



Article

Varifold Lifts of Visibility Graphs: Beyond Fractality and the Geometry of Safe Haven Decoupling in Commodity and Currency Markets

Mehmet Ali Balcı ^{1,*} , Ömer Akgüller ^{1,*} , Deniz Rümeysa Erdoğan ¹ and Lucian Gaban ^{2,*}

¹ Department of Mathematics, Faculty of Science, Mugla Sıtkı Kocman University, 48000 Muğla, Türkiye; mehmetalibalci@mu.edu.tr (M.A.B.); denizrumeysaerdogan@posta.mu.edu.tr (D.R.E.)

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

* Correspondence: oakguller@mu.edu.tr (Ö.A.); lucian.gaban@uab.ro (L.G.)

Abstract

Visibility graphs map time series to networks whose combinatorial structure encodes fractality, recovering the Hurst exponent of self-affine processes. We ask what the visibility construction carries beyond this fractal content. We lift the visibility graph to a 1-varifold, a measure on position and direction space from geometric measure theory, and equip it with a multiscale positive definite kernel. The lift embeds visibility graphs of unequal size in a common Hilbert space and yields a channel-resolved measure of cross-series geometric alignment. On a 25.8-year daily panel of thirteen commodity and currency layers, we define a relative alignment contrast that compares commodity currencies and safe haven currencies in their geometric alignment with the commodity complex. During global risk-off episodes the contrast is large and positive: commodity currencies import commodity shock geometry far beyond a time-shift independence benchmark, while the Japanese yen remains near geometric independence and the franc is confounded by a managed regime. The contrast is significant under three stress definitions with autocorrelation robust inference, holds as a continuous dose response, survives the removal of any single crisis, withstands moment, fractal, and topological controls, is direction-consistent across sixteen specifications, and collapses under a time-shift placebo. Detrended fluctuation analysis explains only two percent of it, so the reconfiguration is geometric information beyond fractality at this horizon, and a scaling exponent of the kernel mass separates a fractal-free component from a fractal-driven one. For investors, financial institutions, and policymakers, the contrast is a real-time structural diagnostic of flight to safety: it marks when commodity currencies stop diversifying the commodity complex while genuine safe havens still do, signaling through a channel that second-moment risk measures are built to miss.



Academic Editor: Palle Jorgensen

Received: 6 June 2026

Revised: 8 July 2026

Accepted: 11 July 2026

Published: 13 July 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: visibility graph; varifold; geometric measure theory; multiscale kernel; Hurst exponent; multifractal detrended fluctuation analysis; safe haven currencies; commodity markets; financial stress; multiplex networks

1. Introduction

The visibility graph (VG) entered the analysis of time series as a bridge between dynamics and network science: each observation becomes a node, and two observations are linked whenever the straight segment connecting them stays above all intermediate observations [1]. The construction was immediately recognized as a fractality encoder. For

fractional Brownian motion, the exponent of the VG degree distribution is an affine function of the Hurst exponent [2], the horizontal variant admits exact combinatorial results that separate chaotic from stochastic dynamics [3], and degree-based VG statistics recover Hurst exponents of exchange rate series in close agreement with the prediction for self-affine processes [4]. A decade of applications, surveyed in [5] alongside parallel network and symbolic representations of dynamics [6,7], has consolidated a clear division of labor. The combinatorial skeleton of the VG is read as a proxy for scaling properties. The geometry of the construction, the actual positions, lengths, and slopes of the visibility segments in the time and value plane, is discarded at the moment the adjacency matrix is formed.

In finance, this combinatorial information has been put to work directly. Horizontal visibility graphs of equity indices track how crises reshape price dynamics [8], and a recent study turns backward-visibility relations into a relative-strength-style technical indicator that produces profitable trading signals across equity, currency, and gold instruments [9]. These results show that the visibility construction carries information of practical value, and they motivate asking what further content is available once its discarded geometry is restored.

This discarded geometry is the subject of the present paper. Our motivating question is deliberately posed against the lineage of the method: what does the visibility construction know that fractality does not? The question is sharpened by the stylized facts of financial series, where scaling and multifractality are pervasive [10–12] and standard estimators such as detrended fluctuation analysis (DFA) [13] and its multifractal extension (MF-DFA) [14] are routinely applied to commodity and currency markets [15]. If a new representation built on visibility graphs merely re-measures the Hurst content that the VG skeleton is already known to encode, it adds computational cost without adding information. Any claim of novelty must therefore be tested against fractal benchmarks directly, and this identification problem organizes the entire design of our study.

Our answer is constructed with tools from geometric measure theory. A 1-varifold is a Radon measure on the product of ambient space and the Grassmannian of unoriented directions; it is the natural carrier of first-order geometric information for sets that need not be smooth [16,17]. Varifolds and the closely related currents underpin modern computational shape analysis, where kernel metrics on these measures provide registration invariant, correspondence-free distances between curves and surfaces [18–20], and have recently been combined with learning architectures for point-cloud geometry [21,22]. We observe that a visibility graph is not an abstract graph: it is a geometric object embedded in the normalized time-and-value plane, a finite union of segments, and therefore lifts canonically to a 1-varifold. Each edge contributes an atom located at its midpoint, carrying its unoriented direction and weighted by its length. A multiscale positive definite kernel on the lift then embeds visibility graphs of arbitrary, unequal size into one reproducing kernel Hilbert space, where alignment between two series is a single normalized inner product. The lift retains exactly what the adjacency matrix forgets.

The economic laboratory in which we interrogate the lift is the commodity and currency nexus. Currencies of commodity exporting economies are tied to commodity prices through the terms of trade. The Swiss franc and the Japanese yen, in contrast, appreciate in episodes of global distress, a behavior documented at the level of returns and risk reversals [23–25] and connected to the unwinding of carry trades [26]. Gold plays a parallel role among commodities [27,28]. These roles are visible in our own panel before any modeling. Averaged over the 634 risk-off days on which the VIX exceeds 30, the Japanese yen and gold are the only two assets among the thirteen with a positive mean daily return, +6.1 and +4.5 basis points, respectively. Gold has the highest stress return of the six commodities, the other five all falling, and the yen appreciates while the five commodity currencies decline

by 9 to 14 basis points per day. The franc is essentially flat in stress, -1.2 basis points, a first sign of the managed regime neutrality that the geometric analysis later isolates. Crisis-driven reconfiguration of dependence structures in these markets has been studied with correlation networks and their filtered variants [29–33], spillover decompositions [34], connectedness analyses of the pandemic shock [35], and, closest in spirit to our geometric viewpoint, topological data analysis of sliding windows [36,37]. None of these strands, however, measures whether the shapes of price trajectories align, and none can do so for series whose dependence is invisible to second moments.

