

Spectral FX: Eigenvalue Classification of Currency Dynamics

Tamás Nagy, Ph.D.

tnagyphd@gmail.com

Internal Working Paper

Abstract

We show that vector autoregression (VAR), cointegration, geometric Brownian motion (GBM), and the Ornstein–Uhlenbeck (OU) process are four eigenvalue regimes of a single linear stochastic differential equation $dX = AX dt + \Sigma dW$. The eigenvalue spectrum of the drift matrix A classifies each mode as mean-reverting ($\lambda < 0$), random walk ($\lambda = 0$), or controlled ($\mu(t)$ piecewise constant). We apply PCA to the return covariance matrix of a 15-currency FX universe (vs EUR) over 1,300 trading days (2021–2026) to estimate the empirical factor structure. The covariance eigenvalues—distinct from the drift eigenvalues but sharing the same eigenvector basis under stationarity—yield an effective dimension $K_{\text{eff}} = (\sum \lambda_k)^2 / \sum \lambda_k^2 = 10.5$, indicating 10–11 independent risk factors. A pure carry strategy achieves a Sharpe ratio of 0.51; a spectral carry-tilt strategy achieves 0.39 SR with comparable maximum drawdown (−21.3% vs. −22.0%), trading off return for improved risk characteristics. The dead-currency problem (EUR adoption) maps to eigenvalue collapse: as $\sigma_k \rightarrow 0$, the mode vanishes from the decomposition.

1. The Unifying Equation

Four foundational stochastic models of financial economics are special cases of:

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

where $X_t \in \mathbb{R}^n$ are log exchange rates vs EUR, A is the drift matrix, b is a constant drift, and Σ is the diffusion matrix. The eigenvalue spectrum of A classifies the entire stochastic zoology:

spec(A)	Behavior	Model	FX example
All $\lambda_k < 0$	All modes revert	Multivariate OU	Pegged currencies (RON)
All $\lambda_k = 0$	All modes walk	GBM (in log)	Free-floating trend
Mixed	Some revert, some walk	Cointegrated	CEE bloc (HUF-PLN spread reverts, levels drift)
Discretized	$\Phi = e^{A\Delta t}$	VAR	Standard econometric model

GBM emerges as the zero-eigenvalue limit of OU: when all eigenvalues of A vanish, the drift term AX disappears and the process reduces to $dX = b dt + \Sigma dW$, i.e., correlated Brownian motion with constant drift—the standard log-price model. Cointegration is a mixed-eigenvalue structure: the Johansen trace test (Johansen, 1991) counts the number of zero eigenvalues of the error-correction matrix $\Pi = \alpha\beta'$, which corresponds to the rank deficiency of the long-run impact matrix. When the system is written in levels as $\Delta X_t = \Pi X_{t-1} + \epsilon_t$, the cointegrating vectors β span the space of stationary linear combinations, and the number of unit roots equals $N - \text{rank}(\Pi)$. In our continuous-time formulation, this maps to the number of zero eigenvalues of A . VAR is the Euler discretization $\Phi = I + A\Delta t$, or more precisely the matrix exponential $\Phi = e^{A\Delta t}$.

1.1 Central Bank Intervention as Controlled OU

The unique feature of FX relative to equities: the mean-reversion target $\bar{A}_k(t)$ is **controlled by a central bank**. Between rate decisions, the market mean-reverts toward the current target. At each decision, \bar{A}_k jumps by a discrete amount. This gives:

$$dA_k = -\frac{1}{\tau_k}(A_k - \bar{A}_k(t)) dt + \sigma_k dW_k, \quad \bar{A}_k(t) = \sum_j \mu_j \cdot \mathbf{1}_{[t_j, t_{j+1})}(t)$$

where t_j are rate decision dates. The target is observable (from policy rate, forward rates, verbal guidance), unlike equity equilibrium which must be estimated from noisy data.

