean reversion**: each mode follows an OU process (unique stationary distribution)

From (1) alone: Pythagorean Sharpe, diversification, Kelly additivity.

From (2) alone: profit sign, equilibrium, horizon dependence.

From (1) + (2): risk-return trade-off, CML, EMH, momentum-MR unification.

Any result that holds for independent OU processes in $\mathbb{R}^N$ holds for the spectral market. We claim no additional structure.

The six extensions (Theorems 10–16) arise when we apply domain-specific interpretations to the mode structure: mapping modes to factor loadings yields CAPM/APT, mapping to the risk-neutral density yields the vol smile, mapping to forward rates yields the yield curve, and so on. Each extension adds a domain-specific identification but no new mathematical machinery.

11.3 What Is General and What Is Fourier-Specific

A natural question: do these results require the Fourier-cosine basis, or would Chebyshev polynomials, kernel eigenfunctions, or wavelets work equally well?

The answer separates cleanly into three layers:

Layer 1: The seven theorems are basis-independent. Theorems 4–9 follow from two properties: (i) mode independence (orthogonality) and (ii) OU dynamics (mean reversion with stationary distribution). Any orthogonal basis with independent OU coefficients yields the same Pythagorean Sharpe, Kelly growth, diversification $\sqrt{N}$, CML, EMH, and risk-return trade-off. The proofs never use the fact that the basis functions are $\cos(k\pi x/(b-a))$ rather than, say, Chebyshev polynomials $T_k(x)$ or kernel eigenfunctions $\phi_k(x)$.

Basis	Orthogonal?	7 theorems hold?	Convergence rate
Fourier-cosine	Yes	Yes	$O(\rho^{-N})$ exponential
Chebyshev	Yes	Yes	$O(\rho^{-N})$ exponential
Kernel PCA	Yes	Yes	Depends on kernel
Wavelets	Yes	Yes	$O(N^{-s})$ polynomial
Hermite	Yes	Yes	$O(N^{-s})$ polynomial

Layer 2: The convergence rate is Fourier $\approx$ Chebyshev. For analytic densities, Fourier-cosine and Chebyshev are essentially equivalent: the Joukowski map $z = (w + w^{-1})/2$ transforms one into the other, and both achieve the Bernstein ellipse rate $O(\rho^{-N})$ (Trefethen, 2013). Kernel methods can be faster for specific density classes (e.g., a Gaussian kernel is optimal for Gaussian-like densities) but slower for others. Wavelets give polynomial convergence $O(N^{-s})$ for C^s functions, which is slower than exponential for analytic densities but handles non-smooth cases that Fourier cannot.

Layer 3: The frequency interpretation is Fourier-specific. The mapping of mode k to "frequency" and hence to "timescale $1/k$" is unique to the Fourier basis. In this basis, $k = 0$ captures the mean, $k = 1$ captures skewness, and the power law $\tau_k \sim k^{-\alpha}$ translates naturally into "low frequencies persist, high frequencies revert." In a Chebyshev basis, k is a polynomial degree, not a frequency — the strategy spectrum table (Section 8) would lose its physical interpretation. In a kernel basis, modes are ordered by eigenvalue, not by any intrinsic timescale, and the momentum-MR unification (Theorem 9) would require a different argument.

Summary. The mathematical content (7 theorems) is universal. The intuitive content (strategy spectrum, Maxwell analogy, momentum = mean reversion at different frequencies) is Fourier-specific. The practical content (FFT computation at $O(N \log N)$) is a Fourier advantage over kernel methods ($O(N^2)$) and wavelets ($O(N)$ but with larger constants).

We chose Fourier because it provides all three layers simultaneously. But a reviewer who prefers Chebyshev or kernel methods should know: the theorems survive any basis change. Only the story changes.

11.4 The Stochastic Unification: VAR, Cointegration, GBM, and OU Are One Equation

Four foundational stochastic models of financial economics — vector autoregression (VAR), cointegration, geometric Brownian motion (GBM), and the Ornstein–Uhlenbeck (OU) process — appear in different textbooks, different literatures, and different notations. We show they are all special cases of a single object: the multivariate linear SDE

$$dX_t = (AX_t + b)\,dt + \Sigma\,dW_t$$

classified by the eigenvalue spectrum of the drift matrix A.

