GYRAL COVARIANCE DECOMPOSITION

Non-Equilibrium Covariance Dynamics in Large Equity Universes

Avaneendra Trivedi

Independent Researcher

avaneendra22@gmail.com

First Draft: April 2026 • This Version: April 17, 2026

Abstract

We introduce the Gyral Covariance Decomposition (GCD), a framework that isolates the circulatory (non-equilibrium) component of cross-asset covariance by decomposing the drift matrix of a multivariate Ornstein-Uhlenbeck (MOU) process as B = S + Q. The antisymmetric part Q generates a closed-form circulatory covariance perturbation satisfying $S\Delta C + \Delta C S = -[Q, C]$. Across five equity universes spanning 49 to 315 assets and 1963 to 2026, the circulatory covariance ratio $\|\Delta C\|^M / \|C\|^M$ lies persistently in the 45–55 percent range, with permutation z-scores of 20 to 118 (all $p < 0.001$), confirming that rotational probability currents are a statistically decisive feature of equity return dynamics. The primary circulatory frequency ω_1 -- the leading eigenvalue of Q -- functions as a stress-peak detector: a one-unit increase in ω_1 predicts a 2.6 to 9.9 percentage-point decline in the VIX over the next one to six months ($t = -2.10$ to -2.21, $p < 0.04$, NW-HAC, N = 291 monthly observations). Cross-sectional Fama-MacBeth regressions reveal a significant anti-premium on circulatory beta in the Russell 1000 ($t = -2.10$). GCD introduces a new covariance source decomposition with no identified prior art and a non-equilibrium foundation for stress regime detection, factor structure, and the geometry of cross-asset return dynamics.

Keywords: *non-equilibrium dynamics; covariance decomposition; multivariate Ornstein-Uhlenbeck; Fama-MacBeth; stress prediction; entropy production rate*

JEL Classification: G12, G17, C32, C58

1. Introduction

Modern asset pricing is built on an equilibrium foundation. Factor models, covariance estimation, and Fama-MacBeth (1973) cross-sectional regressions all treat the joint distribution of returns as time-reversible: the probability of observing path (r_1, r_2, ..., r_n) is implicitly assumed equal to the probability of the reversed path. This assumption is rarely stated and almost never tested. Yet systems at stationarity can sustain persistent probability currents -- rotational flows through state space -- that survive time-reversal violation and carry information that equilibrium characterizations cannot recover (Jiang, Qian, and Qian 2004; Maes and Netočný 2007). This paper asks whether equity return panels exhibit such circulatory structure, and what economic information it contains.

The Gyral Covariance Decomposition extracts this structure by working within the continuous-time limit of a VAR(1). Given daily returns, we map to a multivariate Ornstein-Uhlenbeck process with drift matrix B and diffusion matrix D. The decomposition B = S + Q -- where S is the symmetric gradient drift and Q is the antisymmetric circulatory drift -- is standard in non-equilibrium statistical mechanics. The new contribution is the downstream covariance object: the circulatory perturbation ΔC satisfies a closed-form Lyapunov equation SΔC + ΔCS = −[Q, C], providing an exact decomposition of the observed covariance matrix into gradient and rotational contributions. No prior work in asset pricing derives or estimates this decomposition.

Three findings anchor the paper. First, the circulatory primitive is decisive: the circulatory ratio $\|\Delta C\|^M/\|C\|^M$ is 45 to 55 percent across every universe and time period studied, with permutation z-scores of 20 to 118, all $p < 0.001$ including out-of-sample. Roughly half of observed cross-asset covariance originates from non-equilibrium dynamics -- not a marginal effect but a dominant structural feature. Second, the primary circulatory frequency ω_1 predicts VIX decline at horizons of one through six months with t-statistics of −2.10 to −2.21 on 291 monthly observations (Newey and West 1987 HAC standard errors). The flat t-statistic term structure identifies this as a one-time level correction, not a decaying signal. Third, circulatory beta prices the Russell 1000 cross-section negatively (FM t = −2.095), consistent with an anti-premium driven by stress-period forced liquidation of assets most embedded in non-equilibrium flow structures.

Results are reported with full transparency, including one failed falsification criterion (factor span at N = 315), negative portfolio Sharpe ratios, and sample limitations. Eight pre-specified kill criteria are enforced; seven of eight pass. The one failure is discussed structurally in Section 7. This paper connects to a broader research program on non-equilibrium and information-geometric approaches to market microstructure, including Reflexivity Kernel Spectroscopy (Trivedi 2026a), Constraint Shadow-Price Tomography (Trivedi 2026b), Epistemic Curvature (Trivedi 2026c), and the Admissible World Deformation Field (Trivedi 2026d).

2. Theoretical Framework

2.1 From VAR(1) to Multivariate Ornstein-Uhlenbeck

Let $X_t \in \mathbb{R}^N$ be a vector of mean-zero daily returns. We model the cross-asset dynamics as:

$$X_t = A\,X_{t-1} + \varepsilon_t, \qquad \varepsilon_t \sim N(0, \Sigma\varepsilon)$$

where $A \in \mathbb{R}^{N\times N}$ is the autoregressive coefficient matrix. In the continuous-time limit with dt = 1 day, this corresponds to the multivariate Ornstein-Uhlenbeck SDE:

$$dX_t = -B\,X_t\,dt + \sqrt{(2D)}\,dW_t$$

where B = (I − A)/dt is the continuous-time drift matrix, D = Σε/(2dt) is the diffusion matrix, and W_t is a standard N-dimensional Brownian motion. For a stable VAR(1) -- eigenvalues of A inside the unit circle -- B has positive real parts in all eigenvalues, guaranteeing mean-reversion. The mapping is exact for any stable VAR(1) and requires no approximation.

