The Currency Network

A Spectral Graph Framework for Detecting G10 FX Network Fractures

Alexander Robbins

Department of Mathematics | Department of Economics | Department of Computer Science

University of Florida | Quant Research Lead, AlgoGators

Working Paper — May 2025

Abstract

We introduce the **Currency Network**, a physics-inspired framework that models the G10 foreign exchange market as a mechanical system in a scalar displacement field. Each currency is a node with inertial mass calibrated from trading volume, relative volatility, and trend strength; bilateral exchange-rate correlations determine spring constants through a weighted signed Laplacian. We prove equilibrium existence, uniqueness, and Lipschitz continuity (Theorem 1) and establish a tight displacement bound controlled by the Fiedler eigenvalue (Theorem 2). A critical design note, stated upfront: the equilibrium `u*(t) = L⁺f(t)` is a contemporaneous mapping of day-t observed displacements, not a forecast. The model carries no predictive content for tomorrow's returns; using lagged inputs produces average $R^2 = -0.20$, confirming that the spectral transform of yesterday's shock actively misdirects rather than predicts. The framework's value is structural diagnosis, not return prediction. The primary empirical application is the **Currency Network Stress Index (CNSI)**: total elastic potential energy stored in the displaced network. Backtested on G10 FX data from Yahoo Finance (2015–2024, 2,605 trading days), the CNSI detects structural network fractures at z-scores of 9–16σ for three confirmed episodes — Brexit (June 2016, 15.2σ), the COVID synchronized selloff (March 2020, 15.8σ), and the Japan carry-trade unwind (August 2024, 9.1σ). Crucially, a simple G10 realized-volatility index correlates with CNSI-Z at only 0.21 (full-sample), and the Japan 2024 carry unwind — CNSI 9.1σ vs. RVol 2.2σ — demonstrates that the network structure detects structural fractures that high-volatility measures miss. The realized-volatility index also produces false positives on Russia/Ukraine (2.6σ), SVB (3.0σ), and Truss/Yen (5.7σ) that the CNSI correctly suppresses. We formally distinguish two crisis regimes by the sign of ΔFiedler at CNSI peaks: synchronization crises (COVID, Brexit), in which correlations converge and Fiedler rises; and fragmentation crises (Japan 2024 carry unwind), in which the carry-currency cluster splits and Fiedler declines. The false-positive rate at 2.5σ is 1.6% (32 false positives out of 2,055 non-crisis days) vs. 2.1% for the realized-volatility baseline.

Keywords: Foreign exchange, spring-mass networks, graph Laplacian, spectral graph theory, mechanical equilibrium, systemic risk, Fiedler value, crisis detection, G10 FX.

1. Introduction

Foreign exchange markets constitute the largest and most liquid financial markets on Earth, with daily turnover exceeding $7.5 trillion. Despite decades of empirical and theoretical investigation, a unified mechanical framework capturing how pressure propagates across the G10 currency network during stress episodes remains elusive. Existing approaches — uncovered interest parity, purchasing power parity, principal component factor models, and network-theoretic graph models — illuminate individual facets of FX dynamics while leaving others unexplained. In particular, none simultaneously accounts for the restoring forces that link correlated currencies and the inertial resistances that make some currencies harder to displace than others, nor do they provide an analytically grounded, real-time stress index that distinguishes *structural network fractures* from idiosyncratic currency shocks.

This paper proposes such a framework. Inspired by the classical spring-mass system from Newtonian mechanics, we embed N currencies as nodes in a weighted network where each edge carries a spring constant proportional to the rolling Pearson correlation of the corresponding log-return pair. Each node possesses a gravitational inertia — a scalar mass calibrated from three observable market quantities — that encodes its resistance to displacement. Under an external shock vector whose components sum to zero, the network relaxes to a mechanical equilibrium whose solution we characterize analytically using the Moore-Penrose pseudoinverse of the weighted Laplacian.

The principal empirical contribution is the **Currency Network Stress Index (CNSI)**: the total elastic potential energy stored in the displaced network. Backtested over a 10-year window (2015–2024), the CNSI reliably identifies episodes in which the G10 FX network itself fractured — not merely episodes in which individual currencies moved. We document three such events with CNSI z-scores between 9 and 16 standard deviations, and we show that events commonly labeled as “currency crises” but which did not propagate across the G10 network (Russia/Ukraine, SVB, the 2022 US election) produce near-zero CNSI readings. This selectivity is not a limitation but a structural feature: the CNSI measures elastic energy stored in the correlation network, so only events that reorganize cross-currency correlations simultaneously trigger it.

A secondary contribution is the formal typology of crisis regimes. We show that detected events divide into two classes by the sign of the Fiedler eigenvalue change at peak CNSI: synchronization crises (COVID, Brexit), in which all G10 correlations converge under shared flight-to-safety pressure and the Fiedler value increases; and fragmentation crises (Japan 2024 carry unwind), in which the carry-currency cluster splits from the safe-haven bloc and the Fiedler value declines. We prove that these two regimes are distinguishable from the sign of $\Delta\delta$ at the CNSI peak.

We are explicit about the model's limits. The static equilibrium $u^* = L^+f$, by construction, is a contemporaneous function of today's observed displacements; it is not a forecast of tomorrow's returns. No statistically meaningful one-step-ahead R^2 is produced. The mass exponents (β_1, β_2, β_3) are treated as a structural prior; calibration is not possible because u^* is invariant to M (Proposition 1). The model uses Yahoo Finance mid-close prices rather than institutional Refinitiv tick data. Within these constraints, the CNSI result is robust and empirically grounded.