Eigendecompose A: write $A = P\Lambda P^{-1}$ where $\Lambda = \mathrm{diag}(\lambda_1, \dots, \lambda_n)$. In the eigenbasis $Z = P^{-1}X$, the system decouples into n independent scalar processes:

$$dZ_k = \lambda_k Z_k\, dt + \bar{b}_k\, dt + \bar{\sigma}_k\, dW_k$$

Each Z_k is a 1D OU process with mean-reversion rate $|\lambda_k|$ and half-life $\tau_k = -1/\lambda_k$ (for $\lambda_k < 0$). These are precisely the spectral modes A_k of Section 1.2.

The eigenvalue spectrum classifies the entire stochastic zoology:

Eigenvalue regime	Mode behavior	Model name
All $\lambda_k < 0$	All modes mean-revert	**Multivariate OU**
All $\lambda_k = 0$	All modes are random walks	**GBM** (in log-space)
Mixed: some $\lambda_k < 0$, some $= 0$	Some stationary, some nonstationary	**Cointegrated system**
Any of the above, discretized	$X_{t+1} = \Phi X_t + c + \varepsilon_t$	**VAR**

GBM as zero-eigenvalue OU. Geometric Brownian motion $dS = \mu S\, dt + \sigma S\, dW$ lives on $(\mathbb{R}^+, \times)$. Setting $Y = \log S$ gives $dY = (\mu - \sigma^2/2)\, dt + \sigma\, dW$, which is the linear SDE with $A = 0$. GBM is the special case where the drift matrix has only zero eigenvalues. The logarithm is the Lie algebra isomorphism $\log : (\mathbb{R}^+, \times) \to (\mathbb{R}, +)$ that maps multiplicative dynamics to additive.

Cointegration as mixed eigenvalue structure. The Granger Representation Theorem writes the VECM as $\Delta X_t = \alpha\beta' X_{t-1} + \varepsilon_t$, where $\Pi = \alpha\beta'$ has rank $r < n$. In the continuous-time limit, $\Pi \to A\,\Delta t$. The r nonzero eigenvalues of A correspond to cointegrating relationships (stationary, OU-type modes), while the $n-r$ zero eigenvalues correspond to common stochastic trends (nonstationary, GBM-type modes). The cointegrating vectors β are precisely the eigenvectors of A with negative eigenvalues.

This reframes cointegration as a spectral statement: a cointegrated system is one where the drift matrix has a nontrivial kernel. The cointegration rank equals the number of negative eigenvalues. The error-correction speed α is the magnitude of those eigenvalues.

VAR as discretized OU. Setting $\Phi = e^{A\Delta t} \approx I + A\Delta t$ gives $X_{t+1} = \Phi X_t + c + \varepsilon_t$, which is VAR(1). The eigenvalues of Φ are $\phi_k = e^{\lambda_k \Delta t}$: a unit root ($\phi_k = 1$) corresponds to $\lambda_k = 0$; a stable root ($|\phi_k| < 1$) corresponds to $\lambda_k < 0$. The Johansen (1991) trace test for cointegration rank is testing how many eigenvalues of A are zero.

The Fokker–Planck tower. The probability density $p(x, t)$ of X_t satisfies:

$$\frac{\partial p}{\partial t} = -\nabla \cdot (Ax\, p) + \frac{1}{2}\nabla \cdot (\Sigma\Sigma^T \nabla p)$$

In the eigenbasis, this separates into n independent 1D Fokker–Planck equations:

$$\frac{\partial p_k}{\partial t} = -\frac{\partial}{\partial z_k}(\lambda_k z_k\, p_k) + \frac{\bar{\sigma}_k^2}{2}\frac{\partial^2 p_k}{\partial z_k^2}$$

For $\lambda_k < 0$, each p_k converges to a Gaussian stationary distribution. The COS expansion of p_k in Fourier modes is the spectral representation of this paper. Thus the three levels form a tower:

$$\underbrace{dX = AX\,dt + \Sigma\,dW}_{\text{stochastic dynamics}} \xrightarrow{\text{Fokker–Planck}} \underbrace{\partial_t p = \mathcal{L}^* p}_{\text{density evolution}} \xrightarrow{\text{spectral decomposition}} \underbrace{p(x,t) = \sum_k c_k(t)\,\psi_k(x)}_{\text{mode coefficients}}$$

The 130 Fourier coefficients of the Spectral Fenton Distribution are projections onto the eigenmodes of this Fokker–Planck operator. The seven core theorems of Sections 4–5 are properties of the eigenvalue spectrum $\{\lambda_k\}$ that hold regardless of which level of the tower one works in.