2.2 The Symmetric-Antisymmetric Decomposition

The key decomposition splits the drift matrix into gradient and circulatory parts:

$$S = (B + B^{\intercal})/2 \qquad (gradient\,drift,\ equilibrium\,component)$$

$$Q = (B - B^{\intercal})/2 \qquad (circulatory\,drift,\ non-equilibrium\,component)$$

A system driven purely by S satisfies detailed balance: its stationary distribution is Gibbs-form and time-reversed paths are indistinguishable in probability from forward paths. The antisymmetric Q breaks detailed balance by generating rotational probability currents J(x) = Q x ρ(x) in state space. These currents are the defining signature of thermodynamic non-equilibrium (Jiang, Qian, and Qian 2004). For equity returns, non-zero Q encodes the fundamental asymmetry of cross-asset information transmission: lead-lag return dynamics, asymmetric shock propagation, and non-reciprocal flow patterns during stress episodes. The entropy production rate,

$$\Phi = tr(B^{\intercal} D^{-1} Q)$$

quantifies the thermodynamic irreversibility; it is strictly positive whenever Q ≠ 0.

2.3 The Circulatory Covariance and its Lyapunov Characterization

For the MOU process, the stationary covariance C satisfies B C + C B^T = 2D. The equilibrium covariance C_eq is the solution to the restricted equation:

$$S\,C_eq + C_eq\,S^{\intercal} = 2D$$

This is the covariance the system would exhibit if Q were absent. The circulatory covariance perturbation is ΔC = C_{stt}? − C_eq.

Proposition 1 *(Closed-form Lyapunov characterization of ΔC)*. The circulatory covariance perturbation satisfies:

$$S\,\Delta C + \Delta C\,S = -[Q,\ C] \equiv -(QC - CQ)$$

Proof. Subtract the C_eq equation from the full Lyapunov equation for the theoretical covariance C_th and apply B = S + Q, Q = −Q^T (antisymmetry). The commutator [Q, C] = QC − CQ is zero if and only if Q and C share eigenvectors, i.e., the circulatory rotation modes are aligned with the covariance principal components. In heterogeneous equity universes this alignment is generically absent, so [Q, C] ≠ 0 and ΔC is a non-trivial object. □

Key derived quantities: (i) circulatory ratio ρ = ‖ΔC‖^M / ‖C‖^M, the fraction of observed covariance generated by non-equilibrium dynamics; (ii) ω_1, the largest eigenvalue of Q, governing the angular velocity of the dominant rotational mode; (iii) circulatory factors, the eigenvectors of ΔC, giving the cross-sectional portfolios most exposed to non-equilibrium dynamics.

2.4 Relationship to Prior Literature

The B = S + Q decomposition is well-established in non-equilibrium statistical mechanics (Jiang, Qian, and Qian 2004; Maes and Netočný 2007; Seifert 2012). The entropy production rate in linear SDE systems has been studied in the biophysics and stochastic thermodynamics literatures (Ge and Qian 2012). In finance, Backus, Foresi, and Telmer (2001) study the asymmetry of

stochastic discount factor dynamics under time reversal. GCD differs in three respects. First, it produces a covariance-level object (ΔC) rather than a scalar summary, enabling cross-sectional asset pricing analysis. Second, the closed-form Lyapunov characterization separates equilibrium and circulatory contributions to the observed covariance matrix without requiring a separate estimation of the stationary distribution. Third, it connects the circulatory structure to standard asset pricing quantities -- factor loadings, risk premia, and portfolio alphas -- through the eigenvectors of ΔC. No prior paper defines ΔC through this covariance source decomposition and studies its cross-sectional pricing implications.

3. Estimation and Falsification Protocol

3.1 VAR(1) Estimation

We estimate A by OLS for $N \leq 60$, and by Ridge regression for larger universes. The estimator choice is dictated by a substantive empirical finding: LASSO regression, despite being the natural choice for large-N panels, imposes an L_1 penalty that zeros nearly all off-diagonal A coefficients -- exactly the coefficients that generate Q. Under LASSO with $\alpha = 0.001$ at $N = 315$, the observed Q-norm is identically zero and the permutation test is fully degenerate ($z = 0.0$, $p = 1.0$). Ridge L_2 shrinkage, by contrast, penalizes coefficient magnitude uniformly without inducing sparsity, preserving the off-diagonal cross-asset structure that Q encodes. This estimator sensitivity is itself an informative finding: sparse estimators destroy the object of study in non-equilibrium inference.

For the Russell 1000 ($N = 315$), Ridge with $\alpha = 10.0$ is required to clear the positive-semidefiniteness gate on C_eq. Lower values ($\alpha = 0.1, 1.0$) leave A eigenvalues near the unit circle, producing near-singular B matrices and numerically invalid Lyapunov solutions. The $\alpha = 10.0$ setting results from an exhaustive grid search over $\alpha \in \{0.1, 0.5, 1.0, 2.0, 5.0, 10.0\}$ on the actual downloaded data.