The paper makes four contributions. First, we define the varifold lift of a visibility graph, prove positive definiteness and invariance properties of the associated multiscale kernel, and obtain a stability statement for the induced alignment functional (Section 3). Second, we validate the machinery on a synthetic ground truth: geometric descriptors of the lift track the Hurst exponent of fractional Brownian motion monotonically, the alignment functional recovers a true coupling parameter, and, critically, it detects a sign independent common volatility coupling to which linear correlation is blind by construction (Section 5). Third, on a daily panel of six commodities, five commodity currencies, and two safe haven currencies spanning August 2000 to June 2026, we document a channel-resolved reconfiguration of cross-layer geometry during global risk-off episodes. We condense the finding into a single pre specified primary endpoint, the relative alignment contrast D . The contrast survives autocorrelation robust inference under three stress definitions, holds as a continuous dose response, survives the removal of any single crisis, withstands moment, fractal, and topological controls, is direction-consistent across sixteen specifications, and collapses under a time-shift placebo that pins the mechanical component of the effect at zero (Sections 4 and 6). Fourth, we close the identification loop that the VG lineage demands. Window-level DFA and MF-DFA estimates explain two percent of the variance of D , and stress does not move window-level Hurst exponents. The classic combinatorial degree channel captures only part of the effect while loading on fractality. A newly defined scaling exponent of the varifold kernel mass is flat across Hurst regimes yet contracts in stress. Together these results decompose visibility graph geometry into a fractal-driven component and a fractal-free component, and they locate the economically meaningful risk-off signal in the latter (Section 6.4).

We also report, deliberately and in detail, what failed. Supervised learning on the lift, whether with a learnable kernel embedding, a frozen embedding, or a linear varifold machine in the spirit of [22], does not transfer across market eras. A hedging-based exhibit of economic significance returned a null. The raw decoupling ratio that a naive analysis would headline is partly a mechanical artifact, exposed by our own placebo. The honest accounting of these negatives, and the reformulation they forced, is part of the methodological contribution.

The remainder of the paper is organized as follows. Section 2 reviews the related literature and the theoretical background. Section 3 develops the lift, the kernel, the alignment functional, the contrast endpoint, and the fractal instruments. Section 4 describes the panel and the empirical protocol. Section 5 reports synthetic validation. Section 6 contains the empirical anatomy, the primary inference, the fractal decomposition, and the robustness program. Section 7 discusses mechanisms, practical interpretation, limitations, and the negative results. Section 8 concludes.

2. Related Work and Theoretical Background

2.1. Visibility Graphs in Finance

The visibility graph was introduced as a fractality encoder, with the degree distribution exponent of the natural graph affine in the Hurst exponent for self-affine processes [1,2],

exact results for the horizontal variant [3], and accurate recovery of Hurst exponents for exchange rates [4]. In finance, the construction has been used to characterize how crises reshape equity index dynamics through horizontal visibility graphs [8], to study irreversibility and multivariate coupling [38], and most recently to build a relative-strength-style trading indicator that produces profitable signals on equity, currency, and gold instruments [9]. Across these strands, the object actually analyzed is the combinatorial graph, summarized by degree-based or path-based statistics; the embedded geometry of the visibility segments is discarded, and no existing visibility graph method measures whether the shapes of two distinct series are aligned.

2.2. Dependence Structure and Safe Havens

The reconfiguration of cross-asset dependence in stress has been studied with correlation networks and their filtered variants [29–33], with variance decomposition spillovers [34,35], and, closest to our geometric viewpoint, with topological data analysis of sliding windows [36,37]. In parallel, the safe haven literature establishes that the Swiss franc and the Japanese yen appreciate in global distress [23–25], that this is tied to the unwinding of carry trades funded in low yield currencies [26], and that gold plays a comparable role among commodities [27,28]. These studies describe when and how strongly assets co-move, but they operate on returns and second moments and are therefore silent about the geometry of the co-movement.

2.3. Fractal Market Hypothesis as Theoretical Foundation

The fractal market hypothesis provides the theoretical frame that connects these threads to our construction [39]. It models a market as a population of heterogeneous agents distinguished by investment horizon, and attributes market stability to the diversity of those horizons: when many horizons coexist, a seller at one horizon meets a buyer at another, and liquidity is maintained. In a crisis, the horizon distribution collapses toward the short end, agents become homogeneous, liquidity evaporates, and the scaling structure of prices changes, a sequence that has been documented empirically around the global financial crisis [40]. Two implications shape our design. First, stress should leave a multiscale signature, which is why our kernel reads geometry across a range of scales rather than at one. Second, a safe haven is, in this language, an asset whose horizon structure decouples from the risk complex when horizons collapse, which is precisely the relative geometric decoupling our contrast is built to detect.

2.4. Positioning and Novelty

Against this background, the contribution is specific. We retain the embedded first-order geometry that the combinatorial visibility graph discards, we lift it to a varifold and compare series through a single positive definite multiscale kernel, and we use this to measure cross-series shape alignment, an object that correlation networks, spillover decompositions, topological summaries, and existing visibility graph statistics do not form. The economic question we pose with it, whether safe haven currencies decouple geometrically from the commodity complex in a way that second moments cannot see, is new, and the identification of a fractal-free component of the effect closes the loop with the fractal market hypothesis that motivates the multiscale design.

3. The Varifold Lift of a Visibility Graph

3.1. Visibility Graphs and Their Embedded Geometry

Let $x = (x_0, \dots, x_{n-1})$ be a window of n consecutive observations of the logarithmic price of one market layer.