Implication for the paper's framework. The spectral mode $A_k(t)$ with dynamics $dA_k = -(1/\tau_k)(A_k - \bar{A}_k)\,dt + \sigma_k\,dW_k$ is simultaneously: - An OU process (continuous time, one mode) - A VAR(1) component (its discretization) - A GBM factor (if $\tau_k \to \infty$) - A cointegrating direction (if $\tau_k < \infty$ while other modes have $\tau_k = \infty$)

The eigenvalue $\lambda_k = -1/\tau_k$ is the single number that determines which textbook the mode belongs to. The four stochastic models are not four theories. They are four eigenvalue regimes of one theory.

11.5 Limitations

1. **Independence is approximate.** Real modes may have cross-correlations, especially during crises. The framework extends to correlated modes (replacing $\sum SR_k^2$ with $\mathbf{SR}^\top \Sigma^{-1} \mathbf{SR}$), but the proofs are more complex.
2. **The OU model is Gaussian.** Mode dynamics may have jumps, stochastic volatility, or regime switches. These extensions are left for future work.
3. **The power law is empirical.** Conjecture 1 ($\tau_k \sim k^{-\alpha}$) requires validation against data. The seven theorems hold for any positive τ_k sequence; the power law is needed only for the momentum-MR unification.
4. **Basis optimality is open.** While the seven core theorems are basis-independent, the *number* of modes needed for a given accuracy depends on the basis. An adaptive basis (e.g., kernel PCA trained on the specific market) might achieve the same accuracy with fewer modes. Whether a "canonical" basis exists for financial markets — analogous to Fourier being canonical for translation-invariant systems — is an open question.
5. **No empirical validation.** The numerical examples in Section 9 use synthetic data with hand-picked parameters. Estimation of $(\bar{A}_k, \tau_k, \sigma_k)$ from real market data (e.g., S&P 500 returns, Treasury yields, FX pairs) and out-of-sample testing of the spectral trading strategy are essential next steps. The strategy spectrum assignments (Section 8) — value at $k = 0$–1, momentum at $k = 2$–10, HFT at $k = 30$–128 — are conjectural and require empirical confirmation.

---

12. Conclusion

We have derived thirteen fundamental results in financial economics — seven core theorems and six domain extensions — from a single spectral decomposition:

#	Theorem	Classical source	Spectral form
4	Risk-return trade-off	Markowitz (1952)	$\sigma^2 = \mu^2 / \sum SR_k^2$
5	Capital Market Line	Sharpe (1964)	Slope $= \sqrt{\sum SR_k^2}$
6	Sharpe max + diversification	Markowitz (1952)	$SR^2 = \sum SR_k^2 = N \cdot s^2$
7	Kelly criterion	Kelly (1956)	$g^* = \sum SR_k^2 / 2$
8	Efficient Market Hypothesis	Fama (1970)	$\sum w_k SR_k = 0$ at equilibrium
9	Momentum = Mean reversion	Jegadeesh–Titman / Lo–MacKinlay	Same mode, different h vs τ_k
10	CAPM / Beta	Sharpe (1964)	$\beta_i = \sum b_{i,k} \sigma_k^2 / \sum \sigma_k^2$
11	APT / Fama–French	Ross (1976) / FF (1992)	Factors = spectral modes
12	Vol smile	Cont–da Fonseca (2002)	Smile $= (A_0, A_1, A_2)$
13	Yield curve	Nelson–Siegel (1987)	3 factors $= k = 0, 1, 2$
14	Bubbles / crashes	Kindleberger (1978)	Bubble $= \lvert A_0 - \bar{A}_0 \rvert > c\sigma$
15	Behavioral finance	Kahneman–Tversky (1979)	Loss aversion $= \tau^+ < \tau^-$
16	Microstructure	Kyle (1985)	$\lambda_k \propto \sigma_k^2 / \tau_k$

The unifying principle: **the market is a spectrum.** Each mode k is an independent source of risk and return. The Sharpe ratio adds in quadrature (Pythagorean theorem). Kelly growth adds linearly. The efficient frontier is a parabola whose curvature is determined by the spectral Sharpe vector. Momentum and mean reversion are the low-frequency and high-frequency limits of the same dynamics.

Shannon showed that a signal is a spectrum. Markowitz showed that a portfolio is a vector. We show that a market is a spectrum of vectors — and from this, everything follows.

Acknowledgements

The author used Large Language Models for assistance with exposition. The mathematical framework, theorems, and proofs are the author's own work. The Lean 4 type checker provides independent verification.