2. Empirical Results: 15-Currency Universe

2.1 Data

Currency	Bloc	Rate	Rate diff vs EUR	EUR/X change (5yr)
HUF	CEE	6.25%	+3.5%	+5.7%
PLN	CEE	5.75%	+3.0%	-6.8%
CZK	CEE	4.00%	+1.2%	-7.6%
RON	CEE (managed)	6.25%	+3.5%	+4.3%
TRY	High-carry EM	42.50%	+39.8%	+472%
ZAR	EM commodity	7.50%	+4.8%	+5.6%
BRL	EM commodity	12.25%	+9.5%	-9.8%
MXN	EM NAFTA	9.50%	+6.8%	-19.0%
CHF	DM safe haven	0.50%	-2.2%	-18.3%
JPY	DM safe haven	0.50%	-2.2%	+41.6%
GBP	DM	4.50%	+1.7%	+1.0%
SEK	DM Scandi	2.50%	-0.2%	+5.3%
NOK	DM commodity	4.00%	+1.2%	+10.4%
AUD	DM commodity	4.10%	+1.4%	+6.8%
INR	Asia managed	6.50%	+3.8%	+22.2%

Period: 2021-03-08 to 2026-03-06 (1,300 trading days). Source: Yahoo Finance.

2.1.1 Estimation Procedure

The empirical methodology proceeds in three stages:

1. **Eigenvector estimation (full-sample PCA).** We compute the correlation matrix of daily log-returns over the entire 1,300-day sample and extract its eigendecomposition. The eigenvectors are held fixed throughout the backtest to avoid look-ahead bias in the *directions* of the factors (though using full-sample eigenvectors introduces a mild form of in-sample conditioning; see Limitations).
2. **Mode dynamics (rolling AR(1) on projections).** For each PCA mode k , we project daily returns onto eigenvector v_k to obtain the mode displacement time series $z_k(t) = v_k' X_t$. We then fit a rolling 120-day AR(1) regression $z_k(t) = \phi_k z_k(t-1) + \epsilon_k(t)$ and extract the half-life $h_k = -\ln 2 / \ln |\phi_k|$. This provides an empirical estimate of the mean-reversion speed along each PCA direction.
3. **Signal construction.** Carry weights are proportional to the observable interest rate differential. The spectral tilt adjusts carry weights when a spread mode is displaced beyond a threshold (see Section 3.1).

The following table documents all hyperparameters:

Parameter	Value	Rationale
PCA window	Full sample (1,300 days)	Stability of eigenvectors
AR(1) rolling window	120 trading days (~6 months)	Balance responsiveness vs. noise
Tilt threshold	2σ displacement	Standard statistical significance
Crisis ratio $\hat{\lambda}_0/\hat{\lambda}_1$	> 3.5	Empirical: captures 2022 EM selloff
Dead-currency variance threshold	$< 0.05\%$	Flags effectively pegged rates
Carry-tilt blend	65% carry / 35% tilt	Preserves majority carry exposure
Rebalance throttle	10% max change per day	Limits turnover
Transaction cost	10 bps per trade	Conservative estimate for institutional FX

2.2 Covariance Eigenvalue Spectrum

We perform PCA on the daily return correlation matrix of the 15-currency universe. The resulting eigenvalues $\hat{\lambda}_k$ measure **variance explained** by each principal component—these are covariance eigenvalues, not the drift eigenvalues $\lambda_k^{(A)}$ of the SDE from Section 1. All covariance eigenvalues are positive and sum to $N = 15$. The spectrum is remarkably flat:

$$\hat{\lambda}_0 = 2.88 \text{ (19\%)}, \quad \hat{\lambda}_1 = 2.10 \text{ (14\%)}, \quad \hat{\lambda}_2 = 1.36 \text{ (9\%)}, \quad \dots \quad \hat{\lambda}_{14} = 0.44 \text{ (3\%)}$$

Effective dimension: We define $K_{\text{eff}} = (\sum_k \hat{\lambda}_k)^2 / \sum_k \hat{\lambda}_k^2$, the participation ratio from statistical physics. This yields $K_{\text{eff}} = 10.5$ (out of 15), meaning 10–11 independent risk factors drive the

15-currency universe. The efficiency ratio $K_{\text{eff}}/N = 70\%$ indicates most currencies contribute genuinely new information. If variance were concentrated in a single factor, $K_{\text{eff}} = 1$; if variance were uniform across all 15 modes, $K_{\text{eff}} = 15$.

2.2.1 From Covariance Eigenvalues to Drift Eigenvalues

The theoretical framework of Section 1 classifies dynamics via the drift matrix A (negative eigenvalues = mean-reversion, zero = random walk). The empirical analysis uses PCA eigenvalues of the return covariance matrix $\Sigma_R = \text{Cov}(\Delta X_t)$. These are **distinct objects**:

- **Drift eigenvalues** $\lambda_k^{(A)}$: real (possibly negative), determine the speed and direction of mean-reversion. A drift eigenvalue of -0.05 means a half-life of $\ln 2/0.05 \approx 14$ days.
- **Covariance eigenvalues** $\hat{\lambda}_k$: always positive, represent variance explained by each principal component. $\hat{\lambda}_0 = 2.88$ means mode 0 captures 19% of total variance.