Definition 1 (Natural visibility graph [1]). *The natural visibility graph $G(x) = (\{0, \dots, n-1\}, E(x))$ contains the edge (a, b) , $a < b$, if and only if*

$$x_k < x_a + (x_b - x_a) \frac{k - a}{b - a} \quad \text{for all } a < k < b. \quad (1)$$

The criterion is invariant under positive affine transformations of the value axis and under affine reparametrizations of time, so the combinatorial object $G(x)$ is unchanged by the normalization that follows. To expose the geometry, we map the window to the plane by

$$t_i = \frac{i}{n-1}, \quad \tilde{x}_i = \alpha \frac{x_i - \min_j x_j}{\max_j x_j - \min_j x_j}, \quad (2)$$

with aspect parameter $\alpha = 0.3$. Some unit convention here is unavoidable rather than incidental: time and logarithmic price are dimensionally incommensurable, so any embedding of the visibility structure into a metric plane must fix their relative scale, exactly as a kernel density estimate must fix a bandwidth. The visibility criterion itself is invariant to the choice; only the lifted geometry is calibrated by it, and the calibration is treated as part of the specification grid, swept over a factor of four and replaced by a z-score alternative, rather than as a tuned constant. Each edge $e = (a, b) \in E(x)$ is now a segment with endpoints $p_a = (t_a, \tilde{x}_a)$ and $p_b = (t_b, \tilde{x}_b)$, midpoint c_e , length $w_e = \|p_b - p_a\|$, and unoriented direction $\theta_e \in [0, \pi)$, the angle of the line through p_a and p_b .

3.2. The Lift

Before the formalism, the idea of the lift can be stated in plain terms for a reader who knows visibility graphs but not geometric measure theory. Draw the price window as a curve, draw every visibility edge as a straight segment, and then forget the node labels. What remains is a collection of line segments scattered across the time-and-value plane, and each segment carries three attributes: a location, a tilt, and a length. The varifold lift is simply the bookkeeping of these attributes. It places a small mass at the midpoint of each segment, tags that mass with the segment's tilt, and makes it proportional to the segment's length. The visibility graph thereby becomes a cloud of located, tilted, weighted points rather than a table of zeros and ones. Two windows are then compared not edge by edge, which would force them to have the same number of nodes, but by how similar these clouds are. The role of the varifold formalism is to supply the right notion of similarity between such clouds, one that is invariant to relabeling and behaves well under resampling and noise; the object being compared is just a weighted cloud of segment positions and directions. Recall that a 1-varifold on $\mathbb{R}^2$ is a Radon measure on $\mathbb{R}^2 \times G(2, 1)$, where the Grassmannian $G(2, 1)$ of lines through the origin is identified with $[0, \pi)$ [16,17].

Definition 2 (Varifold lift). *The varifold lift of the window x is the discrete 1-varifold*

$$V(x) = \sum_{e \in E(x)} w_e \delta_{(c_e, \theta_e)}. \quad (3)$$

The lift is the atomic approximation of the varifold of the embedded edge set viewed as a 1-rectifiable union of segments, with each segment represented by its midpoint atom; mass corresponds to total edge length. Two windows with different numbers of observations or edges produce measures on the same space, which is the property that the adjacency representation lacks.

3.3. Multiscale Kernel and Alignment

Fix scales $\sigma_1, \dots, \sigma_S > 0$ and concentration $\kappa > 0$, and define on $\mathbb{R}^2 \times [0, \pi)$ the kernel

$$k((c, \theta), (c', \theta')) = \left(\sum_{s=1}^S \exp \left(- \frac{\|c - c'\|_s^2}{2\sigma_s^2} \right) \right) \exp \left(\kappa (\cos^2(\theta - \theta') - 1) \right). \quad (4)$$

The benchmark configuration uses $\sigma \in \{0.02, 0.05\}$ and $\kappa = 4$; the direction factor is a von Mises kernel in the doubled angle, sharper than the classical Cauchy and Binet factor $\cos^2(\theta - \theta')$ used in shape analysis [19], which we retain as a robustness variant.

Three free parameters enter the construction, and none of them is tuned to the outcome. The aspect parameter α fixes the relative scale of the two incommensurable axes, as explained above, and is set to 0.3. The two positional scales are chosen to read the geometry at fine and intermediate separations rather than at the coarse scale where windows are nearly indistinguishable, and the concentration κ is set so that the direction kernel is selective without being degenerate. Rather than defend any single value, we treat all three as a specification grid and report the primary endpoint across it in Section 6.5: the aspect parameter is varied over 0.15, 0.30, and 0.60 and replaced by a z-score normalization; the positional scale is swept over the full range 0.01 to 0.20 both singly and in combination, and the von Mises factor is replaced by the $\cos^2$ factor. The stress minus calm difference of the contrast remains positive with bootstrap intervals excluding zero in every one of these configurations and its scale profile shows that the effect is present at every scale and merely peaks at fine to intermediate separations. The main conclusions are therefore not sensitive to the parameter choices.

Proposition 1 (Positive definiteness). *The kernel (4) is positive definite on $\mathbb{R}^2 \times [0, \pi)$, and*

$$\langle V, V' \rangle = \sum_e \sum_{e'} w_e w_{e'} k((c_e, \theta_e), (c_{e'}, \theta_{e'})) \quad (5)$$

defines an inner product on the linear span of varifold lifts, with norm $\|V\| = \langle V, V \rangle^{1/2}$.

Proof. Sums of Gaussian kernels are positive definite on $\mathbb{R}^2$. Writing $\phi = 2\theta$, the direction factor equals $e^{\kappa(\cos \phi - 1)/1}$ composed with the doubling map; the von Mises kernel $e^{\kappa \cos(\phi - \phi')}$ has a non-negative Fourier expansion in $\cos m(\phi - \phi')$ with modified Bessel coefficients $I_m(\kappa) \geq 0$, hence is positive definite on the circle, and positive definiteness is preserved under the pullback by the doubling map and under multiplication by the positive constant $e^{-\kappa}$. Products and sums of positive definite kernels are positive definite, and (5) is the induced inner product of the kernel mean embeddings of the discrete measures. $\square$

Definition 3 (Geometric alignment). *The alignment of two windows x, y observed over the same dates is*

$$\mathcal{A}(x, y) = \frac{\langle V(x), V(y) \rangle}{\|V(x)\| \|V(y)\|} \in [0, 1]. \quad (6)$$

Non-negativity follows because every kernel evaluation in (4) is positive. The functional $\mathcal{A}$ is invariant to relabeling of edges, to the number of edges, and, through (2), to positive affine transformations of each series separately. It is differentiable in the atoms (c_e, θ_e, w_e) , a property we exploited when testing learnable variants. It is also stable. By the reproducing property, $|\langle V, W \rangle - \langle V', W \rangle| \leq \|V - V'\|_k \|W\|_k$, and the kernel norm of the difference of two lifts is controlled by the displacements of atoms and weights through the Lipschitz moduli of (4), in direct analogy with the stability results for shape varifolds [19,20].