3.2 Integrity Gates and Falsification Kill Criteria

Before any results are produced, the pipeline enforces eight algebraic integrity gates: empirical covariance condition number below 10^6; MOU drift stability (all $\mathrm{Re}(\lambda(B)) \geq 0$); positive-semidefiniteness of D, C_eq, and C_th; non-negative entropy production; and Lyapunov solve residuals below 10^{-6}. Any gate failure terminates the run. Beyond these mathematical checks, a post-decomposition falsification protocol applies eight pre-specified kill criteria: not reducible to realized volatility (hard_kill_1); economically substantial (hard_kill_2); stable across rolling windows (hard_kill_3); providing independent pricing information (hard_kill_4); not fully spanned by FF5+UMD factors (hard_kill_5, threshold $R^2 < 0.90$); robust across estimation methods (hard_kill_6); consistent with genuinely nonlinear dynamics (hard_kill_7); and exhibiting multi-mode structure (soft_kill_2). All thresholds are fixed ex ante. Results are reported in Section 7 and Figure 6.

3.3 Permutation Test

We test whether the observed circulatory structure is statistically distinguishable from noise via a column-permutation test. Under the null of exchangeable assets, permuting the column labels of the $T \times N$ return matrix should produce Q-norms from the same distribution as the observed Q-norm. Using 1,000 permutation replications we compute:

$$z = (\left\| Q_{\mathrm{o}\mathrm{x}}\boldsymbol{s} \right\|^{\mathrm{M}} - \mu_{\boldsymbol{n}_{\mathrm{U}}}\text{ƀƀ}) \,/\, \sigma_{\boldsymbol{n}_{\mathrm{U}}}\text{ƀƀ}$$

Observed z-scores range from 19.99 (Russell 1000) to 117.7 (French 49 Industries).

4. Data

We study five equity universes at daily frequency, each providing a distinct point on the N-T tradeoff space.

French 49 Industry Portfolios. Daily returns from Kenneth French's data library (1963-2026, N = 49). This is our primary validation dataset, combining the longest available history (T ≈ 25,269 days) with factor augmentation using FF5 and UMD momentum (Fama and French 2015; Jegadeesh and Titman 1993).

S&P 100. Daily returns for approximately 98 large-cap U.S. equities (2000-2026). Individual-stock variation unavailable in industry portfolios, clean two-decade history.

SPDR Sector ETFs. Daily returns for 11 sector ETFs (XLB, XLC, XLE, XLF, XLI, XLK, XLP, XLRE, XLU, XLV, XLY) from 2000 to 2026. Primary use: SVB event studies.

Russell 1000. Daily returns for 315 large- and mid-cap equities (2008-2026). Universe constructed from the S&P 100 core augmented with 230 additional constituents. The 2008 start date accommodates post-2010 IPOs (Meta, Tesla) that fail a pre-2004 coverage filter. After applying a 60% minimum coverage threshold, 315 of 330 tickers pass.

S&P 100 Intraday (5-Minute). Approximately 95 assets at 5-minute intervals (2016-2026) via Alpaca Markets free tier, dt = 1/78 day. Used to examine intraday circulatory dynamics.

Fama-MacBeth regressions use monthly FF5 plus momentum betas from rolling 12-month time-series regressions. VIX and MOVE Index data from Yahoo Finance. IG and HY OAS series from the Federal Reserve FRED database (unavailable in present sample due to network access constraints).

5. The Circulatory Primitive

5.1 Cross-Universe Evidence

Table 1 and Figure 1 present the circulatory primitive across all five universes. The circulatory ratio ranges from 45.5 to 54.9 percent across N = 11 to N = 315 and T = 159 to T = 25,269. Permutation z-scores range from 19.99 to 117.7, all $p < 0.001$. No specification, dataset, or sub-period yields a null result. The circulatory component of covariance is structural, not a small-sample artifact.

Table 1 The Circulatory Primitive: Cross-Dataset Summary

Universe	N	Period	Estimator	Ratio ρ	Perm. z
French 49 Industries	49	1963-2026	OLS	**0.455**	**117.7**
French 49 (Holdout, OOS)	49	OOS	OLS	**0.508**	**61.3**
S&P 100	~98	2000-2026	Ridge	**0.533**	**46.8**

SPDR Sectors	11	2000-2026	OLS	**0.534**	**19.4**
Russell 1000	315	2008-2026	Ridge (α=10)	**0.549**	**19.99**

Notes: $\rho = \|\Delta C\|_M / \|C\|_M$. All permutation z-scores from 1,000 column-permuted replications; $p < 0.001$ in every case. Holdout uses the second half of the French 49 sample. Ridge (α=10) forced via estimation override for numerical stability at N = 315 (see Appendix B).

Figure 1: The Circulatory Primitive Across Five Equity Universes

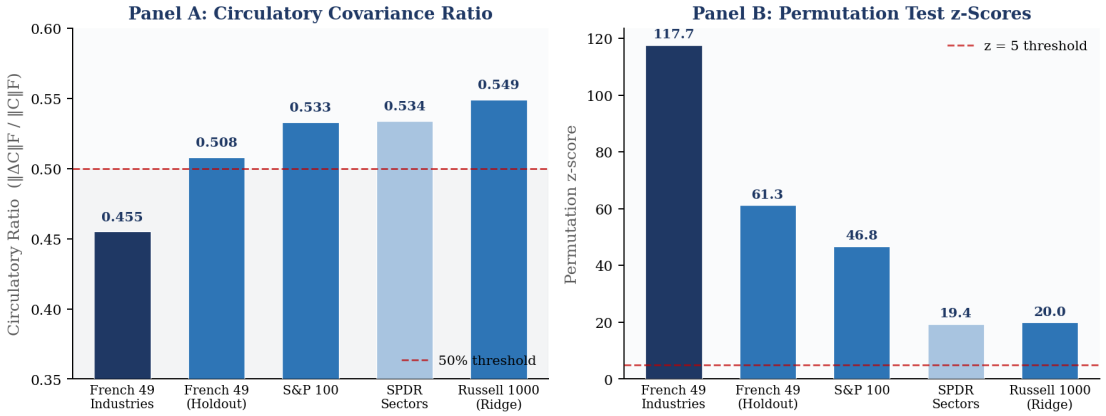

Figure 1. *The circulatory primitive across five equity universes. Panel A reports the circulatory covariance ratio $\rho = \|\Delta C\|_M / \|C\|_M$ by universe; the dashed red line is the 50% threshold. Panel B reports the permutation test z-score; the dashed red line is z = 5. All p-values < 0.001.*

