



Article

Beyond Magnitude: Lacunarity of Cross-Asset Correlation Images as a Structural Measure of Systemic Dependence

Ömer Akgüller ¹, Mehmet Ali Balcı ^{1,*}, Perihan Çetin ¹ and Lucian Gaban ^{2,*}

¹ Department of Mathematics, Faculty of Science, Mugla Sıtkı Kocman University, 48000 Mugla, Turkey; oakguller@mu.edu.tr (Ö.A.); perihancetin@posta.mu.edu.tr (P.Ç.)

² Faculty of Economics, “1 Decembrie 1918” University of Alba Iulia, 510009 Alba Iulia, Romania

* Correspondence: mehmetalibalcı@mu.edu.tr (M.A.B.); lucian.gaban@uab.ro (L.G.)

Abstract

Standard scalar indicators of systemic dependence, such as the mean pairwise correlation, the absorption ratio, and the dispersion of the eigenvalue spectrum, summarise the magnitude of co-movement but are by construction blind to its spatial arrangement. We propose treating the time-varying cross-asset correlation matrix as a greyscale image and quantifying its spatial organisation with the multiscale gliding-box lacunarity. Using a controlled block-factor generative model in which the average correlation is held fixed while the sectoral block strength is varied, we show that lacunarity recovers the planted block structure almost perfectly (partial Spearman $\rho = 0.92$ at fixed mean correlation), a recovery that persists under fat-tailed innovations, time-varying loadings, and overlapping communities, whereas the mean correlation and the absorption ratio remain flat. Applied to twenty years of daily data for sixty-two sector-spanning United States equities, lacunarity tracks a model-free index of block heterogeneity after controlling for correlation magnitude (partial Spearman $\rho = 0.46$, ninety-five percent bootstrap interval $[0.33, 0.58]$) and improves the out-of-sample prediction of block structure beyond the magnitude baselines. We are explicit about two boundaries. A simple permutation-invariant dispersion statistic, the standard deviation of the off-diagonal correlations, tracks block heterogeneity even more strongly than lacunarity, so lacunarity is not the most efficient estimator of that quantity; its distinct role, confirmed by a scrambling test, is that it responds to the spatial arrangement of dependence, which dispersion measures are invariant to, and it remains informative under a canonical clustering or spectral ordering. The measure is descriptive rather than predictive of future drawdowns. The results position correlation-image lacunarity as an interpretable, computationally light, and arrangement-sensitive complement to the existing magnitude and dispersion descriptors of systemic dependence.



Academic Editors: Erveton Pinheiro Pinto and Marcelo Pires

Received: 1 June 2026

Revised: 19 June 2026

Accepted: 21 June 2026

Published: 27 June 2026

Copyright: © 2026 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the [Creative Commons Attribution \(CC BY\)](https://creativecommons.org/licenses/by/4.0/) license.

Keywords: lacunarity; gliding-box algorithm; correlation matrix; systemic risk; spatial heterogeneity; econophysics; image analysis

1. Introduction

Fractal geometry was introduced to describe objects whose irregularity and self-similarity persist across scales, and lacunarity was proposed by Mandelbrot as a complementary descriptor that distinguishes textures sharing the same fractal dimension yet differing in the size and distribution of their gaps [1], whereas the fractal dimension measures how densely a set fills space, lacunarity measures how heterogeneously it does so, i.e., the degree to which the mass of a pattern is clumped rather than evenly spread.

The gliding-box algorithm of Allain and Cloitre [2,3] and its generalisation to quantitative patterns by Plotnick and co-workers [4] turned lacunarity into a practical multiscale statistic that has since been applied to spatial heterogeneity in ecology, geology, medical imaging, and remote sensing [4–7]. The common thread across these applications is that lacunarity reads the spatial organisation of a pattern, which is precisely the dimension of information that magnitude-based summaries discard.

Financial markets are a natural setting for such a structural reading. The cross-correlation matrix of asset returns is the central object of empirical finance and of econophysics [8,9], and its statistical structure has been studied extensively through random matrix theory since the observation that the bulk of its eigenvalue spectrum is consistent with noise [10–12]. The largest eigenvalues escape the Marchenko–Pastur bound and encode a market-wide mode together with sectoral modes, which has motivated community-detection and core–periphery analyses of the dependence structure [13,14]. On the risk-monitoring side, the absorption ratio of Kritzman and co-workers, defined as the fraction of total return variance captured by a small number of leading principal components, has become a widely used indicator of market fragility [15]. Random matrix tools and their machine-learning hybrids continue to be refined for portfolio construction and market-state identification [16–18].

A limitation unites these scalar indicators. The mean pairwise correlation, the absorption ratio, and the dispersion of the eigenvalue spectrum all respond to the same underlying quantity, namely the overall strength of co-movement. When a broad shock arrives, correlations rise together, and every one of these measures climbs in tandem. None of them, however, distinguishes a homogeneous market in which all assets move alike from a blocky market in which dependence is concentrated within sectors. That distinction is a property of the spatial arrangement of the correlation matrix, not of its average level, and it is exactly the kind of property that lacunarity was designed to capture.

Two strands of recent work make this connection plausible. First, encoding a one-dimensional series as a two-dimensional image, through Gramian angular fields or Markov transition fields, has proven effective for time-series classification and forecasting, including in finance [19–22], which shows that the image representation of financial dynamics carries usable information. Second, lacunarity has very recently been promoted from a static texture statistic to a multiscale recurrence quantifier capable of detecting dynamical regime transitions in noisy and short time series, where it succeeds precisely because it reads geometric heterogeneity rather than line structure [23–25]. In parallel, geometry-inspired network measures such as the Ollivier–Ricci and Forman–Ricci curvatures of correlation networks have been advanced as indicators of market fragility, underlining the value of structural rather than magnitude-based descriptors [26–28]. Multifractal detrended fluctuation analysis has likewise established that financial series carry scale-dependent structure beyond second-order statistics [29–32].

The contribution of this paper is to bring lacunarity to bear directly on the spatial structure of cross-asset dependence. We treat the rolling correlation matrix as a greyscale image and characterise it with the differential box-counting fractal dimension and the multiscale gliding-box lacunarity. Our central claim is methodological and is stated and tested as such: lacunarity of the correlation image reads the block, which is a sectoral organisation of dependence as a dimension distinct from its magnitude. We establish this first on a controlled generative model in which the average correlation is held fixed while the block strength is varied, providing a clean ground-truth validation, and then on two decades of equity data. We are deliberate in separating what the measure does, namely structural characterisation, from what it does not do, namely forecasting, and we report a negative forecasting result openly so that the scope of the contribution is not overstated.

2. Background and Related Work

2.1. Cross-Asset Dependence, Contagion, and Risk Transmission

The dependence among asset returns is the empirical foundation of portfolio theory and of systemic-risk analysis. Markowitz [33] established that the covariance structure, not the marginal variances alone, governs the risk of a portfolio, so that the off-diagonal organisation of the correlation matrix is economically primary. A large literature documents that this structure is unstable and tends to intensify during crises. Forbes and Rigobon [34] showed that apparent increases in cross-market correlation during turmoil must be interpreted with care once heteroskedasticity is accounted for, while subsequent work established that genuine contagion and shifts in dependence nonetheless occur, frequently expressed as a collapse of sectoral or regional distinctions into a single common mode. The dynamics of this co-movement have been modelled through multivariate volatility models such as the dynamic conditional correlation specification of Engle [35], which track the time variation of the dependence matrix that our rolling-window construction also targets.

2.2. Connectedness and Systemic-Risk Measures

A parallel literature builds scalar measures of systemic risk and interconnectedness. The absorption ratio of Kritzman et al. [15] quantifies fragility as the fraction of variance captured by a few leading principal components. Diebold and Yilmaz [36] construct connectedness indices from forecast-error-variance decompositions of a vector autoregression, and Billio et al. [37] develop Granger-causality and principal-component measures of connectedness among financial institutions. Institution-level tail-risk measures such as the conditional value-at-risk of Adrian and Brunnermeier [38] and the marginal expected shortfall and SRISK of Acharya et al. [39] capture the contribution of individual firms to system-wide distress, and the broader landscape of such analytics is surveyed by Bisias et al. [40]. These measures are predominantly either magnitude summaries of the dependence matrix or institution-level tail statistics, and, as we discuss, most are coincident descriptions of the prevailing state rather than forecasts. None of them reads the spatial arrangement of the dependence matrix, which is the gap this paper addresses.