The remainder of the paper is organized as follows. Section 2 establishes the mathematical setup. Section 3 develops the equations of motion and proves equilibrium existence and stability. Section 4 introduces the gravitational pressure metric. Section 5 presents the spectral analysis and displacement bound. Section 6 extends the framework to time-varying networks. Section 7 describes the data, implementation, and empirical results. Section 8 defines the CNSI and presents the crisis detection backtest. Section 9 discusses limitations and future work. Section 10 concludes.

2. The Currency Network: Mathematical Setup

2.1 Node Representation and Displacement

Let $C = \{c_1, c_2, \ldots, c_n\}$ be a finite set of N currencies. The state of each currency c_i at time t is characterized by a scalar displacement:

$$u_i(t) \in \mathbb{R} \quad (2.1)$$

representing the deviation of c_i's log exchange rate from a benchmark fundamental trajectory. Operationally, this is the residual after removing a rolling τ-day linear trend from the log mid-price series $\ln S_i(t)$ against USD as numeraire. We collect all displacements into the global state vector:

$$u(t) = (u_1(t), u_2(t), \ldots, u_n(t))^T \in \mathbb{R}^N \quad (2.2)$$

Remark. We work in $\mathbb{R}^N$ rather than a two-dimensional plane. Network layout diagrams may use 2D embeddings for visualization, but all analytic results are stated and proved in $\mathbb{R}^N$.

2.2 Spring Constants and the Correlation Kernel

Let $r_i(t) = \ln(S_i(t)/S_i(t-1))$ denote the log-return of currency i. The rolling Pearson correlation over a window of τ trading days ending at t is:

$$\rho_{ij}(t; \tau) = [\Sigma_s (r_i(s) - \bar{r}_i)(r_j(s) - \bar{r}_j)] / [\hat{\sigma}_i(t,\tau) \cdot \hat{\sigma}_j(t,\tau)] \quad (2.3)$$

For robustness when N is large relative to τ, we replace ρ_{ij} with the Ledoit-Wolf shrinkage estimator $\hat{\rho}_{ij}^{nL}$ [Ledoit & Wolf 2004]. The spring constant is:

$$k_{ij} = k_0 \cdot |\rho_{ij}| \quad (2.4)$$

where $k_0 > 0$ is a global stiffness scale normalized so that the Frobenius norm of K equals N. The directional coupling parameter $\alpha_{ij} = \mathrm{sgn}(\rho_{ij}) \in \{+1,-1\}$ indicates whether the spring is attractive (positive correlation) or repulsive (negative correlation).

2.3 The Stiffness Matrix: Signed Weighted Laplacian

The signed weighted Laplacian captures both attractive and repulsive springs:

$$L_{ij} = \Sigma_{j\neq i}\, k_{ij} \quad (i = j), \qquad -\alpha_{ij}\, k_{ij} \quad (i \neq j) \quad (2.5)$$

For networks with predominantly positive correlations (typical in G10 FX), L is diagonally dominant and positive semi-definite. In all G10 calibrations reported here, L was found to be PSD. For brevity we write L for the signed Laplacian throughout.

2.4 Gravitational Inertia: Currency Mass

Each currency c_i is assigned a scalar mass $m_i > 0$. We operationalize mass as a Cobb-Douglas composite:

$$m_i = V_i^{\beta_1} \cdot \Omega_i^{\beta_2} \cdot IR_i^{\beta_3} \quad (2.6)$$

where: V_i is the 90-day average daily FX trading volume, normalized to [0,1]; $\Omega_i = (1 + Z_i)^{-1}$ where Z_i is the Z-score of realized volatility relative to the network average (low-volatility currencies receive higher mass); and $IR_i = |\hat{\mu}_i| / \hat{\sigma}_i$ is the information ratio (trend strength). The exponents satisfy $\beta_1 + \beta_2 + \beta_3 = 1$, $\beta_i > 0$.

Remark on mass calibration. As established in Proposition 1 (Section 4), the static equilibrium $u^* = L^+f$ is independent of M. Consequently, the exponents (β_1, β_2, β_3) cannot be calibrated by minimizing equilibrium prediction error — the objective is flat across all valid exponent triples. We treat (β_1, β_2, β_3) = (0.47, 0.31, 0.22) as a structural prior motivated by BIS Triennial Survey evidence that trading volume is the strongest predictor of reserve-currency status. Mass enters the model through two channels only: (i) the gravitational pressure metric (Section 4), and (ii) the dynamic return timescale (Section 3.4). Readers interested in sensitivity should note that CNSI is also M-independent through the identity $\mathrm{CNSI} = \frac{1}{2} f^T L^+ f$.

3. Equations of Motion and Equilibrium

3.1 The Coupled ODE System with Rayleigh Damping

The mechanical motion of the Currency Network under external shock vector $f(t) \in \mathbb{R}^N$ is governed by Newton's second law with mass-proportional (Rayleigh) damping:

$$M\ddot{u}(t) + \gamma M\acute{p}(t) + L\,u(t) = f(t) \quad (3.1)$$

The Rayleigh damping term is economically motivated: the mean-reversion force on currency i is proportional to its market footprint, not to the spring topology. This models the empirical observation that high-volume currencies revert toward fundamental value on a timescale determined by liquidity depth rather than correlation structure. Mass-proportional damping allows the ODE to be diagonalized in the mass-weighted eigenbasis (Section 3.4), yielding closed-form solutions.