Perturbing one observation of x moves finitely many atoms by a bounded amount, so $\mathcal{A}$ inherits robustness to observational noise at a fixed window length.

3.4. Blocks, the Primary Endpoint, and the Independence Benchmark

The panel introduced in Section 4 consists of a commodity block $\mathcal{C}$ with six layers, a commodity–currency block $\mathcal{F}$ with five layers, and a safe haven block $\mathcal{S} = \{\text{CHF}, \text{JPY}\}$. For window date τ , write $\mathcal{A}_\tau(\ell, \ell')$ for the alignment of layers ℓ and ℓ' , and for two blocks A, B , define the block coupling $\bar{\mathcal{A}}_\tau(A, B)$ as the average of $\mathcal{A}_\tau(\ell, \ell')$ over distinct pairs.

A naive analysis would track the raw safe haven coupling $\bar{\mathcal{A}}_\tau(\mathcal{S}, \mathcal{C})$ and read its drop in stress as decoupling. Our placebo analysis shows that part of such a drop is mechanical. In stress, commodity windows become geometrically atypical, and atypicality lowers their alignment with any reference series, including a time-shifted copy of the safe haven series that cannot share information with the contemporaneous market. The well-posed quantity is relative.

Definition 4 (Relative alignment contrast). *The primary endpoint of the study is*

$$D_\tau = \bar{\mathcal{A}}_\tau(\mathcal{C}, \mathcal{F}) - \bar{\mathcal{A}}_\tau(\mathcal{S}, \mathcal{C}). \quad (7)$$

The mechanical component of stress-induced atypicality enters both terms through the commodity side and cancels to first order in the difference, a cancellation we verify empirically by applying the same contrast to time-shifted surrogates. The endpoint was fixed before the inferential battery below was run on it; all earlier channel-level findings are reported as anatomy.

3.5. Fractal Instruments

Three instruments connect the lift to fractal analysis. First, window-level Hurst exponents H are estimated by DFA of order one [13] on the increments of each window, with fluctuation scales of between 4 and 28 observations. Second, multifractal width $\Delta h = h(-4) - h(4)$ is estimated by MF-DFA [14] on the same grid. These serve as controls: if the contrast (7) is fractal information in disguise, conditioning on H and Δh of both blocks must absorb it. Third, we define a scaling exponent intrinsic to the lift.

Definition 5 (Varifold scaling exponent). *For a window x with lift (3), let $m(\sigma) = \sum_{e, e'} w_e w_{e'} \exp(-\|c_e - c_{e'}\|^2 / 2\sigma^2)$ denote the positional kernel mass at scale σ . The varifold scaling exponent is the slope*

$$D_{\text{eff}}(x) = \frac{d \log m(\sigma)}{d \log \sigma} \quad (8)$$

estimated by least squares over $\sigma \in [0.01, 0.05]$.

The quantity is a correlation integral-type descriptor of the midpoint measure: for an exactly d dimensional measure, the Gaussian mass scales as σ^d at small σ , so a smaller slope means that a larger share of the kernel mass is already recovered at fine separations, that is, the atoms of the lift pack more tightly in space. Because the slope is estimated over a finite scale range, it should be read as a packing descriptor on that range, not as an asymptotic Hausdorff dimension. Section 6.4 shows that D_{eff} is flat across Hurst regimes of fractional Brownian motion, in contrast with directional dispersion, which tracks H closely; the pair therefore separates the fractal-driven from the fractal-free part of the lift geometry.

Figure 1 summarizes the construction, and Figure 2 displays the embedded visibility geometry of a calm and a stressed window for one currency and one commodity layer.




Figure 1. The varifold lift. (a) A 120-day logarithmic price window with its natural visibility edges. (b) The lift (3): each edge becomes an atom at its midpoint, colored by unoriented direction $\theta_e \in [0, \pi)$ and sized by length weight w_e . (c) Direction kernels: the von Mises factor with $\kappa = 4$ used in the benchmark configuration and the classical $\cos^2$ factor retained as a robustness variant.




Figure 2. Embedded visibility geometry of calm (2005) and stressed (global financial crisis, late 2008) windows for the Japanese yen and copper. The combinatorial skeleton discards the visible differences in edge directions and lengths; the lift retains them.

4. Data and Empirical Protocol

4.1. A Three Block Multiplex Panel

The panel covers 30 August 2000 to 5 June 2026, a span of 25.8 years with 6652 common trading days and, with 120-day windows advanced in steps of 5 days, 1307 analysis windows per layer. Table 1 lists the thirteen layers. Commodity prices are front month futures settlement values; exchange rates are European Central Bank reference rates crossed to United States dollar quotations. All series are converted to logarithmic prices; the single negative West Texas Intermediate settlement of 20 April 2020 is replaced by the previous close, an adjustment affecting one observation and flagged in every robustness variant. The CBOE volatility index VIX is retained outside the layer set as an exogenous stress meter [41]. The quarter-century span is chosen so that the analysis spans a wide range of stress regimes rather than a single episode. It contains the dotcom unwind of the early 2000s, the 2008 global financial crisis, the 2010 to 2012 euro area sovereign debt crisis, the 2014 to 2016 commodity price collapse, the 2020 coronavirus shock with its episode of negative oil prices, the 2022 Russia and Ukraine war and the associated energy and grain disruptions, the late 2022 turmoil in digital asset markets, and the 2024 to 2025 escalation in the Middle East. This breadth is what makes the leave-one-crisis-out and dose-response checks of Section 6.3 meaningful, since no single event dominates the sample.

Table 1. Panel composition. All layers are daily and span 30 August 2000 to 5 June 2026 without missing values after a three-day-holiday alignment rule.