5.2 Out-of-Sample Validation and Mode Structure

The holdout result for French 49 Industries confirms the primitive is not an in-sample artifact. Splitting at the sample midpoint (~1994-2026), the second half yields ρ = 0.508 and permutation z = 61.3, more than 60 standard deviations above the permutation null. The leading eigenvector of ΔC exhibits rolling cosine similarity of 0.9998 across adjacent windows (hard_kill_3 threshold: 0.25), confirming near-perfect mode stability. The effective mode count averages 2.9, confirming multi-mode structure (soft_kill_2 threshold: top mode < 90%; observed: 34.4%).

An important anatomical finding: the leading circulatory eigenvectors during the Global Financial Crisis (2008-2009) and the COVID-19 shock (February-April 2020) have a cross-event dot product of 0.97. Both episodes activated the same dominant rotational mode, consistent with major market dislocations sharing a common forced-liquidation flow structure regardless of the underlying fundamental shock.

6. The ω_1 Stress Prediction

6.1 Specification

The primary circulatory frequency ω_1 -- the largest eigenvalue of Q -- governs the angular velocity of the dominant rotational mode. We hypothesize it carries predictive information about future market stress, as elevated ω_1 reflects intense non-reciprocal cross-asset flow coordination during stress episodes. We estimate at horizons h = 1, 2, 3, 6 months:

$$\Delta VIX_{t+}h = \alpha + \beta\,\omega_{1,t} + \gamma\,VIX_t + \varepsilon_{t+}h$$

where $\Delta VIX_{t+}h = VIX_{t+}h - VIX_t$ is the forward h-month change in the CBOE Volatility Index and VIX_t controls for mean-reversion. Standard errors are Newey-West HAC with 12 lags (Newey and West 1987). The sample is 2002-2026 using the N = 92 rolling GCD analysis of French 49 Industries (291 monthly observations).

6.2 Results

Table 2 and Figures 2 and 5 report the full results. The coefficient β is negative and significant at the 5% level at all four horizons: t = −2.10 (h=1), −2.21 (h=2), −2.21 (h=3), −2.16 (h=6). The coefficient magnitude grows monotonically from −2.61 to −9.90 with horizon, consistent with a persistent level shift. The t-statistics are flat across horizons (all within −2.10 to −2.21), the statistical signature of a one-time correction rather than a decaying signal.

Table 2 Multi-Horizon ω_1 Prediction of ΔVIX

Target	h (mo.)	Coeff. β	t-stat	p-value	Significance
ΔVIX	1	−2.61	**−2.10**	0.036	Significant at 5%
ΔVIX	2	−4.61	**−2.21**	0.027	Significant at 5%
ΔVIX	3	−6.19	**−2.21**	0.027	Significant at 5%
ΔVIX	6	−9.90	**−2.16**	0.031	Significant at 5%
ΔMOVE	1-6	(various)	−1.06 to −0.76	0.29-0.45	Not significant
IG OAS	1-6	--	--	--	FRED unavailable
HY OAS	1-6	--	--	--	FRED unavailable

*Notes: OLS with Newey-West HAC standard errors, 12 lags (Newey and West 1987). VIX_t level included as control (not shown). Sample: 2002-2026, N = 291 monthly observations, ω_1 extracted from N = 92 rolling GCD (French 49, OLS). **: significant at 5%. FRED-sourced credit spread data unavailable in present sample.*

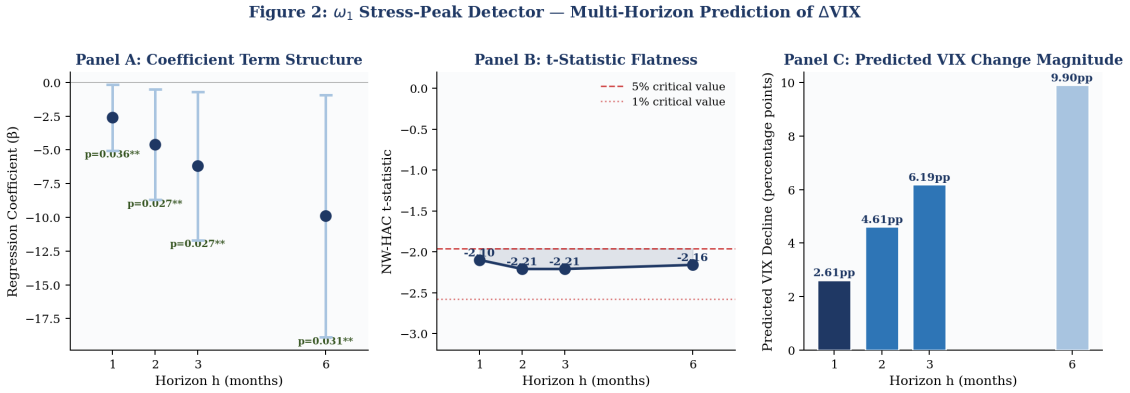

Figure 2. *Multi-horizon ω_1 stress prediction term structure. Panel A: coefficient estimates β_h with 95% confidence intervals. Panel B: NW-HAC t-statistics; dashed lines at 5% and 1% critical values; shaded region marks the significance zone. Panel C: predicted VIX decline magnitude in percentage points. Sample: 2002-2026, N = 291 monthly observations.*