2.3. Correlation Networks, Diversification, and Structural Descriptors

The correlation matrix has long been studied as a network and through random matrix theory. Mantegna [41] introduced hierarchical structure in correlation-based asset networks, and the eigenvalue analyses of Laloux et al. [10] and Plerou et al. [11] separated a market mode and sectoral modes from noise. On the portfolio side, the diversification ratio of Choueifaty and Coignard [42] and the effective number of bets of Meucci [43] formalise the intuition that the benefit of diversification depends on the structure, not merely the average level, of dependence; this is precisely the quantity that a structural descriptor should illuminate. More recently, geometry-inspired measures such as the Ollivier–Ricci and Forman–Ricci curvatures of correlation networks have been proposed as indicators of market fragility [26–28], and lacunarity has been promoted to a multiscale recurrence quantifier for regime transitions [23]. Our contribution sits within this structural strand: we read the spatial organisation of the correlation matrix directly, as an image, with a fractal descriptor, complementing the magnitude, spectral, network curvature, and dispersion descriptors rather than replacing them.

3. Materials and Methods

3.1. Data and Returns

We analyse daily adjusted closing prices $P_{i,t}$ of $N = 62$ large, continuously listed United States equities spanning ten economic sectors over the period 3 January 2005 to 30 April 2025, yielding 5114 trading days. Logarithmic returns are

$$r_{i,t} = \ln\left(\frac{P_{i,t}}{P_{i,t-1}}\right), \quad i = 1, \dots, N. \quad (1)$$

The cross-section spans information technology, financials, energy, health care, consumer staples, consumer discretionary, industrials, materials, utilities, and telecommunications, so that a genuine block structure exists against which a structural measure can be assessed (Table 1). The constituents are, by sector, the following. Information technology: MSFT, AAPL, IBM, ORCL, CSCO, INTC, TXN, QCOM, and HPQ. Financials: JPM, BAC, WFC, C, GS, AXP, USB, and PNC. Health care: JNJ, PFE, MRK, ABT, BMY, LLY, AMGN, and MDT. Industrials: GE, BA, CAT, MMM, HON, UPS, UNP, and LMT. Consumer staples: PG, KO, PEP, WMT, CL, MO, and COST. Energy: XOM, CVX, COP, SLB, OXY, and HAL. Consumer discretionary: HD, MCD, NKE, LOW, SBUX, and TGT. Utilities: DUK, SO, D, EXC, and AEP. Materials: NEM, FCX, and APD. Telecommunications: T and VZ. All series are continuously listed over the sample period and were retained only if their coverage exceeded ninety-eight percent of trading days, after which the panel was intersected to common dates.

Table 1. Sector composition of the 62-asset United States equity cross-section.

Sector	Number of Assets
Information technology	9
Financials	8
Health care	8
Industrials	8
Consumer staples	7
Energy	6
Consumer discretionary	6
Utilities	5
Materials	3
Telecommunications	2
Total	62

The universe is constructed to provide a stable, liquid, and sectorally balanced cross-section over a long horizon rather than to maximise coverage at a single date. We use the ten-sector Global Industry Classification Standard taxonomy in its pre-2018 form, which predates the separation of real estate and the reorganisation of communication services and therefore gives a clean and temporally consistent partition across the whole 2005–2025 window. Within each sector we select large, continuously listed constituents with near-complete price histories, which is what permits a fixed 62-asset panel over five thousand trading days without survivorship gaps. The sector counts are deliberately uneven, reflecting the differing breadth of the sectors among large continuously listed names; materials and telecommunications contribute few such names over the full horizon, whereas information technology and financials contribute more. We retain the financial sector despite its known idiosyncrasies, namely its elevated internal correlation and its central role in transmission, precisely because a structural measure of dependence should be exposed to the most strongly clustered block in the market; we verify in Section 4.7 that

the structural signal is not an artefact of any single sector by recovering it under label-free orderings. Table 2 lists the constituents with their issuer names.

Table 2. Constituents of the cross-section by sector, with ticker and issuer name.

Sector	Constituents (Ticker: Issuer)
Information technology	MSFT: Microsoft; AAPL: Apple; IBM: IBM; ORCL: Oracle; CSCO: Cisco; INTC: Intel; TXN: Texas Instruments; QCOM: Qualcomm; HPQ: HP
Financials	JPM: JPMorgan Chase; BAC: Bank of America; WFC: Wells Fargo; C: Citigroup; GS: Goldman Sachs; AXP: American Express; USB: U.S. Bancorp; PNC: PNC
Health care	JNJ: Johnson & Johnson; PFE: Pfizer; MRK: Merck; ABT: Abbott; BMY: Bristol-Myers Squibb; LLY: Eli Lilly; AMGN: Amgen; MDT: Medtronic
Industrials	GE: General Electric; BA: Boeing; CAT: Caterpillar; MMM: 3M; HON: Honeywell; UPS: United Parcel Service; UNP: Union Pacific; LMT: Lockheed Martin
Consumer staples	PG: Procter & Gamble; KO: Coca-Cola; PEP: PepsiCo; WMT: Walmart; CL: Colgate-Palmolive; MO: Altria; COST: Costco
Energy	XOM: ExxonMobil; CVX: Chevron; COP: ConocoPhillips; SLB: Schlumberger; OXY: Occidental; HAL: Halliburton
Consumer discretionary	HD: Home Depot; MCD: McDonald's; NKE: Nike; LOW: Lowe's; SBUX: Starbucks; TGT: Target
Utilities	DUK: Duke Energy; SO: Southern Company; D: Dominion; EXC: Exelon; AEP: American Electric Power
Materials	NEM: Newmont; FCX: Freeport-McMoRan; APD: Air Products
Telecommunications	T: AT&T; VZ: Verizon

3.2. Rolling Correlation Image

For a window of length $W = 60$ trading days ending at time τ we form the Pearson correlation matrix $C^{(\tau)} \in [-1, 1]^{N \times N}$ with entries

$$C_{ij}^{(\tau)} = \frac{\sum_{t=\tau-W+1}^{\tau} (r_{i,t} - \bar{r}_i)(r_{j,t} - \bar{r}_j)}{\sqrt{\sum_{t=\tau-W+1}^{\tau} (r_{i,t} - \bar{r}_i)^2} \sqrt{\sum_{t=\tau-W+1}^{\tau} (r_{j,t} - \bar{r}_j)^2}}, \quad (2)$$

and slide the window in steps of five days, producing 1011 matrices. The rows and columns are placed in a fixed sector-grouped order π , so that assets belonging to the same sector occupy contiguous indices and any block structure appears as on-diagonal blocks of the image. The matrix is rendered as a greyscale image by the affine map

$$g_{ij}^{(\tau)} = \frac{1}{2}(C_{ij}^{(\tau)} + 1) \in [0, 1]. \quad (3)$$

This affine map preserves the sign of the correlations; because strong negative co-movements are structurally as informative as strong positive ones, we also consider the magnitude map $g_{ij} = |C_{ij}|$, which we find sharpens the structural readout (Section 4.3). The sector ordering requires a partition that a structural measure should arguably discover rather than assume, so we treat the ordering itself as a design choice and assess, in addition to the fixed sector ordering, a data-driven average-linkage hierarchical-clustering ordering of the correlation-distance matrix $d_{ij} = \sqrt{2(1 - C_{ij})}$ and a canonical spectral (Fiedler) ordering obtained by sorting the assets along the second eigenvector of the graph Laplacian of g . The spectral ordering is not arbitrary: the Fiedler vector x is the minimiser of the weighted arrangement objective $\sum_{ij} g_{ij}(x_i - x_j)^2$ subject to $\|x\| = 1$ and $x \perp \mathbf{1}$, so it places strongly correlated assets at nearby coordinates and is the one-dimensional embedding that makes contiguous blocks as coherent as possible. It is uniquely determined, up to an overall reflection that does not affect lacunarity, whenever the algebraic connectivity of

the graph is simple, which holds generically for the empirical matrices. Their stability is examined in Section 4.7.

3.3. Differential Box-Counting Fractal Dimension

For completeness we also evaluated the fractal dimension of the intensity surface g through the differential box-counting method of Sarkar and Chaudhuri [6], in which the image is partitioned into grids of side s , a column of boxes of size $s \times s \times h$ with $h = \lceil Gs/M \rceil$ is stacked along the intensity axis over each grid cell, and the dimension is the scaling exponent of the total box count $N(s)$ against $1/s$. As reported in Section 4.3, this descriptor carries essentially no information for the correlation image, and we therefore do not use it; the explanation is that the greyscale surface of a dense correlation matrix possesses no fractal scaling in its intensity, only a texture, that is a gap structure, which lacunarity rather than the box-counting dimension is designed to read.