Block	Layers	Source	Role
Commodities $\mathcal{C}$	gold, silver, copper, WTI crude, natural gas, wheat	front month futures	shock origin
Commodity currencies $\mathcal{F}$	AUD, CAD, NOK, NZD, ZAR (per USD)	ECB reference rates	terms of trade channel
Safe havens $\mathcal{S}$	CHF, JPY (per USD)	ECB reference rates	flight to safety channel
Auxiliary	VIX	CBOE	stress meter, not a layer

The combination is deliberately distinct from the equity index, oil and gold, and cryptocurrency panels that dominate the VG literature [5]: the three blocks have economically distinct, theory-assigned roles, which is what gives the cross-layer geometry testable structure [23–26].

4.2. Regimes, Inference, and Multiplicity Discipline

The principal stress classification marks a window as risk-off when VIX exceeds 30 at the window end (124 windows) and as calm when VIX is below 20 (830 windows under the strict definition used in channel inference); intermediate windows are excluded from two sample contrasts. Because windows overlap heavily, ordinary standard errors are invalid; all regression-based inference uses Newey and West covariance with 24 lags, matching the window step overlap [42], two sample location statements are accompanied by Mann and Whitney tests [43] reported as descriptive complements, and moving-block bootstrap intervals with block length 24 [44] are reported where regression is unavailable. With a 5-day step and 120-day windows, consecutive windows overlap by 96 percent, so the panel contains roughly 55 fully non-overlapping windows per layer and the effective confirmatory sample is far below the nominal 1307; the HAC lag and the bootstrap block length are set to the full overlap span for this reason, and all location statements should be read against this effective size. The lag choice can only err in the conservative direction, and the inference is insensitive to it: for the primary endpoint, the stress statistic declines monotonically from $t = 3.61$ at lag 6 to $t = 2.63$ at lag 24 and then stabilizes, $t = 2.57$ at lag 36 and $t = 2.54$ at lag 48, so the reported scheme sits on the conservative plateau and longer memory corrections leave the conclusions unchanged. The endpoint (7) is the single pre-specified primary quantity; every other comparison is anatomy or robustness, and the specification grid of Section 6.5 is exhaustive disclosure rather than selection: every configuration that was computed is reported, including the placebo. Two alternative stress definitions, the upper decile of 120-day realized volatility of the S&P 500 index and a 20 day S&P 500 drawdown beyond eight percent, remove the dependence of the conclusions on the VIX.

5. Synthetic Validation

Three ground-truth experiments establish that the machinery measures what it claims to measure; results are condensed in Figure 3.

First, fractality tracking. For fractional Brownian motion generated by circulant embedding at $H \in \{0.3, \dots, 0.8\}$ with $n = 120$, the mass-weighted directional dispersion of the lift decreases monotonically from 0.634 to 0.132 as H rises, and edge counts increase concordantly; the lift therefore inherits the fractality sensitivity that the VG skeleton is known to possess [2,4].

Second, coupling recovery. For pairs of correlated random walks with true coupling $\lambda \in [0, 1]$, the alignment (6) rises monotonically from 0.216 at independence to 1 at identity, tracking the absolute correlation benchmark.

Third, and decisively for our identification, correlation blind coupling. Two return series share a common log volatility factor with weight λ , while their signs are independent, so the population linear correlation is zero for every λ . Empirical absolute correlation indeed stays flat between 0.06 and 0.12, while the alignment rises from 0.199 at $\lambda = 0$ to 0.250 at $\lambda = 0.8$ with confidence bands separated from the independent case. Sign-free, volatility-timed dependence is precisely the channel through which the lift can see what second moments cannot.




Figure 3. Synthetic ground truth. (a) Directional dispersion of the lift tracks the Hurst exponent of fractional Brownian motion. (b) For correlated random walks, both the alignment (6) and absolute correlation recover the true coupling λ . (c) For sign-independent common-volatility coupling, correlation is blind by construction while the alignment detects the dependence. Bands are 95% intervals over 100 replications.

6. Empirical Results

6.1. Anatomy of Cross-Layer Geometry in Stress

Everything in this subsection is exploratory anatomy that motivates the endpoint; confirmatory inference is reserved for the pre-specified contrast of Section 6.3. Figure 4 shows the average alignment matrices in calm and risk-off windows and their difference. Three features organize the picture. Within-block alignment of the commodity currencies rises sharply in stress, the commodity-to-commodity currency coupling rises, and the rows of the yen and of gold turn negative against the commodity complex. Channel-level statistics (Table 2) quantify the pattern and compare it with absolute rolling correlation on identical windows and pairs.




Figure 4. Average alignment matrices over (a) calm windows, (b) risk-off windows, and (c) their difference. White lines separate the commodity, commodity currency, and safe haven blocks. The difference panel displays the joint signature: commodity and commodity currency coupling rises (red blocks), while yen and gold rows decouple (blue bands).

Table 2. Channel anatomy: stress-to-calm ratios of mean coupling, with Newey and West t statistics for the stress regressor where the channel is a single time series, and the corresponding absolute correlation channel on identical windows. Stress: VIX > 30; calm: VIX < 20.

Channel	Varifold Ratio	Varifold p	Correlation Ratio	Correlation p
$\mathcal{C} \leftrightarrow \mathcal{F}$	1.117	6.2×10^{-5}	1.305	1.5×10^{-9}
$\{\text{CHF, JPY}\} \leftrightarrow \mathcal{C}$	0.876	1.0×10^{-5}	1.03	0.2
$\text{JPY} \leftrightarrow \mathcal{C}$	0.778	0.013	covered above covered above	
$\text{CHF} \leftrightarrow \mathcal{C}$	0.964	0.71		
gold $\leftrightarrow$ other commodities	0.905	0.030	0.96	0.08
safe havens $\leftrightarrow \mathcal{F}$	0.89	2×10^{-3}	0.83	4×10^{-7}

Two asymmetries deserve emphasis because they discipline the interpretation. First, the decoupling of the safe haven block from the commodity complex itself is invisible to correlation (ratio 1.03, $p = 0.2$), while the safe haven-to-commodity currency channel is visible to both representations. The geometric novelty is channel-specific, and we claim it only where it exists. Second, within the safe haven block, the raw effect is carried by the yen (ratio 0.778, $p = 0.013$). The franc alone is silent (0.964), and a Swiss National Bank floor era analysis (September 2011 to mid 2015, excluded windows including overlap) shows that the franc contributes nothing outside the managed regime. The raw channel anatomy is therefore yen led, consistent with the carry unwinding narrative [25,26].