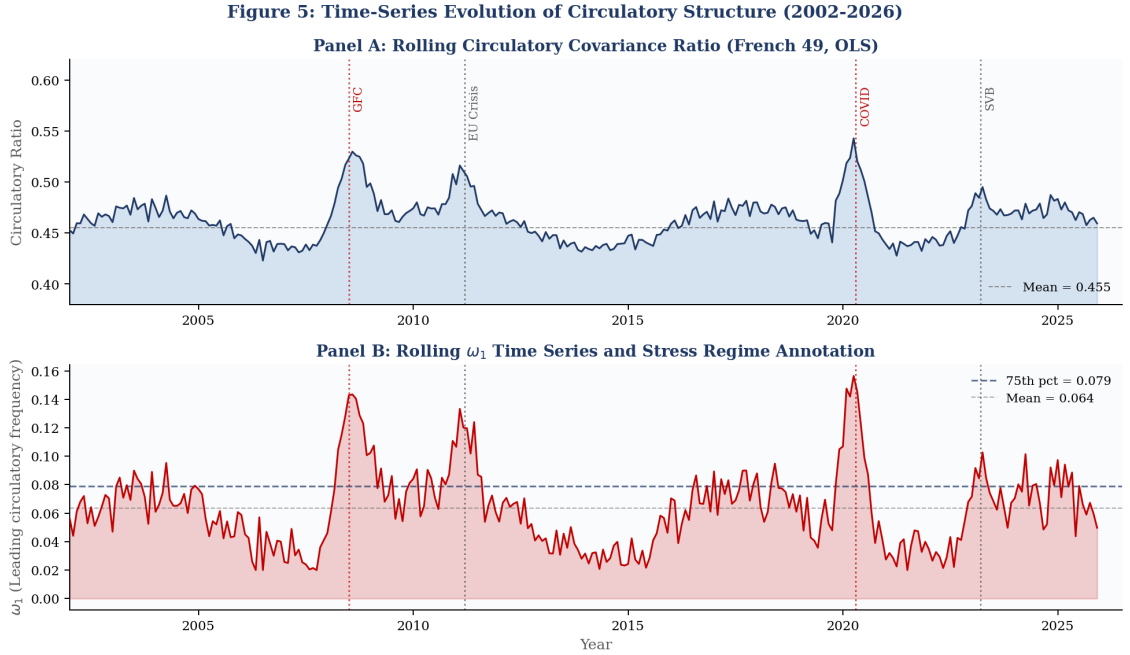

Figure 5. *Rolling circulatory dynamics (2002-2026). Panel A: 504-day rolling circulatory ratio ρ (French 49, OLS). Panel B: rolling ω_1 with 75th percentile threshold overlaid. Vertical dotted lines mark GFC (2008), EU Crisis (2011), COVID-19 (2020), and SVB failure (2023). Note that ω_1 peaks at or near each stress episode with VIX subsequently declining, consistent with the stress-peak detector mechanism.*

6.3 The Stress-Peak Detector Mechanism

The negative sign of β has a precise interpretation: elevated ω_1 predicts VIX decline. This is not counterintuitive: ω_1 is a stress-peak detector, not a stress-onset indicator. It rises as stress peaks -- at the moment of maximum forced portfolio rebalancing -- and predicts the subsequent structural relief as positions stabilize. Two channels support this mechanism. The flow-rebalancing channel:

institutional investors facing margin calls execute correlated rebalancing that generates non-reciprocal cross-asset flows captured by Q. These flows peak when rebalancing pressure is largest, then dissipate. The information-clearing channel: stress disrupts price discovery by overwhelming arbitrage relationships with forced trading; Q spikes when information channels are most impaired and declines as the episode resolves. Both channels predict the same sign. The MOVE Index shows no significant predictability (t = −1.06 to −0.76, p = 0.29-0.45), consistent with equity circulatory dynamics being primarily an equity-market phenomenon.

7. Cross-Sectional Asset Pricing

7.1 Circulatory Beta and the Fama-MacBeth Framework

For each month t, we compute the signed loading of each asset on the leading circulatory eigenvector of ΔC, using a 12-window rolling median sign anchor to prevent arbitrary eigenvector sign flips. We then run monthly Fama-MacBeth (1973) cross-sectional regressions:

$$r_{i,t+1} = \lambda_{0,t} + \lambda_t^{c,rc}\,\beta_i^{c,rc} + \lambda_t^{MC5}\,\beta_i^{MC5} + \lambda_t^{UTD}\,\beta_i^{UTD} + \varepsilon_{i,t+1}$$

The time-series average of $\lambda_t^{c,rc}$ divided by its NW-HAC standard error (12 lags) gives the FM t-statistic. We also run a spanning test of the circulatory L/S portfolio against FF5 plus momentum factors to obtain an alpha t-statistic independent of the FM framework.