3.2 Static Equilibrium

Setting $\ddot{u} = \acute{p} = 0$ in (3.1) eliminates the mass matrix entirely, reducing to:

$$L\,u^* = f \quad (3.2)$$

Since L is singular with $\ker(L) = \mathrm{span}(1)$ for a connected network, a solution exists if and only if $f \perp \ker(L)$, i.e., $\Sigma_i\ f_i = 0$. This is the economic statement that the network is closed: total pressure across all currencies sums to zero, consistent with balance-of-payments accounting. The unique minimum-norm solution is:

$$u^* = L^+f \quad (3.3)$$

3.3 Equilibrium Existence, Uniqueness, and Continuity

> ***Theorem 1 (Equilibrium Existence, Uniqueness, and Continuity).*** *Let K be the spring constant matrix induced by a connected correlation graph, so that L is positive semi-definite of rank N−1. Let $f \in \mathbb{R}^N$ satisfy $\Sigma_i f_i = 0$. Then: (1) the equation $Lu^* = f$ has a solution in $\mathbb{R}^N$; (2) the minimum-norm solution $u^* = L^+f$ is unique in (span(1))⊥; and (3) u* is Lipschitz-continuous in f and in the spring constants k_{ij}.*
>
> *Proof.* (1) Existence follows from the Fredholm alternative: $f \perp \ker(L)$ is the necessary and sufficient condition. (2) Uniqueness on (span(1))⊥: if u_1 and u_2 are both solutions, then $L(u_1 - u_2) = 0$, so $u_1 - u_2 \in \ker(L) = \mathrm{span}(1)$. Restricting to (span(1))⊥ forces $u_1 = u_2$. (3) Continuity: L^+ is continuous in L on the manifold of

symmetric matrices with fixed rank N−1 [Golub & Van Loan 2013]. Since $k_{ij} = k_0|\rho_{ij}|$ is Lipschitz in ρ_{ij}, continuity of u* follows by the chain rule. ■

3.4 Dynamic Stability and Modal Decomposition

Define the mass-normalized Laplacian $\tilde{L} = M^{-½} L M^{-½}$, with eigenvalues $0 = \tilde{\mu}_1 \leq \tilde{\mu}_2 \leq \cdots \leq \tilde{\mu}_n$ and orthonormal eigenvectors $\tilde{q}_v$. Setting $v(t) = M^{½} u(t)$, equation (3.1) transforms into N decoupled scalar ODEs:

$$\ddot{v}_v(t) + \gamma \dot{v}_v(t) + \tilde{\mu}_v v_v(t) = (M^{-½} f)_v \quad (3.4)$$

Each modal equation is a damped harmonic oscillator with characteristic roots $\lambda_{v,\pm} = [-\gamma \pm \sqrt{(\gamma^2 - 4\tilde{\mu}_v)}] / 2$. For all nonzero modes, $\mathrm{Re}(\lambda_{v,\pm}) = -\gamma/2 < 0$, so the system is Lyapunov stable on $(\mathrm{span}(M^{½}\cdot{}^{R}))^{\perp}$. The system is neutrally stable in the zero mode, corresponding to a uniform global FX regime shift.

4. Gravitational Pressure and the Role of Mass

4.1 The Gravitational Analogy

Define the correlation distance $d_{ij} = 1/k_{ij}$. The gravitational pressure that currency j exerts on currency i is:

$$G_{ij} = G_0 \cdot m_i m_j / d_{ij}^2 = G_0 m_i m_j k_0^2 \rho_{ij}^2 \quad (4.1)$$

The total gravitational influence of currency i on the rest of the network is:

$$P_i^{nGrcp} = \Sigma_{j\neq i} G_0 m_i m_j k_0^2 \rho_{ij}^2 \quad (4.2)$$

4.2 Reserve Currency Identification

A currency with high gravitational pressure dominates the network: it exerts large restoring forces on many other currencies, is highly correlated with the rest of the network, and has high mass. On the 9-currency non-USD G10 network (USD serves as numeraire and is mechanically excluded from the node set), we assess ranking stability via block bootstrap: 500 draws of contiguous 252-day windows sampled with replacement from 2015–2024, with gravitational pressure rankings recomputed at each draw using the mass prior $(\beta_1, \beta_2, \beta_3) = (0.47, 0.31, 0.22)$. Table 4 reports the results.

Currency	Median Rank	p5 Rank	p95 Rank	Frac Top-3	Stability
EUR	1	1	8	77.6%	Robust #1 (72.2% of draws)
GBP	2	1	7	83.2%	Robust top-3; most consistent
AUD	3	1	7	62.6%	Robust top-3; volatile rank within it
CHF	4	2	8	30.4%	Unstable; IQR spans ranks 2–8
JPY	4.5	2	9	28.2%	Unstable; IQR spans ranks 2–9

Currency	Median Rank	p5 Rank	p95 Rank	Frac Top-3	Stability
CAD	5	3	8	15.6%	Mid-tier; order not stable
NZD	6	4	7	1.4%	Mid-tier; order not stable
SEK	6	4	8	1.0%	Mid-tier; order not stable
NOK	9	8	9	0.0%	Robustly last in all draws

Table 4: Bootstrap gravitational pressure rankings (500 block-bootstrap draws, 252-day windows). Median rank, 5th/95th percentile, and fraction of draws in top 3. EUR is robustly #1 (72.2% of draws); GBP and AUD are robustly top-3. Mid-tier ordering (ranks 4–8) is unstable and should not be interpreted.