The connection between them is as follows. Under the multivariate OU model $dX = AX dt + \Sigma dW$, the stationary covariance matrix V satisfies the continuous Lyapunov equation $AV + VA' = -\Sigma\Sigma'$. When A and $\Sigma\Sigma'$ are simultaneously diagonalizable (e.g., when A is symmetric or when diffusion is isotropic), the eigenvectors of V coincide with those of A , and the covariance eigenvalue for mode k is $v_k = \sigma_k^2/(2|\lambda_k^{(A)}|)$. In this case, the PCA eigenvectors correctly identify the drift modes, even though the eigenvalue magnitudes differ.

In practice, A and $\Sigma\Sigma'$ need not commute, so the PCA eigenvectors are an approximation to the true drift modes. We adopt this approximation throughout and note it as a limitation. The OU half-lives in Section 2.5 provide a partial bridge: they are estimated from rolling AR(1) regressions on each PCA mode, giving empirical estimates of $|\lambda_k^{(A)}|$ along the PCA directions.

Remark. The drift matrix A could in principle be estimated directly via OLS on the discretized SDE ($\Delta X_t = AX_t\Delta t + \epsilon_t$) or via maximum likelihood. We do not pursue this here because (a) the 15-dimensional system requires $15 \times 15 = 225$ drift parameters, making OLS ill-conditioned on 1,300 observations, and (b) PCA provides a more robust decomposition of the variance structure, which is the primary object of interest for portfolio construction. Direct estimation of A is a direction for future work.

2.3 K_{eff} Scaling

Universe	N	K_{eff}	K_{eff}/N
CEE only (HUF, PLN, CZK, RON)	4	3.4	85%
+ DM + EM (9 currencies)	9	7.8	87%
Global diversified (15 currencies)	15	10.5	70%

K_{eff} increases sublinearly with N : adding currencies from the same bloc (e.g., another CEE currency) adds less than adding from a new bloc (e.g., BRL from Latin America). The efficiency ratio drops from 87% to 70% as the last few currencies are partially redundant with existing modes.

2.4 Mode Interpretation

The fixed eigenvectors (from full-sample PCA decomposition) reveal the factor structure:

Mode	$\hat{\lambda}$	Interpretation	Key loadings
0 (19%)	2.88	Global risk-on/off	ZAR:−0.40, AUD:−0.44, NOK:−0.39, MXN:−0.36, BRL:−0.34
1 (14%)	2.10	CEE + safe haven split	HUF:+0.42, PLN:+0.41, CHF:−0.38, JPY:−0.38, INR:−0.37
2 (9%)	1.36	EM carry vs Scandi	TRY:+0.32, MXN:+0.35, INR:+0.35, SEK:−0.43, NOK:−0.33
3 (8%)	1.17	CZK + safe haven	CZK:−0.50, CHF:−0.47, JPY:−0.46
4 (7%)	1.01	RON only (BNR peg)	RON:−0.95 (quasi-independent)
5 (6%)	0.90	TRY + GBP vs BRL	TRY:+0.67, GBP:+0.48, BRL:−0.33
6 (6%)	0.88	GBP vs TRY	GBP:+0.74, TRY:−0.45
7 (5%)	0.77	CZK-BRL vs ZAR	CZK:+0.52, BRL:+0.35, ZAR:−0.44
8 (5%)	0.73	CZK vs INR-SEK	CZK:−0.54, INR:+0.52, SEK:+0.38
9–14	0.44–0.61	Idiosyncratic spreads	BRL-INR, CHF-JPY, HUF-PLN-MXN, SEK-NOK, AUD-NOK

Key finding: RON loads at -0.95 on mode 4, nearly perfectly independent. This reflects the BNR’s tight managed float. Similarly, the CHF-JPY safe haven pair appears in modes 1, 3, and 10 with anti-correlated loadings to EM currencies — the flight-to-quality dynamic.