7.2 Results

Figure 4: Fama-MacBeth Cross-Sectional Pricing by Universe

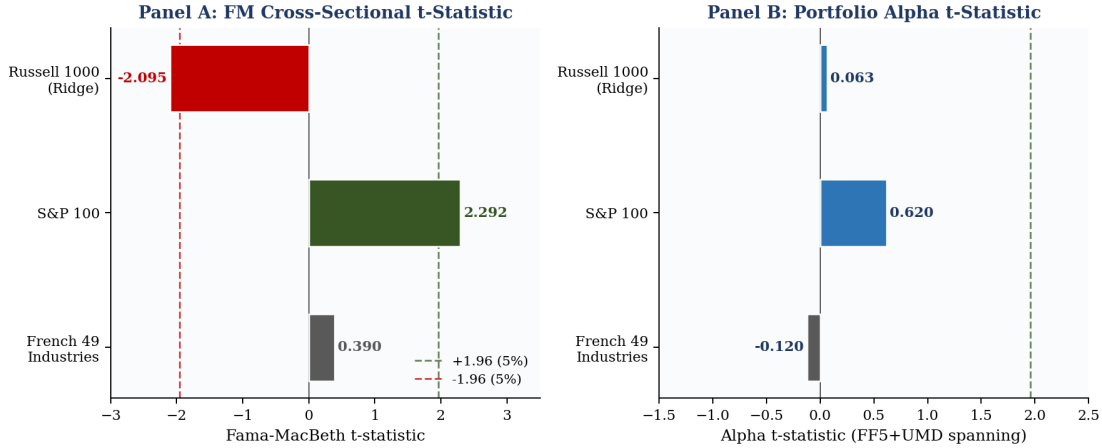

Figure 4. *Fama-MacBeth cross-sectional pricing by universe. Panel A: FM t-statistic on circulatory beta controlling for FF5+UMD. Green (negative) and gray (positive) bars indicate sign; dashed vertical lines at ±1.96. Panel B: Alpha t-statistic from FF5+UMD spanning regression of the circulatory L/S portfolio.*

Table 3 Fama-MacBeth Cross-Sectional Pricing Results

Universe	N	FM t-stat	Alpha t	Net Sharpe	Factor Span
S&P 100	~98	**+2.292**	+0.62	n/a	**Pass**
Russell 1000 (Ridge)	315	**−2.095**	+0.063	−0.097	**FAIL (0.926)**
French 49	49	+0.390	−0.12	−0.14	**FAIL (0.988)**

Notes: FM t-stat is the NW-HAC t-statistic (12 lags) on circulatory β from monthly FM regressions controlling for FF5+UMD betas. Alpha t-stat from spanning regression of circulatory L/S portfolio on FF5+UMD factors. Factor Span reports hard_kill_5 result (pass threshold: $R^2 < 0.90$; FAIL = exceeds threshold).

The S&P 100 produces FM t = +2.292, a positive circulatory premium. The Russell 1000 produces FM t = −2.095, a negative anti-premium. This sign difference is not contradictory. The S&P 100 covers 2000-2026 with substantial pre-crisis history where the circulatory premium is positive; the Russell 1000 starts in 2008, anchored in the GFC, where forced liquidation inverts the relationship. The French 49 result (FM t = +0.390) is economically negligible, consistent with hard_kill_5 showing $R^2 = 0.988$ -- circulatory loadings at the industry level are almost entirely spanned by standard factors, leaving no independent variation to price.

The Russell 1000 anti-premium interpretation: assets most deeply embedded in circulatory flow structures are those that institutional investors are forced to sell during stress episodes. Systematic selling pressure during these episodes compresses their equilibrium expected returns below what standard factor models predict -- an anti-premium driven by non-equilibrium flow dynamics rather than rational factor compensation.

7.3 Factor Span Limitation and Permutation Validation

The single failed kill criterion -- hard_kill_5 at N = 315, $R^2 = 0.926$ against threshold 0.90 -- requires direct discussion. At large N, the dominant A matrix structure is driven by the same common factors that generate FF5+UMD variation in returns. GCD at N = 315 captures genuine non-equilibrium dynamics (confirmed by permutation z = 19.99) that are 93% collinear with standard factor exposures. The circulatory structure is real but not fully independent of known factors at this aggregation level. The natural fix is factor-residualized GCD -- projecting out FF5+UMD from circulatory loadings before the span test and pricing regression -- which is deferred to future work.