EUR's dominance is consistent with BIS Triennial Survey data on reserve currency holdings and spot market share. GBP's bootstrap stability (top-3 in 83.2% of draws) is a stronger result than the terminal-window point estimate, which had placed GBP at #7 — a snapshot artifact of the specific 2024 correlation window. AUD's top-3 presence (62.6% of draws) reflects its high spot-market volume and strong commodity-currency correlation structure. Mid-tier rankings (CHF, JPY, CAD, NZD, SEK, ranks 4–8) are unstable across bootstrap draws and should not be interpreted as a stable ordering. NOK is robustly last across all 500 draws. Because USD serves as numeraire and is mechanically excluded from the 9-node network, no claim about USD's gravitational ranking can be made from this formulation.

4.3 Mass Governs Dynamic Return Speed, Not Static Displacement

Proposition 1 (Mass Governs Dynamic Return Speed, Not Static Displacement). *For a unit shock $f = e_j$ (concentrated at node j), the equilibrium displacement of node i is $u^*_i = (L^+)_{ij}$. This quantity is independent of m_j and independent of m_i: the static equilibrium $u^* = L^+f$ depends only on L and f. The role of mass is entirely dynamic: the decay envelope of the transient response of mode v is governed by $e^{(-\gamma t/2)}$, and the oscillation frequency in the underdamped regime is $\omega_v = \sqrt{(\tilde{\mu}_v - \gamma^2/4)}$, where $\tilde{\mu}_v$ depends on the mass matrix through the normalized Laplacian $\tilde{L} = M^{(-1/2)} L M^{(-1/2)}$. Thus mass strictly governs the speed of return to equilibrium, not the magnitude of displacement at equilibrium. Heavy currencies return to equilibrium more slowly (larger m_i lowers $\tilde{\mu}_v$, slowing transient dynamics), but their equilibrium displacement $u^*_i = (L^+)_{ij}$ is identical to that of a massless system sharing the same Laplacian.*

Proof. From the static equation (3.2), setting $\ddot{u} = \dot{p} = 0$ in (3.1) eliminates the mass matrix entirely: $Lu^* = f$, so $u^* = L^+f$ is independent of M by inspection. For the dynamic claim: from (3.4), each modal coordinate $v_v(t)$ satisfies a damped harmonic oscillator with eigenfrequency $\tilde{\mu}_v$, an eigenvalue of $\tilde{L} = M^{(-1/2)} L M^{(-1/2)}$. Scaling all masses by a factor $\alpha > 1$ replaces $\tilde{\mu}_v$ by $\tilde{\mu}_v/\alpha$, which lowers the oscillation frequency (underdamped case) and increases the return timescale, while leaving the steady-state u^* unchanged. ■

5. Spectral Analysis of the Network

5.1 Eigenstructure of the Laplacian

Since L is real symmetric and PSD, it admits the spectral decomposition $L = Q\Lambda Q^T = \Sigma_v \mu_v q_v q_v^T$. The pseudoinverse is:

$$L^+ = \Sigma_{v\geq 2} (1/\mu_v)\, q_v\, q_v^T \quad (5.1)$$

The equilibrium is a superposition of spectral modes:

$$u^* = L^+f = \Sigma_{v\geq 2} (q_v^T f / \mu_v)\, q_v \quad (5.2)$$

The *Fiedler vector* q_2 (corresponding to the smallest nonzero eigenvalue $\mu_2 = \delta$, the Fiedler value) identifies the minimum-cut partition of the network — the most fragile division of the currency universe [Fiedler 1973].

5.2 Spectral Gap and Network Resilience

The spectral gap $\delta = \mu_2$ controls network resilience: large δ means a strongly connected network where shocks dissipate quickly; small δ means the network is near disconnection along the Fiedler cut, concentrating shocks in one cluster.

> ***Theorem 2 (Displacement Bound).*** *For any shock vector f with $\|f\|_2 = F$ and $\Sigma_i f_i = 0$: $\|u^*\|_2 \leq F / \mu_2 = F / \delta$, with equality when $f \propto q_2$.*
>
> *Proof.* By the spectral decomposition (5.2) and Bessel's inequality: $\|u^*\|_2^2 = \Sigma_{v\geq 2} (q_v^T f)^2 / \mu_v^2 \leq (1/\mu_2^2) \Sigma_{v\geq 2} (q_v^T f)^2 \leq \|f\|_2^2 / \mu_2^2$. Taking square roots yields the bound. The bound is tight when f is proportional to q_2. ■

Theorem 2 provides the theoretical motivation for monitoring $\delta(t)$ as a structural fragility indicator. However, empirical calibration (Section 8) reveals that the relationship between δ and CNSI depends critically on the crisis regime, as formalized in Section 5.3.

5.3 Two Crisis Regimes: Synchronization vs. Fragmentation

A key empirical finding is that CNSI peaks divide into two structurally distinct regimes, distinguished by the sign of $\Delta\delta$ at the CNSI peak.

Synchronization crises (COVID March 2020, Brexit June 2016): A large simultaneous shock drives all G10 currencies in correlated directions — a flight-to-safety synchronized selloff. This increases all pairwise correlations ρ_{ij}, raising all spring constants k_{ij}, which raises the Fiedler value δ (the network becomes more tightly coupled, not more fragile). The CNSI spikes because the forcing vector f is enormous, not because δ is small. In these events, $\Delta\delta > 0$ at peak CNSI.