3.4. Multiscale Gliding-Box Lacunarity

Lacunarity is computed with the gliding-box algorithm of Allain and Cloitre [2] in the greyscale form of Dong [5]. A box of side r glides over the image in unit steps; at each of the $(M - r + 1)^2$ positions the box mass is the sum of the enclosed intensities,

$$S_r = \sum_{(u,v) \in \text{box}} g_{uv}. \quad (4)$$

Let $Z^{(1)}(r) = \mathbb{E}[S_r]$ and $Z^{(2)}(r) = \mathbb{E}[S_r^2]$ be the first two moments of the box-mass distribution over all positions. The lacunarity at scale r is

$$\Lambda(r) = \frac{Z^{(2)}(r)}{[Z^{(1)}(r)]^2} = 1 + \frac{\text{Var}(S_r)}{\mathbb{E}[S_r]^2}. \quad (5)$$

Equation (5) makes the interpretation explicit. A perfectly homogeneous image, in which every box of a given size carries the same mass, has zero variance and hence $\Lambda(r) = 1$; the more the local mass varies across positions, the larger $\Lambda(r)$. We summarise the lacunarity curve by its mean over the analysed scales

$$\bar{\Lambda} = \frac{1}{|\mathcal{R}|} \sum_{r \in \mathcal{R}} \Lambda(r), \quad \mathcal{R} = \{2, 3, 4, 6, 8, 12, 16\}, \quad (6)$$

which serves as the scalar structural descriptor used throughout. Figure 1 illustrates the pipeline.




Figure 1. (a) The rolling correlation matrix of sector-ordered returns, with each entry mapped to the intensity $g = (C + 1)/2 \in [0, 1]$; for visibility it is shown with a diverging red–blue colour map (red: high correlation; blue: low), and the colour bar reports the mapped value. (b) The same matrix as the greyscale image on which a gliding box of side r samples every position, recording the enclosed intensity mass; the three coloured squares (red, blue, green) mark three representative positions of the gliding box. (c) The resulting multiscale lacunarity curves $\Lambda(r)$ for a calm market (green, 2006) and for the 2008 crisis (red); the crisis curve lies lower, reflecting a more homogeneous image.

3.5. Baseline Measures

We compare lacunarity against two families of baselines. The first family summarises the magnitude of co-movement. The mean off-diagonal correlation is

$$\bar{\rho} = \frac{1}{N(N-1)} \sum_{i \neq j} C_{ij}. \quad (7)$$

The absorption ratio of Kritzman et al. [15] is the share of total variance carried by the leading eigenvalues,

$$\text{AR}_k = \frac{\sum_{m=1}^k \lambda_m}{\sum_{m=1}^N \lambda_m}, \quad \lambda_1 \geq \lambda_2 \geq \dots \geq \lambda_N, \quad (8)$$

with $k = 1$ in the reported results. The eigenvalue dispersion is $\text{ED} = \sigma_\lambda / \bar{\lambda}$. We further include three spectral descriptors that quantify the concentration of the spectrum: the participation ratio $\text{PR} = (\sum_m \lambda_m)^2 / \sum_m \lambda_m^2$, the effective rank $\exp(-\sum_m p_m \ln p_m)$ with $p_m = \lambda_m / \sum_{m'} \lambda_{m'}$, and the number of eigenvalues exceeding the Marchenko–Pastur edge $\lambda_+ = (1 + \sqrt{N/W})^2$ [12].

The second family captures the dispersion of the correlation values without reference to any ordering and provides the strong, simple structural baselines that a fractal descriptor must be measured against. These are the standard deviation of the off-diagonal correlations, σ_ρ , and their Gini coefficient. Both are permutation-invariant: they depend only on the distribution of the correlation values, not on their arrangement in the matrix.

3.6. A Model-Free Index of Block Structure

To provide a ground-truth target that does not itself depend on any image construction, we define the block heterogeneity of a correlation matrix using the known sector partition $\{S_1, \dots, S_K\}$ as the gap between the average within-sector and average between-sector correlation

$$H_B = \frac{1}{|\mathcal{W}|} \sum_{(i,j) \in \mathcal{W}} C_{ij} - \frac{1}{|\mathcal{B}|} \sum_{(i,j) \in \mathcal{B}} C_{ij}, \quad (9)$$

where $\mathcal{W}$ and $\mathcal{B}$ are the sets of within-sector and between-sector off-diagonal pairs. A large H_B signals a blocky, sector-differentiated dependence structure, and a near-zero H_B a homogeneous one. The methodological question of the paper is whether lacunarity, which uses no partition labels in its computation, recovers H_B beyond what the magnitude baselines explain.

3.7. Block-Factor Generative Model

For a controlled validation we simulate returns from a one-market-factor plus sector-factor model

$$r_{i,t} = \sqrt{a} F_t^m + \sqrt{b} F_t^{s(i)} + \sqrt{1-a-b} \varepsilon_{i,t}, \quad (10)$$

where F^m is a common market factor, $F^{s(i)}$ is the factor of the sector to which asset i belongs, and all factors and the idiosyncratic terms $\varepsilon_{i,t}$ are independent standard normals. Under (10) the population correlation between two assets is $a + b$ if they share a sector and a otherwise, so the block heterogeneity is exactly b and the mean off-diagonal correlation is

$$\bar{\rho} = a + b \varphi, \quad \varphi = \frac{\sum_s n_s (n_s - 1)}{N(N-1)}, \quad (11)$$

with n_s the size of sector s . Holding the target mean correlation at a constant ρ^* while increasing the block strength is achieved by setting $a = \rho^* - b\varphi$. This lets us vary structure

at fixed magnitude, isolating exactly the dimension that distinguishes lacunarity from the scalar baselines.

3.8. Statistical Assessment

Associations are quantified with the Spearman rank correlation ρ . To remove the influence of magnitude, we use the partial Spearman correlation, obtained by ranking the variables, regressing the ranks of two targets on the ranks of a control, and correlating the residuals. The incremental explanatory value of lacunarity is assessed by ordinary least squares, comparing the coefficient of determination of a regression of H_B on the scalar baselines with that of the same regression augmented by lacunarity, and by the t -statistic of the lacunarity coefficient. All real-data associations are computed on the full panel of 1011 windows.

3.9. Computational Pipeline

Algorithm 1 summarises the computation applied to each rolling window. The cost per window is dominated by the gliding-box convolutions, each of which is evaluated in $O(M^2)$ time through a summed-area (integral-image) implementation of the uniform filter [44], so that the full lacunarity curve over $|\mathcal{R}|$ scales costs $O(|\mathcal{R}| M^2)$, while the correlation matrix itself is $O(N^2 W)$. With $N = 62$, $W = 60$, and a five-day step, the entire twenty-year panel of 1011 windows is processed in a few seconds on a single core, so the measure is inexpensive enough to recompute routinely over long panels.

Algorithm 1 Correlation-image lacunarity for one window

Require: return matrix $R \in \mathbb{R}^{W \times N}$, ordering π , box scales $\mathcal{R}$

- 1: $C \leftarrow \text{corr}(R)$ ▷ $N \times N$ Pearson correlation
- 2: $C \leftarrow C[\pi, \pi]$ ▷ apply row/column ordering
- 3: $g \leftarrow (C + 1)/2$ ▷ map to greyscale image in $[0, 1]$
- 4: **for** $r \in \mathcal{R}$ **do**
- 5: $S_r \leftarrow \text{boxsum}(g, r)$ ▷ gliding-box masses via integral image
- 6: $\Lambda(r) \leftarrow \mathbb{E}[S_r^2] / \mathbb{E}[S_r]^2$
- 7: **end for**
- 8: $\bar{\Lambda} \leftarrow \text{mean}_r \Lambda(r)$, $\beta \leftarrow \text{slope}\{\log \Lambda(r) \text{ vs } \log r\}$
- 9: **return** $\bar{\Lambda}, \beta$

4. Results

4.1. Synthetic Validation: Lacunarity Isolates Structure at Fixed Magnitude

Table 3 reports the block-factor experiment, in which the block strength b is increased from 0 to 0.30 while the target mean correlation is held at $\rho^* = 0.30$. The mean correlation and the absorption ratio stays essentially constant across the whole range, with Spearman correlations against b of only 0.05, confirming that the magnitude baselines are insensitive to structure by construction. The block heterogeneity rises linearly with b , as it must, and the lacunarity rises monotonically from 1.085 to 1.327, with a Spearman correlation against b of 0.94. Crucially, after partialling out the mean correlation, the association between lacunarity and block strength is essentially unchanged, with a partial Spearman correlation of 0.94. Figure 2 visualises the mechanism: the correlation image is featureless when $b = 0$ and displays pronounced on-diagonal blocks when $b = 0.30$, yet the two images share the same mean correlation, and only lacunarity registers the difference.