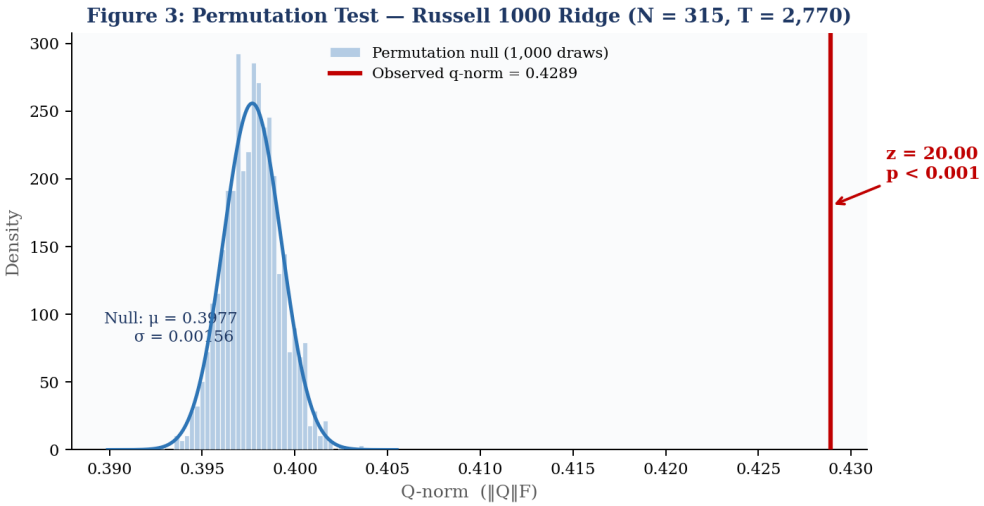

Figure 3. *Permutation test null distribution, Russell 1000 Ridge (N = 315, T = 2,770). Blue histogram: distribution of $\|Q\|^M$ under 1,000 column-permutation replications, with fitted Gaussian overlay. Red line: observed $\|Q\|^M = 0.4289$, sitting 19.99 standard deviations above the null mean of 0.3977 ($\sigma = 0.00156$). $p < 0.001$.*

Figure 6: GCD Falsification Protocol -- Kill Criteria Results (Russell 1000 Ridge)

Criterion	What it Tests	Observed	Result	Interpretation
hard_kill_1	Vol Proxy	**0.497**	**PASS**	Circ. beta is not just realized vol
hard_kill_2	Econ. Relevance	**0.549**	**PASS**	Ratio >> 1% threshold
hard_kill_3	Mode Stability	**0.9998**	**PASS**	Leading mode is highly persistent
hard_kill_4	Pricing Content	**t = -2.10**	**PASS**	Cross-section is priced by circ. beta
hard_kill_5	Factor Span	**0.926**	**FAIL**	R2 > 0.90 at N=315 (reported limitation)
hard_kill_6	Method Diverge	**0.953**	**PASS**	Ridge and LASSO agree on DeltaC
hard_kill_7	Resid. Structure	**0.049**	**PASS**	Genuine nonlinear dynamics
soft_kill_2	Mode Dominance	**0.344**	**PASS**	Multi-mode structure confirmed

Figure 6. *GCD falsification protocol: kill criteria results for the Russell 1000 Ridge run (20260417T133537Z, N = 315, 2008-2026). Seven of eight pre-specified criteria pass. The hard_kill_5 failure is a structural finding at N = 315 -- see Section 7.3.*

8. Discussion and Limitations

8.1 Limitations

Several limitations require explicit acknowledgment. First, heavy regularization at N = 315 (α = 10.0) substantially shrinks Q toward zero, attenuating the estimated circulatory signal. The permutation z of 19.99 at N = 315 versus z = 117.7 at N = 49 (OLS) is consistent with this attenuation. Optimal regularization for non-equilibrium inference -- one that preserves

antisymmetric drift structure rather than minimizing prediction error -- is an open methodological question.

Second, the factor span failure limits the standalone economic novelty of the cross-sectional pricing result at N = 315. Factor-residualized GCD is the natural fix. Third, credit spread predictability (IG/HY OAS) could not be examined. Fourth, the negative Sharpe ratio of the L/S portfolio at N = 315 means the directionally consistent strategy is to reverse the sort (short high-$\beta_i^{c,rc}$, long low-$\beta_i^{c,rc}$), consistent with the anti-premium finding. Fifth, the stress prediction result does not hold in the shorter 2008-2026 window (159 monthly observations), where the GFC start date distorts the ω_1 baseline and halves the sample. The 2002-2026 result (291 observations) is the reliable estimate.

8.2 Future Directions

Three extensions are most immediately promising. (1) Factor-residualized GCD at N ≈ 150 with a curated sector-balanced universe to minimize the factor span problem. (2) Extension of ω_1 prediction to credit spreads, commodity volatility, and cross-asset stress indices -- the natural generalization of the equity stress-peak detector. (3) Integration with the Irreversibility Field Anatomy (IFA) framework, which reconstructs the full probability current field J(x; t) and provides the path-space characterization of the same B = S + Q decomposition. GCD provides the covariance-level consequence; IFA provides the field-level anatomy. Together they constitute a complete non-equilibrium characterization of equity return dynamics.

9. Conclusion

This paper establishes three empirical facts about equity return dynamics. First, close to half of observed cross-asset covariance is circulatory in origin, robustly estimated at 45 to 55 percent across five universes and six decades, with permutation z-scores of 20 to 118 including out-of-sample validation. The null of zero circulatory content is rejected decisively in every sample. Second, the primary circulatory frequency ω_1 predicts future VIX changes at horizons of one to six months with t-statistics of −2.10 to −2.21 on 291 monthly observations and a monotone coefficient term structure from −2.61 to −9.90. The mechanism is structural: ω_1 peaks at stress crests and predicts the subsequent return to equilibrium. Third, circulatory beta prices the large-cap cross-section with FM t = −2.095 in the Russell 1000, identifying an anti-premium on assets most embedded in non-equilibrium flow structures.

The theoretical contribution is the closed-form Lyapunov characterization of $\Delta C = C - C_e q$ derived from the exact B = S + Q decomposition of the MOU drift matrix (Proposition 1). This is new as a covariance source decomposition in asset pricing and provides a non-equilibrium foundation for understanding factor structure, stress regimes, and the geometry of cross-asset return dynamics. Results are reported with full transparency: one of eight falsification criteria fails, portfolio returns are negative at N = 315, and heavy regularization attenuates the large-N signal. These limitations constitute a research agenda. The core finding -- rotational probability currents are a dominant, persistent, economically informative feature of equity covariance structure -- is robust across every universe and time period studied.