Fragmentation crises (Japan carry unwind August 2024): A shock drives a subset of currencies in one direction while others move oppositely, splitting the network into subclusters with low or negative cross-cluster correlations. In the Japan 2024 case, the JPY reversal caused AUD, NZD, and CAD (carry-funding recipients) to sell off while JPY and CHF (safe havens) appreciated — widening the spread between the two blocs and reducing Fiedler. Spring constants across the carry/safe-haven cut decline, reducing δ. In these events, $\Delta\delta < 0$ before or at peak CNSI, consistent with the theoretical prediction of Theorem 2. Fragmentation crises are also structurally distinct from synchronization crises in their volatility signature: because individual currencies move in opposite directions, aggregate G10 realized volatility need not spike — explaining why RVol-Z was only 2.2σ for Japan 2024.

Formal criterion: Let $\delta(t)$ denote the Fiedler value at time t and let t* denote the CNSI peak date. We classify:

- **Synchronization crisis:** $\delta(t^*) > \delta(t^* - 60)$, i.e., Fiedler increased over the 60-day pre-crisis window.
- **Fragmentation crisis:** $\delta(t^*) < \delta(t^* - 60)$, i.e., Fiedler declined over the 60-day pre-crisis window.

This typology replaces the earlier, incorrect narrative that the Fiedler value universally declines before crises. The empirical data show the opposite for synchronization events, which are the most severe (highest CNSI z-scores). The Fiedler value is a meaningful indicator only for fragmentation crises, not for the class of events that dominate the top of the CNSI ranking.

5.4 Mode Decomposition and Economic Interpretation

In G10 calibrations on 2015–2024 data, the spectral modes have stable economic interpretations:

- **Mode 2 (Fiedler):** Aligns with the risk-on/risk-off axis — JPY and CHF positive loadings vs. AUD, NZD, and CAD negative loadings.
- **Mode 3:** Captures commodity-currency clustering — AUD/CAD/NZD vs. JPY/CHF safe-haven bloc.
- **Modes 4–5:** Carry-trade groupings and region-specific correlations (EUR/GBP vs. commodity proxies).

6. Dynamic Extension: Time-Varying Networks

6.1 Rolling Window Formulation

In practice, the correlation structure is non-stationary. We define the time-varying Currency Network CN(t) by computing K(t) over a rolling window of length τ ending at t. The equilibrium displacement path:

$$u^*(t) = L(t)^+ f(t) \quad (6.1)$$

evolves continuously. The network velocity $v^*(t) = u^*(t) - u^*(t-1)$ (in discrete time) signals structural breaks in correlation regimes. Large $\|v^*(t)\|_2$ indicates a network phase transition.

6.2 Perturbation Theory for Shock Analysis

For a shock $f = \varepsilon(e_j - (1/N)\cdot{}^R)$ at currency j (symmetrized to satisfy the zero-sum condition), the first-order perturbation to the equilibrium is:

$$\Delta u^* \approx \varepsilon L^+ (e_j - (1/N)\cdot{}^R) = \varepsilon \Sigma_v{\geq}_2 (q_{vj} / \mu_v) q_v \quad (6.2)$$

The displacement of currency i due to a unit shock at j is $\Delta u^*_i = \varepsilon (L^+)_{ij}$ — the Green's function of the mechanical system. The matrix L^+ thus plays the dual role of equilibrium solver (3.3) and shock propagator (6.2).

7. Empirical Calibration

7.1 Data and Implementation

We calibrate the Currency Network on daily G10 FX close prices from Yahoo Finance, January 2015 to December 2024 (2,605 trading days). Currencies: EUR, GBP, JPY, CHF, AUD, CAD, NZD, SEK, NOK — all quoted as USD per unit (or units per USD for JPY), consistent with treating USD as the base numeraire. Log-returns are computed from daily close prices.

All computations use Python 3.11 with NumPy, SciPy (`scipy.linalg.pinv` for pseudoinverse), pandas, and sklearn. The Ledoit-Wolf shrinkage estimator is applied via `sklearn.covariance.LedoitWolf` throughout. The primary rolling window is $\tau = 60$ trading days for tactical-horizon analysis; $\tau = 30$ and $\tau = 252$ are used for sensitivity checks. Code is available at [repository URL].

Data source note. These results use Yahoo Finance close prices rather than Refinitiv Eikon mid-prices. Yahoo Finance provides mid-market indicative prices suitable for academic research; differences from institutional data are expected to be small for G10 major pairs but non-negligible for less liquid pairs (NOK, SEK) during stress episodes. Results should be interpreted accordingly.

7.2 Spring Constant Estimation

Given the $N\times\tau$ log-return matrix R, the Ledoit-Wolf correlation matrix is:

$$\hat{P}^{LW} = D^{-1/2}\, \hat{\Sigma}^{LW}\, D^{-1/2} \quad (7.1)$$

where Σ^{LW} is the shrinkage covariance and $D = \mathrm{diag}(\Sigma^{LW})$. The global stiffness scale is:

$$k_0 = N \,/\, \| |\hat{P}^{LW}| \|_F \quad (7.2)$$

ensuring the expected sum of spring constants per node equals one. All correlations are used in the Laplacian; no sparsification threshold is applied in the baseline (the threshold described in some earlier drafts was not implemented and its effect on CNSI performance is a subject of future work).

7.3 Mass Prior

As established in Proposition 1, the static equilibrium $u^* = L^+ f$ and the CNSI are both independent of the mass matrix M. Mass exponents are therefore set to a structural prior: $(\beta_1, \beta_2, \beta_3) = (0.47, 0.31, 0.22)$, giving volume dominant influence consistent with BIS survey data. Table 6 verifies this invariance empirically across three alternative exponent triples: peak CNSI-Z for all three confirmed crisis events is unchanged to within numerical precision across specifications, confirming that the detection result does not depend on the mass prior.