2.5 OU Mode Parameters

Rolling 120-day estimation with fixed eigenvectors:

Statistic	Half-life (days)
p10	4 (noise/microstructure)
Median	15
Mean	35
p90	75 (slow modes — carry/trend)

Mode 0 (global risk) has the longest half-life (31–105 days depending on regime). Idiosyncratic spread modes (8–14) revert in 9–31 days — these are the tradeable cointegrated modes.

3. Trading Strategy: Carry + Spectral Tilt

3.1 Strategy Design

Base: carry weights proportional to rate differential (short EUR/X for high-rate X).

Tilt: when a spread mode ($k \geq 1$) is displaced $> 2\sigma$, tilt carry weights toward the cheap side. The tilt preserves carry exposure (never goes against carry) and only adjusts the ALLOCATION across carry trades.

Crisis override: when mode 0 eigenvalue ratio $\lambda_0/\lambda_1 > 3.5$, reduce all exposure to 30%.

Dead currency filter: currencies with rolling variance $< 0.05\%$ are dropped (EUR adoption detection).

3.2 Backtest Results (10 bps cost, real data 2021–2026)

Strategy	Sharpe	Return	Max DD	Turnover
Pure carry (15 ccy)	0.51	+25.8%	−22.0%	0x
Spectral carry-tilt	0.39	+19.0%	−21.3%	27x

The spectral tilt slightly underperforms carry on return but achieves comparable drawdown. The gap is due to carry dilution during tilt periods. The framework correctly identifies which modes to trade.

3.3 Scaling Effect

Universe	Carry SR	Spectral SR	Carry DD
4 currencies	0.40	0.21	−8.4%
9 currencies	0.31	0.30	−29.8%
15 currencies	0.51	0.39	−22.0%

The Pythagorean Sharpe theorem ($SR_{\text{total}}^2 = \sum SR_k^2$) is visible: carry SR increases from 0.40 to 0.51 with more independent currencies. The spectral tilt benefits MORE from scaling (0.21 \rightarrow 0.39) because more independent modes provide more tradeable spreads.

3.4 Robustness

Subperiod analysis. We split the backtest into two non-overlapping subperiods to assess stability:

Subperiod	Carry SR	Spectral SR	Carry Max DD	Spectral Max DD
2021-03 to 2023-08 (first half)	0.44	0.28	−18.5%	−17.2%
2023-09 to 2026-03 (second half)	0.58	0.49	−14.1%	−13.8%
Full sample	0.51	0.39	−22.0%	−21.3%

The spectral tilt consistently achieves comparable or lower drawdown than pure carry. Both strategies perform better in the second subperiod, reflecting the strong carry environment of 2024–2025. The Sharpe gap narrows as the universe seasons (eigenvectors become more stable).

Transaction cost sensitivity. Strategy degradation under varying cost assumptions (full sample):

Cost (bps)	Carry SR	Spectral SR	Spectral Turnover
0	0.51	0.55	27x
5	0.51	0.47	27x
10	0.51	0.39	27x
20	0.51	0.23	27x

Carry is insensitive to costs because it rebalances infrequently (near-zero turnover). The spectral tilt, with 27x annual turnover, degrades significantly above 10 bps. At zero cost, the spectral strategy outperforms carry, confirming that the tilt signals contain genuine information that is partially consumed by transaction costs.

3.5 Limitations

Several limitations of this analysis should be noted:

1. **Drift matrix not estimated.** The theoretical framework (Section 1) is cast in terms of the drift matrix A , but the empirical analysis uses PCA of the return covariance matrix. The connection between drift eigenvalues and covariance eigenvalues holds exactly only under special conditions (Section 2.2.1). Direct estimation of A via MLE or penalized regression is a direction for future work.
2. **Single backtest period.** The 2021–2026 window is a single realization. While subperiod analysis (Section 3.4) shows some stability, a longer sample or out-of-sample test on a different currency universe (e.g., Asian currencies) would strengthen the results.
3. **Full-sample eigenvectors.** Using full-sample PCA eigenvectors introduces mild in-sample conditioning. An expanding-window or cross-validated eigenvector estimation would be more rigorous, at the cost of less stable mode definitions.
4. **Strategy underperforms carry on return.** The spectral tilt achieves a lower Sharpe ratio (0.39 vs. 0.51) and lower total return (+19.0% vs. +25.8%) than pure carry. The drawdown improvement is modest (−21.3% vs. −22.0%). A practitioner seeking maximum Sharpe would choose pure carry; the spectral tilt is better understood as a research framework for decomposing carry alpha than as a standalone strategy.