6.2. The Placebo That Reshaped the Endpoint

A time-shift surrogate exposes the limit of raw channel readings. Shifting the safe haven series by approximately two trading years destroys any contemporaneous economic linkage. Yet the surrogate safe haven-to-commodity coupling still falls in stress (ratio 0.778), because stressed commodity windows are geometrically atypical against any reference. The counter surrogate completes the logic. Shifted commodity currencies fall to 0.806 in the same windows in which the true commodity currencies rise to 1.12 and beyond. The coupling side is therefore genuine dependence, while the level of the decoupling side sits near an independence benchmark. The contrast (7) differences out the shared mechanical component: applied to the doubly shifted surrogate pair, it collapses to +0.005 with a bootstrap interval $[-0.014, 0.022]$ straddling zero, while the true contrast difference is +0.068 with interval $[0.048, 0.088]$. The full surrogate distributions confirm the collapse beyond their means: the placebo contrast has mean 0.011 and standard deviation 0.079 in stress windows against mean 0.006 and standard deviation 0.063 in calm windows, with central eighty percent ranges $[-0.05, 0.09]$ and $[-0.07, 0.08]$ that overlap almost completely. All inference that follows concerns D .

6.3. Primary Endpoint Inference

Table 3 assembles the inferential battery on D , and Figure 5 displays its time series against the five canonical crisis onsets of the sample. The stress effect on D is positive and significant under Newey and West inference for the VIX classification ($t = 2.63$, $p = 0.0084$), for the realized volatility classification ($t = 2.37$, $p = 0.018$), and for the drawdown classification ($t = 2.49$, $p = 0.013$). Conditioning on eleven window-level moment controls, absolute and rank correlations of the relevant pair block, lower and upper tail co-exceedance rates, realized volatilities, skewness, and kurtosis of both blocks, strengthens the effect ($t = 3.09$, $p = 0.0020$). The controls jointly explain 21% of the variance of D but none of its stress response. The correlation analogue of the contrast, constructed identically from absolute rolling correlations, is marginal ($t = 1.68$, $p = 0.094$), and the practitioner standard exponentially weighted correlation of the safe haven channel, the

simplest member of the conditional correlation family [45], moves in the opposite direction (ratio 1.10, $p = 0.084$), confirming that second-moment instruments do not carry the signal.

Table 3. Primary endpoint inference. Each row reports the stress coefficient for D under the stated design; HAC denotes Newey and West covariance with 24 lags. Moment controls comprise absolute and rank correlation, two tail co-exceedance rates, realized volatility, skewness, and kurtosis of both blocks. The correlation-based row is expressed in correlation units and is included for comparison of inferential strength, not of magnitudes.

Design	Stress Effect	t (HAC)	p
VIX classification	+0.068	2.63	0.0084
VIX, eleven moment controls	+0.046	3.09	0.0020
VIX, fractal controls (Section 6.4)	+0.067	2.74	0.0061
VIX, moment and fractal controls	+0.038	2.26	0.024
S&P 500 realized volatility decile	+0.066	2.37	0.018
S&P 500 drawdown classification	+0.049	2.49	0.013
VIX, topological persistence controls	+0.058	2.42	0.015
correlation based contrast (benchmark)	+0.032	1.68	0.094




Figure 5. (a) The relative alignment contrast D over 2000 to 2026 (six window moving average) with VIX shading and crisis onsets marked. (b) Its two components: commodity-to-commodity currency coupling and safe haven-to-commodity coupling.

The magnitudes admit a benchmark-referenced reading. The stress response of D is roughly fifteen times the point estimate obtained under the time-shift placebo, and the stress and calm distributions are separated by 0.85 pooled standard deviations. In the units of the independence benchmark of Section 6.2, risk-off windows move the commodity currency block from slightly above geometric independence to far beyond it, while leaving the yen essentially at the benchmark. Across the six pre-specified inferential designs that involve the varifold contrast, all rejections survive a Holm correction at the five percent level, so the conclusion does not depend on selecting a favorable design; the topological control row of Table 3 serves as a complementary shape benchmark and is discussed separately below. The distributional view is given in Figure 6. An event study around the five crisis onsets, Lehman Brothers 2008, the euro area crisis 2011, the commodity crash 2014, the COVID shock 2020, and the Russia and Ukraine war 2022, shows the contrast exiting a permutation null band after onset; the average post onset elevation has two sided permutation $p = 0.013$ over 5000 random onset placements (Figure 7).




Figure 6. Kernel density estimates of the primary endpoint D in calm and risk-off windows.




Figure 7. Event study of D around five crisis onsets, in days relative to onset, against a 95% permutation band from 2000 random onset placements; thin lines are individual episodes, the bold line is their mean relative to the pre event baseline.

The conditional structure of the effect is an interaction, not a main effect (Figure 8). Classifying windows jointly by global risk-off ($VIX > 30$) and by internal commodity turbulence (upper quintile of own panel realized volatility) yields cell means of D equal to 0.029 in calm, 0.036 under pure risk-off, 0.011 under pure commodity turbulence, and 0.145 under the joint shock; the interaction coefficient is $+0.127$ with $t = 3.59$, $p = 0.0003$. Relative safe haven decoupling is the geometric signature of flight to safety in the precise sense that it requires a commodity shock arriving during global risk-off: commodity currencies absorb the shock geometry, the yen does not.




Figure 8. Cell means of D with 95% intervals under the two-way classification by global risk-off and internal commodity turbulence. The contrast activates only under the joint shock.

6.4. Beyond Fractality

The lineage of the visibility graph obliges a direct confrontation with fractal analysis: if the contrast is Hurst content in disguise, conditioning on window-level fractal estimates must absorb it. It does not. Adding the DFA Hurst exponents and MF-DFA widths of both blocks to the stress regression leaves the effect essentially unchanged ($t = 2.74$, $p = 0.0061$); adding them on top of the eleven moment controls leaves it significant ($t = 2.26$, $p = 0.024$); the four fractal covariates explain 2.0% of the variance of D ; and stress itself does not move the window-level Hurst exponents of either block (all $|t| < 0.6$).