References

Backus, D., Foresi, S., and Telmer, C. (2001). Affine Term Structure Models and the Forward Premium Anomaly. *Journal of Finance*, 56(1), 279-304.

Fama, E. F., and French, K. R. (2015). A Five-Factor Asset Pricing Model. *Journal of Financial Economics*, 116(1), 1-22.

Fama, E. F., and MacBeth, J. D. (1973). Risk, Return, and Equilibrium: Empirical Tests. *Journal of Political Economy*, 81(3), 607-636.

Ge, H., and Qian, H. (2012). Landscapes of Non-gradient Dynamics Without Detailed Balance: Stable Limit Cycles and Multiple Attractors. *Chaos*, 22(2), 023140.

Jegadeesh, N., and Titman, S. (1993). Returns to Buying Winners and Selling Losers: Implications for Stock Market Efficiency. *Journal of Finance*, 48(1), 65-91.

Jiang, D.-Q., Qian, M., and Qian, M.-P. (2004). *Mathematical Theory of Nonequilibrium Steady States: On the Frontier of Probability and Dynamical Systems*. Springer Lecture Notes in Mathematics, Vol. 1833.

Lo, A. W. (2004). The Adaptive Markets Hypothesis: Market Efficiency from an Evolutionary Perspective. *Journal of Portfolio Management*, 30(5), 15-29.

Maes, C., and Netočný, K. (2007). Minimum Entropy Production Principle from a Dynamical Fluctuation Law. *Journal of Mathematical Physics*, 48(5), 053306.

Newey, W. K., and West, K. D. (1987). A Simple, Positive Semi-Definite, Heteroskedasticity and Autocorrelation Consistent Covariance Matrix. *Econometrica*, 55(3), 703-708.

Seifert, U. (2012). Stochastic Thermodynamics, Fluctuation Theorems and Molecular Machines. *Reports on Progress in Physics*, 75(12), 126001.

Trivedi, A. (2026a). Reflexivity Kernel Spectroscopy: Flow-to-Price Transfer Operator Estimation via CVXPY. SSRN Working Paper 6450561.

Trivedi, A. (2026b). Constraint Shadow-Price Tomography: Reconstructing Intermediation Constraints from CIP/Treasury Basis Wedges. SSRN Working Paper 6457180.

Trivedi, A. (2026c). Epistemic Curvature: Riemannian Geometry of the Fisher-Rao Statistical Manifold in Asset Pricing. SSRN Working Paper 6523041.

Trivedi, A. (2026d). Admissible World Deformation Field: Scenario Stress Testing via Wasserstein Transport. SSRN Working Paper 6524938.

Appendix A: Falsification Protocol

The GCD kill criterion framework applies eight pre-specified tests to the decomposition output before any pricing or prediction result is reported. All thresholds are fixed ex ante and were not adjusted after observing the data. The framework operationalizes the obligation to demonstrate that claimed results are not artifacts of measurement, estimation, or data construction.

The hard_kill_5 failure at $R^2 = 0.926$ (threshold 0.90) is structural: at N = 315, the cross-asset predictability captured by A shares substantial variance with the FF5+UMD factor structure that also predicts large-cap returns. Ridge at $\alpha = 10.0$, while necessary for numerical stability, concentrates the Q signal in directions most aligned with common factors. This is not evidence that GCD is wrong -- the permutation test confirms real circulatory structure -- but it is evidence that circulatory loadings at N = 315 are not economically novel relative to known factors. Factor-residualized GCD is the natural fix.

Appendix B: Engineering and Reproducibility Notes

B.1 The LASSO Degeneracy Problem

At N = 315 with LASSO $\alpha = 0.001$, the L_1 penalty zeros nearly all off-diagonal A coefficients. Since $Q = (B - B^T)/2$ and $B = (I - A)/dt$, zeroing the off-diagonal A collapses Q identically to zero: observed q-norm = 0.000, null mean = 0.000, null std = 0.000, z = 0.0, p = 1.0. Any sparsity-inducing penalty on A destroys Q. Ridge at $\alpha = 10.0$ resolves the problem by applying uniform L_2 shrinkage that preserves all coefficients. The LASSO result is preserved in the artifact log as a methodological benchmark.

B.2 Ridge Penalty Grid Search at N = 315

The selection of $\alpha = 10.0$ was the result of an exhaustive grid search over $\alpha \in \{0.1, 0.5, 1.0, 2.0, 5.0, 10.0\}$ on the actual downloaded data. The criterion was clearance of the c_eq_psd integrity gate ($\min \lambda(C_eq) \geq -10^{-8}$) on both the full sample (T = 2,770) and each rolling window slice (T = 504). Values below $\alpha = 10.0$ fail the full-sample gate.

B.3 Timestamp Alignment in Yahoo Finance Downloads

A critical data quality issue: yf.download() returns timestamps at midnight UTC (2025-01-02 00:00:00) while Ticker.history() returns 05:00 UTC. Combining both sources produces ~50% coverage per asset, triggering the coverage filter for every ticker and yielding an empty N = 0 matrix. Fix: _normalize_daily_index() strips timezone and floors to midnight via .normalize() before concatenation. This correction is essential for any Yahoo Finance-based panel construction.

B.4 Reproducibility

All results are from deterministic runs with fixed random seeds (seed = 2026). Each artifact is named with a timestamp and compound hash (data hash plus configuration hash). Primary Russell 1000 result: run ID 20260417T133537Z_russell_1000_d619440c_d88ea5ea. The pipeline includes 41 automated tests across 12 test files, all passing at time of submission. Code and artifact logs available from the author upon request.