Mass Prior ($\beta_1, \beta_2, \beta_3$)	Brexit CNSI-Z	COVID CNSI-Z	Japan CNSI-Z	FP Rate	Notes
(0.47, 0.31, 0.22) — baseline	15.2σ	15.8σ	9.1σ	1.6%	Volume-dominant prior
(0.33, 0.33, 0.34) — equal weight	15.2σ	15.8σ	9.1σ	1.6%	No volume preference

Mass Prior (β_1, β_2, β_3)	Brexit CNSI-Z	COVID CNSI-Z	Japan CNSI-Z	FP Rate	Notes
(0.70, 0.20, 0.10) — volume only	15.2σ	15.8σ	9.1σ	1.6%	Extreme volume weight
(0.20, 0.60, 0.20) — vol-inverse dominant	15.2σ	15.8σ	9.1σ	1.6%	Stability-weighted

Table 6: CNSI-Z peak values and false positive rate across alternative mass exponent triples. All results are identical to four significant figures, confirming that CNSI = $\frac{1}{2} f^T L^+ f$ is independent of M by Proposition 1. The mass prior affects only the gravitational pressure ranking (Section 4.2) and the dynamic return timescale (Section 3.4), not the stress detection result.

7.4 Theorem 2 Validation

Theorem 2 predicts that $\|u^*\|_2 \leq \|f\|_2 / \delta$. We test this bound daily over 2015–2024. The bound holds on 91.1% of all trading days; the 8.9% violations are concentrated in windows where numerical rank deficiency of the rolling Laplacian causes pseudoinverse instability at very low δ values (sub-0.01). Excluding windows with $\delta < 0.01$ raises the hold rate to 96.8%. The mean ratio $\|u^*\|_2 \cdot \delta / \|f\|_2$ across all dates is 0.75, confirming the bound is not vacuous — actual displacement uses approximately 75% of the theoretical maximum on average.

8. The Currency Network Stress Index and Crisis Detection

8.1 CNSI Definition

Define the CNSI as the total elastic potential energy stored in the network:

$$\mathrm{CNSI}(t) = \tfrac{1}{2}\, u^*(t)^T L(t)\, u^*(t) = \tfrac{1}{2}\, f(t)^T L^+(t)\, f(t) \quad (8.1)$$

The identity $u^{*T} L u^* = f^T L^+ f$ follows from $u^* = L^+ f$ and the standard pseudoinverse identity $L L^+ L = L$. Note that CNSI is independent of M, confirming that mass specification does not affect detection performance. CNSI rises when the network is under high external pressure ($\|f\|$ large) and/or when the shock aligns with the Fiedler mode (which Theorem 2 shows amplifies displacement most). For real-time monitoring, we compute the rolling z-score:

$$\mathrm{CNSI\text{-}Z}(t) = [\mathrm{CNSI}(t) - \mu(t-1)] / \sigma(t-1) \quad (8.2)$$

where μ and σ are the mean and standard deviation computed over the trailing 252 trading days, using a strictly lagged window (`shift(1)`) to avoid including the current value in its own baseline.

8.2 Crisis Detection Results

Table 1 reports CNSI-Z at peak for all identified stress episodes over 2015–2024, classified by regime type. The threshold for a detection is CNSI-Z $\geq 2.5\sigma$.

Event	CNSI-Z	Regime	Why Detected
Brexit (Jun 2016)	15.2σ	Sync.	GBP crashed 10% overnight, reorganizing the entire G10 correlation structure simultaneously
COVID Shock (Mar 2020)	15.8σ	Sync.	Synchronized global flight-to-safety across all G10 currencies; largest CNSI event in sample
Japan Carry Unwind (Aug 2024)	9.1σ	Frag.	JPY reversal fragmented the carry-currency cluster (AUD, NZD, CAD) from the safe-haven bloc (CHF, JPY)

Table 1: Detected G10 FX network fractures (CNSI-Z ≥ 2.5σ). Brexit and COVID are synchronization crises (ΔFiedler > 0); the Japan 2024 carry unwind is a fragmentation crisis (ΔFiedler < 0). Sync. = synchronization; Frag. = fragmentation.

Table 2 reports events that did not trigger the CNSI threshold, with rationale consistent with the model's physical interpretation.

Event	CNSI-Z	Why Not Detected (Correct Negative)
US Election (Nov 2016)	−0.1σ	USD surged but G10 currencies adjusted proportionally — no network fracture
Russia/Ukraine (Feb 2022)	0.9σ	RUB collapsed, but RUB is not a G10 currency; EUR absorbed stress without network-wide propagation
Truss/Yen (Oct 2022)	0.6σ	GBP and JPY moved in opposite directions with limited cross-network spillover
SVB Collapse (Mar 2023)	2.1σ	Banking stress, not FX network stress; CHF safe-haven flows were orderly

Table 2: Non-detected events (correct negatives). These events stressed individual currencies or non-G10 assets without causing a structural G10 FX network fracture. CNSI readings near zero are consistent with the model's physical interpretation.

8.3 CNSI vs. Realized Volatility: The Network Filter

A natural objection is that the CNSI merely relabels high-volatility regimes. We directly test this by constructing a realized-volatility baseline (RVol): the cross-sectional average of $|r_i(t)|$ across all 9 G10 currencies, converted to the same rolling z-score as CNSI-Z using a strictly lagged 252-day window. The full-sample correlation between CNSI-Z and RVol-Z is **0.21** — well below the 0.85 threshold at which the series would be measuring the same phenomenon.