The clean Gaussian factor model is deliberately simple, so we verified that the recovery is not an artefact of its idealised structure by repeating the experiment under a battery of harder generative processes drawn from the financial network and factor literature (Table 4). These cover heavy tails through Student- t innovations with four degrees of freedom;

non-stationarity through loadings that ramp across the window; ambiguous community membership through overlapping communities, in which every asset loads on a secondary randomly chosen sector factor at half strength; heterogeneous factor exposures through per-asset loadings drawn from a continuous distribution; dynamic network reconfiguration, in which the community membership is reshuffled half-way through the window; and a stochastic block model augmented with random cross-sector links. The partial Spearman correlation between lacunarity and the realised block heterogeneity, computed on the $|C|$ image and controlling for mean correlation, remains high across the whole battery: 0.90 under heavy tails, 0.92 under both time-varying and heterogeneous loadings, 0.91 for the augmented block model, and 0.79 even under mid-window reconfiguration. It falls to a still-substantial 0.61 only when the community structure itself becomes genuinely ambiguous through overlap, which is the expected and interpretable failure mode and which partly accounts for the gap to the weaker real-data association reported below, since real sectors overlap. We further confirmed that lacunarity is sensitive to non-block topologies: under a core-periphery generative model, it tracks the size of the densely correlated core at a partial correlation of magnitude 0.51 controlling for mean correlation, with the negative sign reflecting that a larger core renders the image more uniform.

Table 3. Block-factor validation. As the block strength b increases at a fixed target mean correlation $\rho^* = 0.30$, the magnitude baselines stay flat while lacunarity rises monotonically. Values are averages over 30 replications.

b	Mean Corr. $\bar{\rho}$	Absorption AR_1	Block Het. H_B	Lacunarity $\bar{\Lambda}$
0.00	0.300	0.312	0.000	1.085
0.05	0.303	0.314	0.050	1.119
0.10	0.303	0.314	0.100	1.152
0.15	0.302	0.314	0.149	1.188
0.20	0.305	0.316	0.198	1.226
0.25	0.301	0.313	0.250	1.273
0.30	0.303	0.314	0.298	1.327
Spearman ρ vs. b	0.050	0.052	0.990	0.941

Table 4. Recovery under harder generative models. Partial Spearman correlation between lacunarity (on the $|C|$ image) and realised block heterogeneity, controlling for mean correlation, pooled over structure strengths and 30 replications each. The core-periphery row reports the partial correlation with the core fraction.

Generative Model	$\rho(\bar{\Lambda}, \cdot \bar{\rho})$
Gaussian factor (baseline)	+0.92
Student- $t(4)$ innovations	+0.90
Time-varying loadings	+0.92
Heterogeneous per-asset loadings	+0.93
Stochastic block model with cross-links	+0.91
Dynamic mid-window reconfiguration	+0.79
Overlapping communities	+0.61
Core-periphery (vs. core fraction)	−0.51




Figure 2. Synthetic validation of the structural readout. (a,b) Correlation images with no block structure ($b = 0$) and strong blocks ($b = 0.30$), shown with the red–blue colour map of Figure 1; both share the same mean correlation. (c) Across the experiment the mean correlation (red) and the absorption ratio (purple) remain flat while lacunarity (blue) increases monotonically with block strength.

4.2. Crises Homogenise the Correlation Image

Figure 3 shows representative correlation images. In a calm market, the image is heterogeneous, with visible sector blocks and a mean correlation near 0.25. During the global financial crisis (GFC) of 2008 and the COVID-19 shock of March 2020, the image saturates towards uniform high correlation, the mean rising to 0.62 and 0.71, respectively, and the sectoral blocks are washed out. This is the visual signature of homogenisation: a broad shock raises the magnitude of co-movement while flattening its structure. Figure 4 traces the two quantities over the full sample. The standardised mean correlation spikes upward in every shaded crisis interval, whereas the standardised lacunarity tends to fall in the broad systemic episodes of 2008 and 2020, the two series being negatively associated overall.




Figure 3. Representative correlation images. The calm market shows visible sector blocks and substantial heterogeneity; the 2008 and 2020 crises saturate towards uniform high correlation, washing out the block structure. The mean correlation $\bar{\rho}$ rises across the three states.




Figure 4. Standardised mean correlation and correlation-image lacunarity over 2005–2025. Shaded intervals mark broad systemic episodes (grey: 2008, 2020) and more sectoral stress periods (orange). Mean correlation rises in every episode, while lacunarity falls during the broad homogenising crises.

4.3. Structural Information Beyond Magnitude

On the real panel the magnitude baselines are again mutually almost interchangeable, the mean correlation, the absorption ratio, and the eigenvalue dispersion having pairwise Spearman correlations above 0.99. Lacunarity is distinct from all of them, correlating negatively with the mean correlation (-0.52) and positively with block heterogeneity ($+0.64$), and the association with block heterogeneity survives control for magnitude, with a partial Spearman correlation of 0.46 (ninety-five percent bootstrap interval $[0.33, 0.58]$) that is essentially unchanged when the crisis decile is removed (0.45). It also adds beyond the full spectral battery: controlling jointly for the mean correlation, the absorption ratio, the participation ratio, and the effective rank, the partial correlation is still 0.34. Out of sample, under year-blocked cross-validation, adding lacunarity to the magnitude baselines raises the predictive R^2 for block heterogeneity from 0.24 to 0.35, so the increment is not an in-sample overfitting artefact. We note in passing that the differential box-counting dimension of the image, by contrast, is uninformative, with absolute correlations below 0.08 against every other quantity, for the reason given in the Methods; it is not pursued further. The choice of greyscale map matters: replacing the affine map by the magnitude map $g = |C|$ raises the partial correlation from 0.46 to 0.73, since it stops assigning near-zero correlations to mid-grey and so renders the block structure more sharply.

This structural content must, however, be judged against the simplest available structural statistics, the dispersion of the correlation values themselves. Here, we are explicit about a boundary of the method. The standard deviation and the Gini coefficient of the off-diagonal correlations track block heterogeneity considerably more strongly than lacunarity does, with partial Spearman correlations of 0.83 and 0.82 against 0.46 (Table 5, Figure 5a). Controlling additionally for the standard deviation, the partial correlation of lacunarity with block heterogeneity falls to 0.23, whereas the standard deviation retains 0.77 controlling for lacunarity, and out of sample lacunarity adds only 0.02 to a model that already contains the standard deviation. For the narrow task of estimating the label-defined block heterogeneity, therefore, the simple dispersion statistic is the more efficient descriptor, and we state this plainly rather than claim a predictive advantage that the data do not support.




Figure 5. Lacunarity against simpler structural baselines. **(a)** Partial Spearman correlation with block heterogeneity, controlling for mean correlation, for each descriptor; the red bar marks lacunarity (the proposed measure) and the grey bars the baselines. The dispersion statistics (standard deviation and Gini of the off-diagonal correlations) exceed lacunarity, while the spectral concentration measures (participation ratio, effective rank) carry the opposite sign. Here “# eig > MP” denotes the number of eigenvalues exceeding the Marchenko–Pastur edge. **(b)** Scrambling test: lacunarity tracks block heterogeneity under the sector ordering ($\rho = 0.64$) but not when the same correlation values are randomly arranged ($\rho = 0.05$), whereas the dispersion statistics are unchanged by scrambling. Lacunarity thus measures arrangement, not value dispersion.

Table 5. Comparison of structural descriptors on the real panel ($n = 1011$). **(Top):** partial Spearman correlation with block heterogeneity H_B , controlling for mean correlation. **(Middle):** incremental partial correlations. **(Bottom):** out-of-sample R^2 for H_B under year-blocked cross-validation. A simple correlation-dispersion statistic dominates lacunarity for H_B recovery, but is permutation-invariant and so blind to arrangement (see scrambling test, Figure 5b). Here “# eigenvalues $> \lambda_+$ ” is the number of eigenvalues exceeding the Marchenko–Pastur edge λ_+ .

Partial Spearman with H_B (Control: $\bar{\rho}$)	Value
std. dev. of off-diagonal correlations σ_ρ	+0.83
Gini of off-diagonal correlations	+0.82
# eigenvalues $> \lambda_+$	+0.57
lacunarity $\bar{\Lambda}$ (on $ C $)	+0.73
lacunarity $\bar{\Lambda}$ (on $(C+1)/2$)	+0.46
absorption ratio AR_1	+0.40
Incremental Partial Spearman	Value
$\rho(\bar{\Lambda}, H_B \mid \bar{\rho}, \sigma_\rho)$	+0.23
$\rho(\sigma_\rho, H_B \mid \bar{\rho}, \bar{\Lambda})$	+0.77
$\rho(\bar{\Lambda}, H_B \mid \bar{\rho}, AR_1, PR, \text{eff. rank})$	+0.34
Out-of-Sample R^2 for H_B (Year-Blocked CV)	Value
magnitude baselines only	0.24
+lacunarity	0.35
+ σ_ρ	0.63
+ σ_ρ + lacunarity	0.65

The two descriptors are nonetheless not equivalent because they measure different things. The dispersion statistics are permutation-invariant: they depend only on the distribution of the correlation values and are blind to how those values are arranged in the matrix. Lacunarity reads the spatial arrangement. The distinction is made concrete by a scrambling test (Figure 5b). Under the sector ordering, lacunarity tracks block heterogeneity at Spearman 0.64; when the same correlation values are randomly permuted across asset pairs, which leaves the dispersion statistics exactly unchanged, the association collapses to 0.05. Lacunarity therefore responds specifically to whether strong correlations are organised into contiguous blocks, a property that value dispersion cannot represent. Its role is complementary to, not a replacement for, the dispersion measures: dispersion reports how spread the correlations are, and lacunarity reports whether that spread is spatially structured. Figure 6 shows the underlying relationship that at any fixed magnitude, lacunarity rises with block heterogeneity.