An obvious objection is attenuation: a noisy control cannot absorb what it cannot measure, and DFA exponents at window length 120 are noisy. We quantify the objection and find that it does not carry the weight it appears to. On synthetic fractional Brownian motion of length 120 the DFA estimator has a standard deviation of 0.075 at fixed true H . After the block averaging used in our controls, this implies noise levels of about 0.053 and 0.031 for the currency and commodity blocks, which essentially coincide with the observed cross-window dispersions of 0.059 and 0.040. The true variation of window-level Hurst exponents in this panel is therefore small to begin with. More importantly, attenuation afflicts noisy regressors, not noisy outcomes. The complementary test, whether stress moves the measured exponents, places $\hat{H}$ on the outcome side, where measurement noise only widens confidence bands. That test bounds the stress-induced movement of the currency block exponent at -0.009 with a standard error near 0.005. This is an order of magnitude too small to mediate the contrast, given that our geometric descriptors resolve Hurst differences of 0.1 without difficulty in the same synthetic setting. The reconfiguration documented here is geometric information beyond fractality, where fractality is operationalized, as is standard at this window length, by the DFA exponent together with the MF-DFA width; richer characterizations such as full singularity spectra or pointwise Hölder regularity are not estimable from 120 observations, and the claim is made relative to this operationalization.

A natural question is how commodity windows can become geometrically atypical in stress, as the placebo analysis shows, without any movement in their Hurst exponents. The two statements concern different layers of structure. Self-affinity constrains how fluctuations rescale across horizons; the lift, in contrast, records the spatial organization of the edge measure at fixed horizon, its positional spread, its direction mixture, and its mass distribution across scales. Stress reorganizes the latter, visibly in the realized volatility controls, in the directional statistics, and in the contraction of the scaling exponent below, while leaving the rescaling law intact; atypicality without Hurst shifts is therefore not a tension but a demonstration that the two descriptions are transverse.

Fractality is, moreover, not the only shape benchmark a skeptic might propose. As an alternative non-fractal descriptor family, we computed zero-dimensional sublevel set persistence of every standardized window, the merge tree summary popularized by topological data analysis in finance [36,37], and condensed it into total persistence and persistence entropy of both blocks. These topological covariates are more informative about D than the fractal family, explaining 12% of its variance against 2%, yet conditioning on them leaves the stress effect intact ($t = 2.42$, $p = 0.015$, Table 3). Persistence summaries are best read as a sibling geometric language rather than a competing explanation: they describe shape through merge events as the lift describes it through mass and direction, and the conditioning result locates the stress-relevant content of the lift precisely in the part that the two languages do not share. Taken family by family, none of the three benchmark languages, moments, fractality, and topology, absorbs the effect. Conditioning on all nineteen covariates simultaneously attenuates the coefficient to $+0.032$ with $t = 1.58$, and a partition of the covariates explains why this stack behaves differently from its parts. Six of the nineteen are mechanical responses to the magnitude of the stress shock itself: realized volatilities, tail

co-exceedance rates, and total persistence all scale with realized turbulence. The remaining twelve, dependence measures, distributional shape, fractal exponents, and persistence entropy, are normalized descriptors that do not. Conditioning on the twelve predetermined descriptors leaves the effect significant ($t = 2.27$, $p = 0.023$, Table 4). Even conditioning on the six mediators alone leaves it significant ($t = 2.03$, $p = 0.042$). Only the simultaneous stack of all nineteen collinear regressors pushes the statistic below conventional thresholds, which is the signature of collinearity rather than of absorption. The reducibility question, which the control design is meant to answer, is settled by variance accounting: all nineteen covariates together explain 24% of the variance of D .

Table 4. Event level, dose–response, and control partition checks on the primary endpoint. All rows are full panel HAC regressions with 24 lags. Leave-one-crisis-out rows exclude all windows within one year of the stated onset.

Design	Coefficient	t (HAC)	p
excluding Lehman 2008	+0.038	2.35	0.019
excluding euro area 2011	+0.074	2.60	0.0093
excluding commodity crash 2014	+0.067	2.59	0.0096
excluding COVID 2020	+0.065	2.25	0.025
excluding Russia and Ukraine 2022	+0.066	2.35	0.019
dose response: D on VIX level, all windows	+0.022	3.22	0.0013
dose response: D on log VIX, all windows	+0.020	2.69	0.0072
contrast against the yen alone	+0.091	2.13	0.033
contrast against the franc alone	+0.045	2.33	0.020
stress with twelve predetermined controls	+0.048	2.27	0.023
stress with six mediator controls alone	+0.047	2.03	0.042

The decomposition is sharpened by the scaling exponent of Definition 5 and by the classic combinatorial channel. On fractional Brownian motion, D_{eff} is flat across $H \in [0.3, 0.8]$ (range 1.155 to 1.173 with overlapping intervals), while directional dispersion tracks H strongly, so the two descriptors separate fractal-driven from fractal-free geometry on ground truth; on the panel, their rank correlation across windows and layers is -0.07 . In observed stress, the commodity block D_{eff} contracts from 1.150 to 1.125 ($t = -2.06$, $p = 0.039$): kernel mass accumulates at finer separations, meaning that the midpoint atoms of the lift pack more tightly in the time-and-value plane. Stress, in this reading, spatially concentrates the visibility geometry while leaving its rescaling law untouched, which is the same transversality documented above from the opposite side. The degree sequence channel, the information carrier of the classical VG literature [1,2,4,38], deserves emphasis for a second reason: it involves no embedding, no aspect parameter, and no kernel, so it doubles as a fully calibration-free check on the normalization conventions of Section 3. This channel sees part of the phenomenon: a contrast built from Spearman alignments of degree sequences separates regimes with effect size 0.67 against 0.85 for the varifold contrast, confirming that the direction of the effect does not depend on any geometric calibration at all. Its incremental contribution vanishes once the varifold contrast is included (joint classification coefficients 0.38 against 0.76), and unlike the varifold contrast, it loads on fractality (rank correlation -0.18 with window Hurst). The lift, in short, amplifies the combinatorial signal and purges its fractal loading. Figure 9 condenses the section.




Figure 9. Beyond fractality. (a) On fractional Brownian motion, directional dispersion tracks the Hurst exponent while the varifold scaling exponent D_{eff} is flat. (b) The stress effect on the primary endpoint with 95% HAC intervals, unconditional and under fractal and moment controls. (c) Densities of the commodity block D_{eff} in calm and stress.