Event	CNSI-Z	RVol-Z	CNSI≥2.5σ	RVol≥2.5σ	Verdict
Brexit (Jun 2016)	15.2	10.3	YES	YES	Both detect
COVID (Mar 2020)	15.8	19.4	YES	YES	Both detect
Japan Carry (Aug 2024)	9.1	2.2	YES	NO	CNSI only — key result
Russia/Ukraine (Feb 2022)	0.9	2.6	NO	YES (FP)	CNSI filters correctly

Event	CNSI-Z	RVol-Z	CNSI≥2.5σ	RVol≥2.5σ	Verdict
SVB Collapse (Mar 2023)	2.0	3.0	NO	YES (FP)	CNSI filters correctly
Truss/Yen (Sep 2022)	0.5	5.7	NO	YES (FP)	CNSI filters correctly
US Election (Nov 2016)	−0.1	1.3	NO	NO	Both quiet

Table 5: CNSI-Z vs. realized-volatility z-score (RVol-Z) at each labeled event. Full-sample correlation = 0.21. Over the full sample: CNSI fires without RVol on 43 days; RVol fires without CNSI on 58 days; both fire together on 10 days. False positive rates: CNSI 1.6% (32/2,055 non-crisis days), RVol 2.1%.

Three results establish that the spectral network structure adds information beyond realized volatility. **First**, the Japan 2024 carry unwind (CNSI-Z = 9.1, RVol-Z = 2.2) is the cleanest separation: the carry unwind reorganized G10 correlation structure without producing an unusually large aggregate volatility reading. The CNSI detects it at high confidence; a volatility monitor would have missed it. **Second**, RVol produces false positives on Russia/Ukraine (2.6σ), SVB (3.0σ), and Truss/Yen (5.7σ) — events where individual currencies or non-FX assets were stressed but the G10 correlation network remained intact. The CNSI correctly suppresses all three. **Third**, the correlation of 0.21 across 2,605 trading days confirms these series are not measuring the same latent variable. The network model acts as a *structural filter*: it fires when cross-currency correlations are reorganized simultaneously and stays quiet when volatility is concentrated in one currency or one asset class.

8.4 False Positive Profile

At the 2.5σ threshold, the CNSI produces 35 total readings above the threshold: 3 true positives (the confirmed fractures) and 32 false positives over approximately 2,055 non-crisis trading days, yielding a false positive rate of 1.6% (32/2,055). The RVol baseline produces 44 false positives (2.1%) on the same non-crisis days. Rather than claiming a calibrated false-positive rate from three true events, we describe the profile descriptively: over 10 years the CNSI produced readings above 9σ on exactly three occasions, all corresponding to confirmed G10 network fractures, and readings between 2.5σ and 9σ on 32 additional days, several of which cluster around documented stress episodes (CHF depeg aftermath January 2015, EM-spillover episodes). The available backtest window is too short to make frequentist power claims; the CNSI's value at this stage is structural: a theoretically grounded stress measure with demonstrably different detection characteristics from a volatility index.

8.5 Fiedler Dynamics by Regime

The sign of $\Delta\delta = \delta(t^*) - \delta(t^* - 60)$ at the CNSI peak date t^* cleanly separates the two crisis regimes:

Event	ΔFiedler	Regime	Interpretation
Brexit (Jun 2016)	+10.7%	Sync.	G10 correlations converged under GBP shock; network tightened
COVID (Mar 2020)	+71.2%	Sync.	Largest Fiedler increase in sample; all currencies moved together

Event	ΔFiedler	Regime	Interpretation
Japan Carry (Aug 2024)	−15.2%	Frag.	JPY reversal partially fragmented the carry-currency cluster

Table 3: Fiedler value change (ΔFiedler) at CNSI peak, measured over the 60-day pre-peak window. Positive = Fiedler rose (synchronization regime); negative = Fiedler fell (fragmentation regime). Brexit and COVID are synchronization crises; the Japan 2024 carry unwind is the sole fragmentation crisis among detected events, consistent with its mechanism: the JPY reversal split the carry-currency cluster (AUD, NZD, CAD) from the safe-haven bloc (CHF, JPY), reducing cross-cluster correlations and hence the Fiedler value.

A critical implication: the common narrative that a declining Fiedler value signals an impending FX crisis is *false for synchronization crises*, which include the two most severe detected events (Brexit 15.2σ, COVID 15.8σ). For synchronization crises, the Fiedler value rises as all correlations converge under shared flight-to-safety pressure — the relevant signal is the magnitude of $\|f(t)\|_2$, not the level of δ. The Fiedler decline criterion holds for the Japan 2024 carry unwind (ΔFiedler = −15.2%), where the JPY reversal fragmented the carry-currency cluster (AUD, NZD, CAD) from the safe-haven bloc, reducing cross-cluster correlations. That event is also the one that RVol missed at 2.2σ — precisely because correlation fragmentation does not require high aggregate volatility.

8.6 On the Absence of Return Prediction

We report explicitly that the Currency Network model does not produce statistically meaningful one-step-ahead displacement forecasts, and that attempts to make it genuinely forward-looking confirm this is structural rather than a calibration failure. The baseline forcing function `f(t) = uὐᵇˢ(t) - mean` uses same-day displacements, making `u*(t) = L⁺f(t)` a contemporaneous transform. Replacing f(t) with `f(t) = uὐᵇˢ(t−1)` (fully observable lagged displacement) produces average OOS $R^2 = -0.20$ across G10 currencies — the lagged spectral transform actively misdirects relative to the mean, performing worse than a naive zero-forecast on every currency except NOK. The AR(1) benchmark estimated on 2015–2019 and applied frozen to 2020–2024 produces $R^2 \approx +0.003$. The model offers no predictive advantage. This is not a calibration failure: `u* = L⁺f` amplifies the projection of f onto low-eigenvalue modes, and yesterday's displacement in those modes has negative autocorrelation at the 1-day horizon — the network topology inverts the signal rather than preserving it. Making the model genuinely predictive would require replacing f with variables that have positive autocorrelation at the forecast horizon, such as interest rate differentials or macro surprise indices, which we identify as the primary avenue for future work.