Figure 6. Lacunarity against structure and magnitude. **(a)** At any fixed level of mean correlation (colour), lacunarity increases with block heterogeneity. **(b)** The dependence of lacunarity on mean correlation, with block heterogeneity shown in colour, illustrating that lacunarity separates points of equal magnitude according to their structure.

4.4. Scope: A Coincident Diagnostic, Not a Forecast

Because structural measures are sometimes advanced as early-warning signals, we tested whether correlation-image lacunarity predicts forward market stress, defined as the worst cumulative loss of the equally weighted portfolio over the subsequent twenty trading days. Neither the mean correlation nor lacunarity carries appreciable predictive content for this target, the rank correlations with the forward loss being 0.09 and -0.10 , respectively, the partial correlation of lacunarity after controlling for the mean correlation being -0.06 , and the incremental regression R^2 from adding lacunarity being statistically indistinguishable from zero. We report this openly and draw the appropriate conclusion about scope. Lacunarity is a coincident structural diagnostic, not a forecast. This places it in the same category as most established structural indicators, including the absorption ratio and the correlation-based market-state classifications, which are themselves coincident descriptions of the prevailing dependence regime rather than predictors of returns [14,15]. The value of a coincident structural diagnostic is real but specific: it characterises, in real time and in a single interpretable number, whether the current co-movement is structurally homogeneous or sectorally organised, which is information relevant to conditional diversification, to the interpretation of a developing shock as broad or sectoral, and to the construction of regime-conditioned rules. What we do not claim, and what the data do not support, is that the measure anticipates future stress.

4.5. Episode-Level Homogenisation

Table 6 aggregates the measures over eight labelled market episodes and quantifies the homogenisation mechanism that Figures 3 and 4 display visually. The two broad systemic episodes, the global financial crisis of 2008 and the COVID-19 shock of 2020, together with the European sovereign-debt stress of 2011, combine the highest mean correlations of the sample, between 0.57 and 0.68, with the lowest block heterogeneities, between 0.09 and 0.14, and correspondingly the lowest lacunarities, between 1.013 and 1.020. In contrast, the volatility episode of late 2018 and the rate-driven decline of 2022, both of which were more sectorally differentiated than market-wide, retain high block heterogeneity near 0.26 and the highest lacunarities of the sample, near 1.039, despite only moderate mean correlations. The calm periods sit in between. The pattern is consistent and interpretable: episodes that compress the cross-sectional structure into a single common mode raise the magnitude measures while depressing lacunarity, whereas episodes that preserve sectoral differentiation leave lacunarity elevated. The mean correlation and absorption ratio, which move almost identically across all episodes, cannot make this distinction.

Table 6. Measures averaged over labelled market episodes (window $W = 60$, sector ordering). Broad systemic crises (GFC, Euro, COVID) combine high mean correlation with low block heterogeneity and low lacunarity; sectorally differentiated stress (2018, 2022) keeps block heterogeneity and lacunarity high.

Episode	n	$\bar{\rho}$	AR_1	H_B	$\bar{\Lambda}$
Calm 2005–2007	101	0.264	0.300	0.159	1.031
GFC 2008–2009	29	0.565	0.587	0.141	1.020
Euro 2011	17	0.684	0.698	0.091	1.013
China/oil 2015–2016	25	0.498	0.525	0.182	1.026
Volatility 2018	13	0.276	0.328	0.260	1.036
COVID-19 2020	10	0.647	0.666	0.115	1.019
Rates 2022	38	0.323	0.362	0.260	1.039
Calm 2023–2025	92	0.202	0.248	0.230	1.036

4.6. Window-Length Dependence

Table 7 recomputes the core statistics for window lengths from forty to one hundred and sixty trading days. The partial Spearman correlation between lacunarity and block heterogeneity, controlling for mean correlation, is strongest at the shortest windows, 0.53 at forty days and 0.46 at sixty days and decays smoothly to near zero at one hundred and sixty days. We initially read this as the blurring of transient block structure by long averaging windows, but a control experiment shows the interpretation must be more careful. When the same window sweep is applied to a synthetic series with a fixed, stationary block structure, the partial correlation also decays with window length, from 0.42 at forty days to 0.07 at one hundred and sixty days. Since the true structure is constant in that control, the decay there cannot reflect a loss of real structure; it reflects the fact that, as the window lengthens, the correlation estimate concentrates and the window-to-window variation that the rank association exploits shrinks towards estimation noise. The honest conclusion is therefore twofold. The window-length decay is partly a generic statistical property of the estimator rather than a purely economic phenomenon, and lacunarity is by construction a short-horizon probe of the instantaneous correlation geometry rather than a descriptor that converges as the estimate improves. Encouragingly, the real-data association at $W = 60$, namely 0.46, is of the same order as the 0.42 produced by a genuine stationary block model at the same window, so the real signal is comparable in strength to a true planted-block structure. We adopt $W = 60$ and present the window dependence as an intrinsic property of the measure rather than as a robustness guarantee.

Table 7. Window-length dependence (sector ordering, five-day step), with a stationary-block synthetic control. The control decays despite a constant true structure, indicating that part of the decline is a statistical property of the estimator.

W (Days)	Windows	$\rho(\bar{\Lambda}, \bar{\rho})$	$\rho(\bar{\Lambda}, H_B \bar{\rho})$	Control $\rho(\bar{\Lambda}, H_B \bar{\rho})$
40	1015	−0.50	+0.53	+0.42
60	1011	−0.52	+0.46	+0.44
90	1005	−0.58	+0.31	+0.40
120	999	−0.59	+0.20	+0.22
160	991	−0.67	+0.05	+0.07

4.7. Robustness to Matrix Ordering

Because the lacunarity of a matrix image depends on how its rows and columns are arranged, the ordering deserves scrutiny, and we addressed it in three ways (Table 8). The central methodological point is that the measure does not require the exogenous sector partition at all. The canonical spectral or Fiedler ordering, which minimises the well-defined arrangement objective $\sum_{ij} g_{ij}(x_i - x_j)^2$ and is uniquely determined up to reflection when the algebraic connectivity is simple, recovers the signal without any labels; on the magnitude image $|C|$ it attains a partial correlation of 0.55. A data-driven hierarchical-clustering ordering, also label-free, attains 0.72 on $|C|$, essentially matching the 0.73 of the label-based sector ordering, and the clustering result is stable across the single-, complete-, average-, and Ward-linkage rules, whose lacunarity series have a mean pairwise correlation of 0.91. The exogenous sector ordering is thus an interpretable special case rather than a requirement: principled, label-free orderings deliver the same structural readout. A random ordering, by contrast, scatters correlated assets so that genuine blocks no longer form contiguous regions, and the association collapses and changes sign to −0.24. We do not claim permutation-invariance, which an arrangement-sensitive measure cannot and should not possess; the appropriate standard is that the descriptor be well defined relative to a canonical ordering and stable across principled choices, which it is, while only the

deliberate destruction of spatial coherence removes the signal, confirming that lacunarity reads genuine block topology rather than an order-independent matrix statistic.

Table 8. Robustness to row/column ordering of the correlation image ($W = 60$), reported as the partial Spearman correlation of lacunarity with block heterogeneity controlling for mean correlation. Label-free orderings (canonical spectral, clustering) recover the signal, on the $|C|$ image essentially matching the label-based sector ordering; the random ordering is a negative control. The four clustering linkages yield lacunarity series with mean pairwise correlation 0.91.