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

- Markowitz, H (1952). Portfolio selection. *Journal of Finance*, 7(1), 77-91.
- Sharpe, W.F (1964). Capital asset prices. *Sharpe, W.F.*, 19(3).
- Kelly, J. L (1956). A New Interpretation of Information Rate. *Bell System Technical Journal*, 35(4), 917-926. DOI: 10.1002/j.1538-7305.1956.tb03809.x
- Fama, E.F (1970). Efficient capital markets. *Fama, E.F.*, 25(2).
- Jegadeesh, N. and Titman, S (1993). Returns to buying winners and selling losers. *Jegadeesh, N. and Titman, S.*, 48(1).
- 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
- Lo, A.W. and MacKinlay, A.C (1990). When are contrarian profits due to stock market overreaction? *Review of Financial Studies*, 3(2), 175–205. *Lo, A.W. and MacKinlay, A.C.*, 3(2).
- Fama, E.F. and French, K.R (1992). The cross-section of expected stock returns. *Fama, E.F. and French, K.R.*, 47(2).
- Koijen, R.S.J., Moskowitz, T.J., Pedersen, L.H., and Vrugt, E.B (2018). Carry. *Journal of Financial Economics*, 127(2).
- Nagy, T. (2026). The Fenton Distribution Solved. *Working paper.*
- Nagy, T. (2026). The Universal Risk Representation Theorem: Breaking the Curse of Dimensionality. *Zenodo.* DOI: 10.5281/zenodo.18910566
- The mathlib Community, "Mathlib: a multi-purpose library for formalized mathematics," (2020). –2026. https://leanprover-community.github.io/mathlib4_docs/
- Asness, C.S., Moskowitz, T.J., and Pedersen, L.H (2013). Value and momentum everywhere. *Asness, C.S., Moskowitz, T.J., and Pedersen, L.H.*, 68(3).
- Cont, R (2001). Empirical properties of asset returns. *Cont, R.*, 1(2).
- Bouchaud, J.-P. and Potters, M (2003). Theory of Financial Risk and Derivative Pricing. *Bouchaud, J.-P. and Potters, M.*.
- Baltas, N. and Kosowski, R (2013). Momentum strategies in futures markets. Working paper. *Baltas, N. and Kosowski, R.*.
- Johansen, S (1991). Estimation and hypothesis testing of cointegration vectors in Gaussian vector autoregressive models. *Johansen, S.*, 59(6).
- Hamilton, J.D (1994). Time Series Analysis. *Hamilton, J.D.*.
- Lütkepohl, H (2005). New Introduction to Multiple Time Series Analysis. *New Introduction to Multiple Time Series Analysis.*.
- Engle, R.F. and Granger, C.W.J (1987). Co-integration and error correction: representation, estimation, and testing. *Engle, R.F. and Granger, C.W.J.*, 55(2).
- Nelson, C.R. and Siegel, A.F (1987). Parsimonious modeling of yield curves. *Nelson, C.R. and Siegel, A.F.*, 60(4). DOI: 10.1086/296409
- Ross, S. A (1976). The arbitrage theory of capital asset pricing. *Journal of Economic Theory.* DOI: 10.1016/0022-0531(76)90046-6

- Cont, R. and da Fonseca, J (2002). Dynamics of implied volatility surfaces. *Cont, R. and da Fonseca, J.*, 2(1). DOI: 10.1088/1469-7688/2/1/304
- Litterman, R. and Scheinkman, J (1991). Common factors affecting bond returns. *Litterman, R. and Scheinkman, J.*, 1(1). DOI: 10.3905/jfi.1991.692347
- Diebold, F.X. and Li, C (2006). Forecasting the term structure of government bond yields. *Diebold, F.X. and Li, C.*, 130(2). DOI: 10.3386/w10048
- Kindleberger, C.P (1978). Manias, Panics, and Crashes: A History of Financial Crises. *Kindleberger, C.P.*.
- Kahneman, D. and Tversky, A (1979). Prospect theory: An analysis of decision under risk. *Kahneman, D. and Tversky, A.*, 47(2).
- Kyle, A.S (1985). Continuous auctions and insider trading. *Kyle, A.S.*, 53(6).
- Trefethen, Lloyd N. (2013). Approximation Theory and Approximation Practice. SIAM. DOI: 10.1137/1.9781611975949
- 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
- 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
- Cochrane, J. H (2001). Asset Pricing. *Asset Pricing*, 0-691.
- Gatheral, J (2006). The Volatility Surface: A Practitioner's Guide. *The Volatility Surface: A Practitioner's Guide*.
- Laloux, L., Cizeau, P., Bouchaud, J.-P. and Potters, M (1999). Noise dressing of financial correlation matrices. *Physical Review Letters*, 83(7). DOI: 10.1103/physrevlett.83.1467