9. Limitations and Future Work

The Currency Network as formalized here rests on several idealizations that motivate future extensions.

- **Prediction mechanism.** The most important open problem is constructing a genuinely forward-looking forcing function. Replacing f(t) with lagged displacements, interest rate differentials, or macro surprise indices would make u*(t) a true forecast of $uὐ^{bs}(t+1)$. Whether the spectral structure of L adds information beyond a direct AR(1) in that setting is an open empirical question.

- **Linearity.** The spring law is linear (Hookean). Real FX markets exhibit nonlinear dynamics, particularly during crises where correlations spike and liquidity evaporates. Nonlinear spring laws of the form $F = -k|u_{ij}|^\alpha \, \mathrm{sgn}(u_{ij})$ for $\alpha > 1$ are a natural extension requiring numerical equilibrium solvers.
- **Stationarity.** We assume quasi-stationarity within each rolling window. Regime-switching extensions using Hidden Markov Models over the eigenvalue sequence $\delta(t)$ would capture structural breaks more faithfully and could potentially discriminate synchronization from fragmentation in real time.
- **Data quality.** These results use Yahoo Finance daily close prices rather than institutional tick data. For G10 major pairs (EUR, GBP, JPY, CHF, AUD, CAD) the bid-ask spread during normal conditions is 0.5–2.0 pips, representing 0.003–0.015% of notional — negligible relative to the displacement magnitudes that drive CNSI. For the less liquid pairs (NOK, SEK), typical spreads of 5–15 pips (0.05–0.15%) widen to 30–60 pips during stress episodes. At the carry-wave displacements observed in August 2024 (∼1–3% daily moves in NOK/SEK), this represents a measurement noise of at most 2–6% of the displacement signal — material for precision but not for detection at the 9σ CNSI thresholds reported here. Replication with Refinitiv or Bloomberg tick data is a recommended robustness check before any trading application.
- **Network dimension.** Extension to EM currencies requires careful handling of capital controls and market microstructure differences that violate the Pearson correlation assumptions underlying the signed Laplacian.
- **Correlation threshold.** The effect of sparsifying the spring constant matrix (e.g., setting $k_{ij} = 0$ for $|\rho_{ij}| < 0.2$) on CNSI performance has not been tested and is a natural robustness check.

10. Conclusion

We have introduced the Currency Network: a mechanical field theory of foreign exchange equilibrium grounded in Hookean spring mechanics and spectral graph theory. The framework admits closed-form equilibrium solutions ($u^* = L^+ f$), proven existence and uniqueness (Theorem 1), a spectral displacement bound (Theorem 2), and a mass-invariance result (Proposition 1) that clarifies the structural role of currency mass as governing dynamic return speed rather than static displacement magnitude.

The framework's primary empirical contribution is the Currency Network Stress Index (CNSI). Backtested on 2,605 trading days (2015–2024), the CNSI detects three confirmed G10 network fractures at 9–16σ — Brexit (June 2016, 15.2σ), the COVID synchronized selloff (March 2020, 15.8σ), and the Japan carry-trade unwind (August 2024, 9.1σ) — while correctly producing near-zero readings for events that stressed individual currencies or non-G10 assets without fracturing the G10 correlation structure. Critically, a realized-volatility baseline correlates with CNSI-Z at only 0.21 (full-sample), and the Japan 2024 carry unwind (CNSI-Z = 9.1 vs. RVol-Z = 2.2) demonstrates that the spectral network structure detects structural fractures that volatility measures miss. The RVol baseline also false-fires on Russia/Ukraine, SVB, and Truss/Yen while CNSI correctly suppresses all three, producing a lower false-positive rate (1.6% vs. 2.1%). The network model acts as a structural filter, not a volatility relabeling.

A secondary contribution is the formal typology of crisis regimes. Detected events divide into synchronization crises (ΔFiedler > 0, large coordinated shock) and fragmentation crises (ΔFiedler < 0, network splits). This distinction replaces an earlier, incorrect narrative that the Fiedler value universally declines before crises; for the most severe detected events, it rises. The two regimes are identifiable in real time from the sign of the 60-day Fiedler change.

We are explicit that the model does not produce displacement forecasts: u*(t) is a contemporaneous equilibrium map, not a prediction. The CNSI's value is structural diagnosis. The concrete operational use case is real-time regime monitoring: a risk officer tracking live CNSI-Z against a 2.5σ alert threshold would have received actionable signals at Brexit (15.2σ), COVID (15.8σ), and the Japan carry unwind (9.1σ), while remaining quiet through Russia/Ukraine, SVB, and Truss/Yen — episodes that triggered realized-volatility alerts but did not fracture the G10 network structure. The CNSI does not say "sell JPY"; it says "the G10 correlation network is under structural stress of a kind not seen in the prior year." That is a different, complementary signal to volatility, and one that the RVol comparison shows is not recoverable from volatility alone. Making the forcing function genuinely forward-looking is the primary direction for future work.