Ordering	On $(C+1)/2$	On $ C $
Sector (fixed, label-based)	+0.46	+0.73
Spectral/Fiedler (canonical, label-free)	+0.42	+0.55
Clustering, average linkage (label-free)	+0.41	+0.72
Clustering, complete linkage	+0.45	—
Clustering, Ward linkage	+0.42	—
Clustering, single linkage	+0.35	—
Random (negative control)	−0.24	—

4.8. Multiscale Structure of the Readout

The single summary $\bar{\Lambda}$ aggregates a lacunarity curve $\Lambda(r)$ whose shape carries additional information. Across the sample the calm-period curves lie uniformly above the crisis-period curves at every scale (Figure 7a), confirming that homogenisation suppresses heterogeneity at all resolutions rather than at a single scale. The block signal is, however, concentrated at fine scales: the small-scale lacunarity $\Lambda(2)$ correlates with block heterogeneity at Spearman 0.57, more strongly than the large-scale value at 0.40, and the log-log slope β of the lacunarity curve, a scale-free shape descriptor, correlates negatively with block heterogeneity at -0.57 and at -0.37 after controlling for magnitude (Figure 7b). A flatter curve, that is a less negative slope, accompanies a blockier, more heterogeneous matrix. These multiscale descriptors are consistent with, and reinforce, the conclusion drawn from the mean lacunarity.




Figure 7. Multiscale structure. (a) Mean lacunarity curves $\Lambda(r)$ for calm (low mean correlation) and crisis (high mean correlation) windows; calm windows are more heterogeneous at all scales. (b) The lacunarity slope β against block heterogeneity, coloured by mean correlation.

4.9. Significance and Robustness

To confirm that the structural association is not an artefact of the temporal overlap between adjacent windows, we recomputed the central partial Spearman correlation under a moving-block bootstrap with two thousand resamples. The point estimate of 0.46 is accompanied by a ninety-five percent confidence interval of $[0.33, 0.58]$, which excludes zero by a wide margin, and the interval is insensitive to the block length, the endpoints varying by less than 0.04 as the block length is changed across five, ten, twenty, and forty windows. The result is likewise insensitive to the gliding-box scale set, the partial correlation being

0.46, 0.49, and 0.44 for the scales $\{2, 3, 4, 6, 8, 12, 16\}$ used throughout, all integer scales from two to sixteen, and the dyadic scales $\{2, 4, 8, 16\}$, respectively. Finally, the choice of greyscale map affects only the strength, not the existence, of the association, the partial correlation being 0.46 for the affine map $(C + 1)/2$, 0.60 for the squared map C^2 , and 0.73 for the magnitude map $|C|$, the last being preferable because it does not assign near-zero correlations to mid-grey. The structural content of lacunarity beyond magnitude is therefore statistically robust and stable across these implementation choices.

4.10. Economic Interpretation and Actionability

The preceding results establish that lacunarity measures the spatial arrangement of dependence; this subsection makes the economic content of that arrangement explicit and shows where it is decision-relevant. The key distinction lacunarity draws is between two markets that share the same average correlation but differ in structure. A low-lacunarity market is one in which co-movement is uniform, every asset moving alike, so that the cross-section has collapsed onto a single common mode and diversification across names offers little protection. A high-lacunarity market, at the same average correlation, is one in which strong correlations remain concentrated within sectoral blocks while between-sector correlations stay lower, so that cross-sector diversification continues to reduce risk. This is precisely the quantity that the diversification literature identifies as the driver of portfolio benefit, which depends on the structure rather than the level of dependence [33,42,43].

This interpretation is borne out quantitatively, and it is sharpest exactly when it matters, during stress. Restricting attention to the highest-correlation quintile of windows, the periods a risk manager would most worry about, the partial Spearman correlation between lacunarity and block heterogeneity rises to 0.62. Splitting those stress windows at the median lacunarity, the low-lacunarity half has a within-minus-between correlation gap of only 0.11, whereas the high-lacunarity half retains a gap of 0.18. Translated into a portfolio quantity, we measure the ratio of the variance of a representative single-sector holding to that of the fully diversified equal-weight portfolio, a diversification-benefit ratio that exceeds one when spreading across sectors reduces variance. Across all windows, this ratio tracks lacunarity at a partial Spearman correlation of 0.63 controlling for mean correlation, and within the stress quintile it averages 1.25 for low-lacunarity windows against 1.48 for high-lacunarity windows. In other words, at the same elevated average correlation, cross-sector diversification removes roughly a quarter of the representative single-sector variance when lacunarity is low but closer to a third when lacunarity is high. The actionable reading for a portfolio or risk manager is therefore direct: a fall in lacunarity during a high-correlation episode signals that the usual sectoral diversification is failing and that hedging or de-risking is required, whereas a high lacunarity at the same average correlation signals that cross-sector allocation still works. For a regulator or macroprudential analyst, lacunarity offers a single coincident read on whether a stress episode is broad and systemic, with the sectoral structure dissolved or sectorally contained, a distinction that the absorption ratio and the mean correlation, which move almost identically in both cases, cannot make.

The same point underlies the conceptual answer to the question of what lacunarity buys over a dispersion statistic. The two can be made to disagree exactly when arrangement carries information that values do not. Taking the most block-structured window in the sample and randomly permuting the asset layout leaves the standard deviation of the off-diagonal correlations exactly unchanged at 0.25, yet lacunarity falls from 1.26 to 1.08, almost the full width of its empirical range. A market whose correlations are organised into coherent sectoral blocks and a market whose identical correlation values are scattered without structure are indistinguishable to any dispersion statistic but are read

differently by lacunarity, and they have different diversification implications. Lacunarity is the appropriate tool whenever that spatial organisation, rather than the spread of correlation values alone, is what one needs to measure.

5. Discussion

The findings support a single, carefully delimited claim. The lacunarity of the cross-asset correlation image is a structural descriptor that captures the spatial organisation, specifically the sectoral blockiness, of dependence as a dimension distinct from its magnitude. The synthetic experiment establishes this cleanly, since the mean correlation and absorption ratio are held fixed by construction while lacunarity tracks the planted block strength almost perfectly, and this recovery persists under heavy tails, time-varying loadings, and, with attenuation, overlapping communities. The real-data analysis confirms that the same distinctness from magnitude holds in observed markets. We are, however, deliberately precise about what lacunarity is and is not better at. For the specific task of estimating how strongly the market is organised into sector blocks, a simple permutation-invariant statistic, the dispersion of the off-diagonal correlations, is the more efficient descriptor and largely subsumes lacunarity. What lacunarity uniquely contributes is sensitivity to arrangement: the dispersion statistics depend only on the multiset of correlation values and are invariant to how those values are placed in the matrix, whereas lacunarity responds to whether the strong correlations are spatially contiguous, as the scrambling test demonstrates. The honest positioning is therefore not that lacunarity replaces the magnitude or dispersion measures, but that it adds an arrangement-sensitive axis that neither of them possesses, and that this axis is interpretable, multiscale, and inexpensive.

The standard scalar indicators of systemic dependence are mutually redundant: the mean correlation, the absorption ratio, and the eigenvalue dispersion respond to the same magnitude factor and rise together in every broad shock. Lacunarity reports instead whether that co-movement is homogeneous or blocky. The episode-level evidence makes the distinction concrete, since the broad crises of 2008 and 2020 raise magnitude while depressing lacunarity through homogenisation, whereas the sectorally driven stress of 2018 and 2022 leaves lacunarity elevated, so that magnitude and lacunarity together separate market-wide from sector-specific stress in a way that neither does alone.

This structural reading connects naturally to two recent literature. The promotion of lacunarity to a multiscale recurrence quantifier for regime transitions rests on the same property exploited here, namely that lacunarity quantifies geometric heterogeneity rather than a level or a line structure [23]. The use of discrete curvatures of correlation networks as fragility indicators is likewise a structural rather than magnitude-based approach [26,27]; correlation-image lacunarity offers a closely related but computationally lighter and label-free alternative since it reads block organisation directly from the image geometry without constructing a network or specifying a threshold. The image representation of financial dynamics that underlies Gramian-angular and Markov-transition encodings [19,21,22] is here applied to the dependence matrix rather than to a single series, which is what gives the two-dimensional object its genuine structural content.

The homogenisation signature documented here admits a natural reading in terms of market integration. In quiescent periods the correlation matrix retains a rich sectoral organisation, with assets cohering most strongly within their economic groups, and lacunarity is correspondingly high. As a systemic shock propagates, a single common mode comes to dominate the return variation, the distinction between sectors dissolves, and the matrix approaches a rank-one structure in which every pair of assets is similarly correlated. The mean correlation and the absorption ratio both climb in this regime, but they report only that co-movement has intensified; lacunarity additionally reports that the co-movement has

become structurally uniform, collapsing the heterogeneity of the image. This is precisely the transition from a differentiated to a unified market that the financial states and fragility literature describe through correlation eigenstructure [14,15], recovered here as a change in the spatial texture of a single image. The episode-level evidence sharpens the point, since the sectorally driven stress of 2018 and 2022 raises lacunarity while leaving the magnitude measures comparatively muted, the opposite of the broad-crisis pattern, so that lacunarity and the magnitude measures together separate market-wide from sector-specific stress in a way that neither does alone.