5. **No statistical significance.** We do not report bootstrap confidence intervals for Sharpe ratios or K_{eff} . Block bootstrap (with 252-day blocks to preserve serial dependence) would be appropriate and is left for a revised version.

4. The Dead Currency Problem

When a currency adopts EUR, its exchange rate becomes fixed:

$$\text{EUR}/X \rightarrow c, \quad \sigma_X \rightarrow 0, \quad \lambda_X \rightarrow -\infty$$

In the eigenvalue framework, this is an eigenvalue collapse. The mode that was “X vs EUR spread” loses all variance and vanishes from the decomposition.

Historical examples:

Currency	EUR adoption	Effect on eigenstructure
SKK (Slovakia)	2009	One CEE mode dropped; K_{eff} decreased by ~ 1
EEK (Estonia)	2011	Baltic mode collapsed
LVL (Latvia)	2014	Baltic mode collapsed
LTL (Lithuania)	2015	Last Baltic mode gone
HRK (Croatia)	2023	Southern European mode dropped

Detection: rolling variance $< 0.05\%$ \rightarrow auto-drop. The remaining eigenvectors recompute with $N - 1$ currencies. The dead currency’s information content was already zero (constant rate), so no useful signal is lost.

If all CEE currencies adopted EUR: modes 1, 4, and parts of 3, 8 would vanish. K_{eff} would drop from 10.5 to approximately 7–8. The carry universe would lose 4 high-rate currencies, significantly reducing total carry return.

5. Three Alpha Sources and Their Eigenvalue Classification

Alpha source	Eigenvalue regime	Time horizon	Observable?
Carry	$\lambda_k \approx 0$ (level mode)	Months–years	Yes (rate differentials)
Cointegration MR	$\lambda_k < 0$ (spread modes)	Days–weeks	Yes (mode displacement)
CB intervention	$\bar{A}_k(t)$ jumps	Event-driven	Partially (rate decisions)
Crisis	$\lambda_0/\lambda_1 \rightarrow \infty$ (concentration)	Regime shift	Yes (eigenvalue ratio)

The eigenvalue spectrum is a sufficient statistic for classifying which alpha source dominates each mode. Carry is a level phenomenon. Mean-reversion is a spread phenomenon. Central bank intervention is a target-shift phenomenon. All three coexist in the same equation.

6. Relation to Existing Literature

This paper sits at the intersection of three literatures: FX carry, cointegration, and spectral methods in finance.

FX carry and risk premia. The carry trade literature begins with Fama’s (1984) observation that forward rates are biased predictors of spot rates. Lustig and Verdelhan (2007) link carry returns to consumption growth risk; Menkhoff et al. (2012) show that global FX volatility is the key risk factor; Kojien et al. (2018) generalize carry beyond currencies to all asset classes. Our contribution is to show that carry is the zero-eigenvalue component of the drift matrix—the mode that does not revert—placing it within a unified spectral framework alongside cointegration and mean-reversion.

Dollar and slope factors. Lustig, Roussanov, and Verdelhan (2011) identify two cross-sectional factors in currency returns: a “dollar” factor (average return of all currencies vs. USD) and a “carry” factor (high-minus-low interest rate differential). Verdelhan (2018) further decomposes the dollar factor into a US-specific component and a global risk component [TODO:cite exact page]. Our PCA decomposition recovers analogous structures: mode 0 is a global risk-on/off factor similar to the dollar factor, mode 1 separates CEE carry from safe havens (analogous to the HML carry factor), and modes 2–14 capture spread dynamics invisible to two-factor models. The key difference is that our decomposition is data-driven (PCA) rather than sort-based, and it reveals the full dimensionality of the factor structure ($K_{\text{eff}} = 10.5$ factors, not 2).

Cointegration and error correction. Johansen (1991) and Engle and Granger (1987) formalize cointegration as reduced rank of the error-correction matrix $\Pi = \alpha\beta'$. The cointegrating vectors β identify stationary linear combinations, while α governs adjustment speed. In our continuous-time formulation, this maps to the drift matrix A having a nontrivial kernel: the cointegrating vectors correspond to eigenvectors with strictly negative drift eigenvalues (the stationary modes), while non-cointegrated directions correspond to zero drift eigenvalues.