The robustness analyses delimit the conditions under which the structural readout is informative. The window-length results show that the signal is a short- to medium-horizon phenomenon, strongest near sixty days and decaying at long windows. A stationary-block control experiment clarifies that this decay is partly a generic property of the rank estimator, since the window-to-window variation it exploits shrinks as the correlation estimate concentrates, rather than being wholly an economic blurring of transient structure. The honest reading is that lacunarity is by construction a probe of the instantaneous correlation geometry at a chosen horizon, not a quantity that converges as the estimate improves. The ordering results are more reassuring. A canonical spectral ordering and several clustering linkages all recover the signal with mutually consistent values, so the descriptor can be made label-free and is not at the mercy of an unstable clustering choice, while the collapse under random ordering verifies that the measure responds to genuine spatial block topology. We do not claim permutation-invariance; an arrangement-sensitive measure cannot be permutation-invariant without discarding the very information it is meant to capture, and the appropriate standard is stability across principled orderings, which the data support.

The position of the measure relative to random matrix theory is complementary rather than competitive. Spectral methods filter the correlation matrix into market, sector, and noise components and quantify the strength of each through eigenvalues [10,11,13]; they are the natural tools when the object of interest is the magnitude and rank of the common factors. Lacunarity instead reads the spatial arrangement of the filtered structure once a layout is fixed and is therefore well suited to questions about how concentrated or dispersed the dependence is across the cross-section. The two views can be combined; for example, by applying the image descriptors to the sector-mode component isolated by spectral filtering, which we leave for future work. In the same spirit, the discrete-curvature indicators of market fragility [26–28] and the recurrence-lacunarity quantifiers of regime transitions [23] are structural rather than magnitude-based, and the correlation-image construction offers a particularly economical member of this family, requiring neither a network nor a threshold nor an embedding, only an ordering and a gliding box.

It is worth stating plainly how the measure relates to the established systemic-risk toolkit, since our use of the language of systemic dependence could otherwise be misread. We position correlation-image lacunarity as a structural image-analysis descriptor of the dependence matrix, not as a competitor to, or a forecaster within, the systemic-risk literature. The institution-level tail measures, conditional value-at-risk, marginal expected shortfall, and SRISK [38,39] answer a different question, namely each firm's contribution to system-wide losses, and require balance-sheet and market-capitalisation inputs that our return-only construction does not use. The connectedness indices of Diebold and Yilmaz and of Billio et al. [36,37] and the absorption ratio [15] are closer in spirit, being functions of the dependence structure, but they too are predominantly coincident summaries, and they capture the magnitude and directionality of linkages rather than the spatial concentration of the correlation field. Lacunarity is complementary to all of these: it adds a single, model-free, coincident reading of whether the dependence is spatially organised or homogeneous,

the dimension that the diversification result of Section 4.10 shows to be economically consequential. We make no claim that it forecasts distress, and we have reported the null forecasting result openly so that the contribution is positioned as a structural diagnostic rather than an early-warning indicator.

Several limitations bound the contribution and motivate future work. First, the lacunarity of a matrix image depends on the ordering of its rows and columns. Our ordering analysis shows that this reliance can be removed, since a canonical spectral ordering and data-driven clustering orderings recover the signal stably without sector labels; nonetheless the descriptor is not permutation-invariant, and a fully canonical, provably stable formulation, perhaps based on a seriation that minimises a bandwidth functional, remains an open methodological problem. Second, and most consequentially for practice, a simple permutation-invariant dispersion statistic, the standard deviation of the off-diagonal correlations, recovers block heterogeneity more efficiently than lacunarity and largely subsumes it for that particular task; lacunarity's defensible contribution is therefore not a better block-heterogeneity estimator but its sensitivity to spatial arrangement, which dispersion lacks, and a reader whose sole aim is to quantify block heterogeneity should prefer the dispersion statistic. The absolute lacunarity values are also compressed near unity because the correlation image is dense, so it is the relative ordering across windows rather than the raw level that carries the information. Third, the measure is descriptive rather than predictive; the negative forecasting result is a genuine boundary, not a tuning failure, so the measure is a coincident diagnostic of dependence structure rather than an early-warning signal. Fourth, the analysis is confined to a single developed equity cross-section of sixty-two assets; the resolution of the correlation image, and hence the range of usable box scales, grows with the size of the universe, so larger panels would permit a finer multiscale analysis. Extending the study to emerging-market panels and to cross-country dependence, where regional block structure and contagion are salient, would test the generality of the homogenisation signature. Further methodological directions include a multifractal generalisation of the lacunarity spectrum in the spirit of multifractal detrended fluctuation analysis [29,31,32], the application of the image descriptors to the spectrally filtered sector mode rather than the raw matrix, and the use of alternative dependence kernels such as rank or tail-dependence matrices in place of the Pearson correlation.

6. Conclusions

We introduced the lacunarity of the cross-asset correlation image as a structural measure of systemic dependence and validated it against a controlled block-factor model, where it recovers the planted sectoral structure at fixed correlation magnitude and continues to do so under heavy tails, time-varying loadings, and overlapping communities. On two decades of equity data the measure reads the block heterogeneity of dependence beyond the mutually redundant magnitude-based indicators, improves the out-of-sample prediction of that structure, and traces the homogenisation of the correlation matrix during broad crises. We have been explicit about the boundaries that the analysis reveals. A simple dispersion statistic recovers block heterogeneity more efficiently than lacunarity, so the distinctive contribution of lacunarity is not estimation accuracy but its sensitivity to the spatial arrangement of dependence, a property that permutation-invariant statistics cannot provide and that a scrambling test isolates cleanly. The measure is a coincident structural diagnostic rather than a forecast, and it is well defined only relative to an ordering, for which canonical spectral and clustering choices give stable results. Within these bounds, the measure is interpretable, computationally inexpensive, and complementary to existing magnitude, dispersion, spectral, and network-geometric tools. By bringing the lacunarity construction from spatial pattern analysis to bear on the geometry of financial dependence,

the paper adds an arrangement-sensitive axis to the structural description of systemic dependence and illustrates the broader usefulness of fractal image descriptors for matrix-valued financial data.

The practical implications follow from the economic reading of the measure. For portfolio construction and risk management, lacunarity provides a single, real-time indicator of whether sectoral diversification remains effective: at a given average correlation, a high value signals that cross-sector allocation still reduces variance materially, while a low value warns that co-movement has become uniform and that diversification has lost its protective power, so that explicit hedging or de-risking is warranted. This distinction is most valuable precisely during high-uncertainty episodes driven by major shocks, when average correlations are elevated and conventional magnitude indicators saturate and can no longer discriminate between a broad systemic collapse and a sectorally contained stress. For macroprudential surveillance, the same coincident reading helps characterise the nature of an unfolding episode. We emphasise, however, that the indicator is diagnostic rather than predictive and should be used to interpret the contemporaneous state of dependence, not to forecast returns or losses.

Several limitations frame these conclusions and motivate further work. The descriptor is defined relative to an ordering and is not permutation-invariant; although canonical spectral and clustering orderings recover the signal stably, a provably unique seriation remains open. A simple dispersion statistic estimates block heterogeneity more efficiently, so lacunarity should be used for its distinct arrangement sensitivity rather than as a block-heterogeneity estimator. The analysis is confined to a single developed-market equity cross-section, and the measure is coincident rather than predictive. Future research could extend the construction to emerging-market and cross-country panels, where regional contagion and reconfiguring blocks are salient; apply the image descriptors to the spectrally filtered sector mode; substitute rank or tail-dependence kernels for the Pearson correlation to capture asymmetric and extreme co-movements relevant to hedging; develop a multifractal generalisation of the lacunarity spectrum; and integrate the measure with connectedness and network-curvature indicators into a combined structural panel for systemic surveillance.

Author Contributions: Conceptualization, Ö.A. and M.A.B.; methodology, Ö.A. and M.A.B.; software, Ö.A. and M.A.B.; validation, M.A.B., P.Ç. and L.G.; formal analysis, Ö.A. and M.A.B.; investigation, Ö.A., M.A.B. and P.Ç.; writing—original draft preparation, Ö.A. and M.A.B.; writing—review and editing, P.Ç. and L.G.; visualization, Ö.A.; supervision, M.A.B. and L.G.; project administration, L.G. All authors have read and agreed to the published version of the manuscript.

Funding: This research received no external funding.

Data Availability Statement: The daily price data analysed in this study are publicly available from standard market-data providers. The analysis code is available from the corresponding authors on reasonable request.

Conflicts of Interest: The authors declare no conflicts of interest.