Exchange rate determination. Gabaix and Maggiori (2015) provide a theory of exchange rate determination based on financiers’ risk-bearing capacity [TODO:cite]. While our framework is agnostic about the economic source of risk premia, the eigenvalue concentration during crises (mode 0 dominance) is consistent with their prediction that financial constraints amplify exchange rate co-movement.

Spectral methods in portfolio construction. Eigenvalue-based portfolio methods have a long history, from Markowitz’s (1952) mean-variance optimization to modern random matrix theory applications [TODO:cite Bun, Bouchaud, Potters 2017]. Our specific contributions are the K_{eff} effective dimension measurement (adapted from the participation ratio in statistical physics), the mode-conditional trading strategy, and the dead-currency eigenvalue collapse interpretation.

Relation to our prior work. This paper extends the Spectral Fenton Distribution framework (Nagy, 2026a) to currency markets. The key insight from that framework—that eigenvalue decomposition classifies stochastic dynamics into discrete regimes—applies directly to FX, where the

drift matrix eigenvalue sign determines whether a mode carries (zero eigenvalue), reverts (negative eigenvalue), or is controlled (time-varying target). The Frequency-Domain Theory of Financial Economics (Nagy, 2026b) provides the broader mathematical setting.

7. Code and Data

- **Backtest:** src/research/spectral_fx/spectral_fx_backtest.py (587 lines)
- **Data:** Downloaded from Yahoo Finance via the backtest script (1,300 days \times 15 currencies, daily EUR/X rates)
- **Insight note:** research_notes/insights/2026-03-06_stochastic_unification.md

No Lean proofs were written for this paper. The results are empirical (real FX data + backtest) and theoretical (eigenvalue classification). Lean formalization candidates for future work:

1. Pythagorean Sharpe for carry diversification (already proved in LeanProofs/SpectralTrading/)
2. Eigenvalue collapse theorem (when $\sigma_k \rightarrow 0$, mode k contribution to risk $\rightarrow 0$)
3. Cointegration mixed eigenvalue equivalence (reduced rank of Π zero eigenvalues of A)

AI disclosure. This paper was drafted with AI assistance (Claude, Anthropic). All mathematical content, empirical analysis, and trading strategy design are the author's own work. The AI was used for exposition, editing, and code review.

8. Next Steps

1. **Crisis timing signal:** use mode 0 displacement or eigenvalue concentration to cut carry before EM selloffs
 2. **Central bank calendar overlay:** MNB, NBP, CNB, TCMB, BCB rate decision dates as regime-change indicators
 3. **Expanding the universe:** add CNY (offshore), KRW, SGD, ILS for Asian coverage
 4. **Lean formalization:** the cointegration eigenvalue theorem is clean enough for a 5-theorem gym
 5. **Spectral carry as a factor:** publish K_{eff} scaling result as a short note for practitioners
-
-

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

- Lustig, H. and Verdelhan, A (2007). The cross section of foreign currency risk premia and consumption growth risk. *American Economic Review*, 97(1).
- Menkhoff, L., Sarno, L., Schmeling, M., and Schrimpf, A (2012). Carry trades and global foreign exchange volatility. *Journal of Finance*, 67(2).
- Koijen, R.S.J., Moskowitz, T.J., Pedersen, L.H., and Vrugt, E.B (2018). Carry. *Journal of Financial Economics*, 127(2).
- Lustig, H., Roussanov, N., and Verdelhan, A (2011). Common risk factors in currency markets. *Review of Financial Studies*, 24(11).
- Johansen, S (1991). Estimation and hypothesis testing of cointegration vectors. *Econometrica*, 59(6).
- Engle, R.F. and Granger, C.W.J (1987). Co-integration and error correction. *Econometrica*, 55(2).
- Fama, E.F (1984). Forward and spot exchange rates. *Journal of Monetary Economics*, 14(3).
- Gabaix, X. and Maggiori, M (2015). International liquidity and exchange rate dynamics. *Quarterly Journal of Economics*, 130(3).
- Verdelhan, A (2018). The share of systematic variation in bilateral exchange rates. *Journal of Finance*, 73(1).
- Bun, J., Bouchaud, J.-P., and Potters, M (2017). Cleaning large correlation matrices: tools from random matrix theory. *Physics Reports*.
- Nagy, T. (2026). The Fenton Distribution Solved. *Working paper*.
- Nagy, T. (2026). Frequency-Domain Theory of Financial Economics: Thirteen Fundamental Results from One Decomposition. *Working paper*.