References

1. Mandelbrot, B.B. *The Fractal Geometry of Nature*; W.H. Freeman: San Francisco, CA, USA, 1982.
2. Allain, C.; Cloitre, M. Characterizing the lacunarity of random and deterministic fractal sets. *Phys. Rev. A* **1991**, *44*, 3552–3558.
3. Cheng, Q. The gliding box method for multifractal modeling. *Comput. Geosci.* **1999**, *25*, 1073–1079.
4. Plotnick, R.E.; Gardner, R.H.; Hargrove, W.W.; Prestegard, K.; Perlmutter, M. Lacunarity analysis: A general technique for the analysis of spatial patterns. *Phys. Rev. E* **1996**, *53*, 5461–5468.
5. Dong, P. Test of a new lacunarity estimation method for image texture analysis. *Int. J. Remote Sens.* **2000**, *21*, 3369–3373.

6. Sarkar, N.; Chaudhuri, B.B. An efficient differential box-counting approach to compute fractal dimension of image. *IEEE Trans. Syst. Man Cybern.* **1994**, *24*, 115–120.
7. Smeds, E.A.; Cooper, Z.; Bentley, L.P. lacunr: Efficient 3D lacunarity for voxelized LiDAR data from forested ecosystems. *Methods Ecol. Evol.* **2025**, *16*, 1906–1913. <https://doi.org/10.1111/2041-210X.70126>.
8. Bouchaud, J.P.; Potters, M. *Theory of Financial Risk and Derivative Pricing*, 2nd ed.; Cambridge University Press: Cambridge, UK, 2003.
9. Hurst, H.E. Long-term storage capacity of reservoirs. *Trans. Am. Soc. Civ. Eng.* **1951**, *116*, 770–799.
10. Laloux, L.; Cizeau, P.; Bouchaud, J.P.; Potters, M. Noise dressing of financial correlation matrices. *Phys. Rev. Lett.* **1999**, *83*, 1467–1470.
11. Plerou, V.; Gopikrishnan, P.; Rosenow, B.; Amaral, L.A.N.; Guhr, T.; Stanley, H.E. Random matrix approach to cross correlations in financial data. *Phys. Rev. E* **2002**, *65*, 066126.
12. Marchenko, V.A.; Pastur, L.A. Distribution of eigenvalues for some sets of random matrices. *Math. USSR-Sb.* **1967**, *1*, 457–483.
13. MacMahon, M.; Garlaschelli, D. Community detection for correlation matrices. *Phys. Rev. X* **2015**, *5*, 021006.
14. Münnix, M.C.; Shimada, T.; Schäfer, R.; Levvraz, F.; Seligman, T.H.; Guhr, T.; Stanley, H.E. Identifying states of a financial market. *Sci. Rep.* **2012**, *2*, 644.
15. Kritzman, M.; Li, Y.; Page, S.; Rigobon, R. Principal components as a measure of systemic risk. *J. Portf. Manag.* **2011**, *37*, 112–126.
16. Molero-González, L.; Trinidad-Segovia, J.E.; Sánchez-Granero, M.A.; García-Medina, A. Market Beta is not dead: An approach from random matrix theory. *Financ. Res. Lett.* **2023**, *55*, 103816.
17. Vanni, F.; Hitaj, A.; Mastrogiacomo, E. Enhancing portfolio allocation: A random matrix theory perspective. *Mathematics* **2024**, *12*, 1389.
18. Wang, G.J.; Xie, C.; Chen, S.; Yang, J.J.; Yang, M.Y. Random matrix theory analysis of cross-correlations in the US stock market. *Phys. A Stat. Mech. Appl.* **2013**, *392*, 3715–3730.
19. Wang, Z.; Oates, T. Imaging time-series to improve classification and imputation. In Proceedings of the 24th International Joint Conference on Artificial Intelligence (IJCAI), Buenos Aires, Argentina, 25–31 July 2015; pp. 3939–3945.
20. Vargas, V.M.; Ayllón-Gavilán, R.; Durán-Rosal, A.M.; Gutiérrez, P.A.; Hervás-Martínez, C.; Guijo-Rubio, D. Gramian angular and Markov transition fields applied to time series ordinal classification. In *Advances in Computational Intelligence (IWANN 2023)*; Springer: Cham, Switzerland, 2023; pp. 505–516.
21. Xu, Z.; Wang, Y.; Feng, X.; Wang, Y.; Li, Y.; Lin, H. Quantum-enhanced forecasting: Leveraging quantum Gramian angular field and CNNs for stock return predictions. *arXiv* **2023**, arXiv:2310.07427.
22. Long, H.W.; Ho, O.I.; He, Q.Q.; Si, Y.W. Transfer learning in financial time series with Gramian angular field. In Proceedings of the International Joint Conference on Neural Networks (IJCNN), Rome, Italy, 30 June–5 July 2025; pp. 1–10.
23. Braun, T.; Unni, V.R.; Sujith, R.I.; Kurths, J.; Marwan, N. Detection of dynamical regime transitions with lacunarity as a multiscale recurrence quantification measure. *Nonlinear Dyn.* **2021**, *104*, 3955–3973.
24. Eckmann, J.P.; Kamphorst, S.O.; Ruelle, D. Recurrence plots of dynamical systems. *Europhys. Lett.* **1987**, *4*, 973–977.
25. Marwan, N.; Romano, M.C.; Thiel, M.; Kurths, J. Recurrence plots for the analysis of complex systems. *Phys. Rep.* **2007**, *438*, 237–329.
26. Sandhu, R.S.; Georgiou, T.T.; Tannenbaum, A.R. Ricci curvature: An economic indicator for market fragility and systemic risk. *Sci. Adv.* **2016**, *2*, e1501495.
27. Samal, A.; Pharasi, H.K.; Ramaia, S.J.; Kannan, H.; Saucan, E.; Jost, J.; Chakraborti, A. Network geometry and market instability. *R. Soc. Open Sci.* **2021**, *8*, 201734.
28. Sánchez García, J.; Gherge, S. On the Ollivier–Ricci curvature as fragility indicator of the stock markets. *arXiv* **2024**, arXiv:2405.07134.
29. Kantelhardt, J.W.; Zschiegner, S.A.; Koscielny-Bunde, E.; Havlin, S.; Bunde, A.; Stanley, H.E. Multifractal detrended fluctuation analysis of nonstationary time series. *Phys. A Stat. Mech. Appl.* **2002**, *316*, 87–114.
30. Peng, C.K.; Buldyrev, S.V.; Havlin, S.; Simons, M.; Stanley, H.E.; Goldberger, A.L. Mosaic organization of DNA nucleotides. *Phys. Rev. E* **1994**, *49*, 1685–1689.
31. Gu, R.; Chen, H.; Wang, Y. Multifractal analysis on international crude oil markets based on the multifractal detrended fluctuation analysis. *Phys. A Stat. Mech. Appl.* **2010**, *389*, 2805–2815.
32. Wang, Y.; Wei, Y.; Wu, C. Analysis of the efficiency and multifractality of gold markets based on multifractal detrended fluctuation analysis. *Phys. A Stat. Mech. Appl.* **2011**, *390*, 817–827.
33. Markowitz, H. Portfolio selection. *J. Financ.* **1952**, *7*, 77–91.
34. Forbes, K.J.; Rigobon, R. No contagion, only interdependence: Measuring stock market comovements. *J. Financ.* **2002**, *57*, 2223–2261.
35. Engle, R. Dynamic conditional correlation: A simple class of multivariate GARCH models. *J. Bus. Econ. Stat.* **2002**, *20*, 339–350.
36. Diebold, F.X.; Yilmaz, K. On the network topology of variance decompositions: Measuring the connectedness of financial firms. *J. Econom.* **2014**, *182*, 119–134.
37. Billio, M.; Getmansky, M.; Lo, A.W.; Pelizzon, L. Econometric measures of connectedness and systemic risk in the finance and insurance sectors. *J. Financ. Econ.* **2012**, *104*, 535–559.
38. Adrian, T.; Brunnermeier, M.K. CoVaR. *Am. Econ. Rev.* **2016**, *106*, 1705–1741.
39. Acharya, V.V.; Pedersen, L.H.; Philippon, T.; Richardson, M. Measuring systemic risk. *Rev. Financ. Stud.* **2017**, *30*, 2–47.

40. Bisias, D.; Flood, M.; Lo, A.W.; Valavanis, S. A survey of systemic risk analytics. *Annu. Rev. Financ. Econ.* **2012**, *4*, 255–296.
41. Mantegna, R.N. Hierarchical structure in financial markets. *Eur. Phys. J. B* **1999**, *11*, 193–197.
42. Choueifat, Y.; Coignard, Y. Toward maximum diversification. *J. Portf. Manag.* **2008**, *35*, 40–51.
43. Meucci, A. Managing diversification. *Risk* **2009**, *22*, 74–79.
44. Tolle, C.R.; McJunkin, T.R.; Gorsich, D.J. An efficient implementation of the gliding box lacunarity algorithm. *Phys. D Nonlinear Phenom.* **2008**, *237*, 306–315.

Disclaimer/Publisher’s Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.