6.5. Robustness Program

Figure 10 and Table 5 summarize the specification battery on the primary endpoint. The stress-to-calm difference of D is positive with bootstrap intervals excluding zero in all twelve kernel and normalization configurations: aspect parameters 0.15, 0.30, and 0.60; single scales $\sigma \in \{0.01, 0.02, 0.05, 0.10, 0.20\}$; the multiscale benchmark and its three scale extension; the $\cos^2$ direction kernel; within window detrending; and z-score value normalization. It survives window lengths of 60 and 250 days, the parameter-free horizontal visibility variant [3], and weekly sampling, where both component channels retain sign and significance. Dropping any single commodity preserves the effect, with the strongest value when gold, itself a safe haven [27,28], is removed from the commodity block; the VIX threshold response is monotone (differences 0.041, 0.068, 0.100 at thresholds 25, 30, 35); both the 2000 to 2012 and the 2013 to 2026 subsamples carry the effect, the early one marginally ($p = 0.061$) and the late one strongly ($p < 10^{-4}$). The time-shift placebo row sits at zero. The scale profile of the contrast peaks at intermediate to fine kernel scales, between $\sigma = 0.05$ and $\sigma = 0.10$, and decays toward the coarsest scale (Figure 11); over-smoothing was in fact the reason a preliminary single-scale analysis at $\sigma = 0.1$ with the $\cos^2$ kernel understated the entire phenomenon. The profile should not be read as parameter sensitivity of the conclusion: the contrast is positive with intervals excluding zero at every scale in the sweep, and the location of its maximum is itself a substantive finding about where in scale space the reconfiguration lives, a region that the multiscale benchmark spans.

Table 5. Specification grid for the primary endpoint: stress-to-calm difference of D with moving-block bootstrap 95% intervals and Mann and Whitney p values on the stratified window set (124 stress, 124 calm).

Specification	Difference	95% Interval	p
benchmark ($\sigma \in \{0.02, 0.05\}$, von Mises, minmax)	0.068	[0.048, 0.088]	$<10^{-4}$
aspect 0.15	0.058	[0.037, 0.079]	$<10^{-4}$
aspect 0.60	0.055	[0.038, 0.071]	$<10^{-4}$
$\sigma = 0.01$	0.035	[0.024, 0.045]	$<10^{-4}$
$\sigma = 0.02$	0.054	[0.039, 0.070]	$<10^{-4}$
$\sigma = 0.05$	0.075	[0.053, 0.097]	$<10^{-4}$
$\sigma = 0.10$	0.074	[0.052, 0.096]	$<10^{-4}$
$\sigma = 0.20$	0.063	[0.043, 0.084]	$<10^{-4}$
three scales $\sigma \in \{0.02, 0.05, 0.10\}$	0.070	[0.049, 0.091]	$<10^{-4}$
$\cos^2$ direction kernel	0.057	[0.036, 0.078]	$<10^{-4}$
detrended windows	0.060	[0.039, 0.082]	$<10^{-4}$
z-score normalization	0.078	[0.057, 0.099]	$<10^{-4}$
window 60 days	0.072	[0.042, 0.104]	10^{-4}
window 250 days	0.067	[0.037, 0.098]	2×10^{-4}
horizontal visibility graph	0.068	[0.051, 0.085]	$<10^{-4}$
time-shift placebo	0.005	[−0.014, 0.022]	0.059




Figure 10. Forest view of the specification grid and ablations for the primary endpoint, with the time-shift placebo in red at zero.




Figure 11. Scale profile: the stress-to-calm difference of D as a function of the positional kernel scale σ , with the 95% bootstrap confidence band shown as the shaded region. The contrast lives at intermediate to fine geometric scales.

Two further checks address the event-limited nature of the sample directly (Table 4). First, a leave-one-crisis-out analysis removes, one at a time, all windows within one year of each of the five canonical onsets and re-estimates the stress effect on the remainder: the coefficient retains its sign and significance in every case, including the removal of the Lehman episode, which eliminates a third of the stress windows. No single crisis carries the result. Second, the binary regime classification is dispensable: regressing D on the standardized level of the VIX across all 1307 windows, with no thresholds and no window selection, yields a continuous dose-response with $t = 3.22$ ($p = 0.0013$), and a logarithmic specification agrees. The contrast is a monotone function of stress intensity, not an artifact of where the regime lines are drawn. The same table reports the contrast built against each safe haven currency separately, the yen alone and the franc alone, both of which preserve

the effect in contrast form. The franc-alone contrast is significant only through its coupling leg, which rises in stress independently of the reference safe haven, not through any franc decoupling; this is consistent with the managed regime neutrality of the franc reported in Section 6.1.

6.6. External Validity and Specificity: An Equity Sector Panel

A second panel tests both the generality of the geometric machinery and the specificity of the contrast. Nine US-sector exchange-traded funds over 1999 to 2026 (6898 trading days, 1356 windows under the identical protocol) are assigned the same three roles: energy and materials as the shock origin, industrials, consumer discretionary, and financials as the exposed block, and the classic defensive sectors, consumer staples, utilities, and health care, in the seat of the safe block. What generalizes out of asset class is the stress sensitivity of geometric alignment itself, and it generalizes for every block: in risk-off windows the alignment of the exposed block with the shock block rises from 0.416 to 0.547 ($t = 3.29$, $p = 0.0010$), and the alignment of the defensive block with the shock block rises almost identically, from 0.293 to 0.438 ($t = 3.28$, $p = 0.0010$). Precisely because both blocks respond, the contrast does not arise: the equity analogue of D is flat in stress (-0.014 , $t = -0.42$, $p = 0.67$). This is the prediction of the mechanism, not a failure of it. Defensive sectors are low-beta claims on the same equity factor and import the shock geometry along with every other sector; they are not funding instruments whose flows reverse in risk-off. The contrast therefore requires genuine safe haven assets and does not materialize from an arbitrary three-block partition. In the same spirit as the time-shift placebo, this closes the most natural mechanical objection to the design.

7. Discussion

The findings speak to several active debates. The connectedness and spillover literature has converged on the view that cross-asset linkages intensify in stress [34,35], and our coupling results agree on the side where second moments are informative, the commodity-to-commodity–currency channel. The contribution is to show that this intensification has a geometric counterpart that is channel specific: on the safe haven side, where correlation reports no change, the shapes nonetheless decouple, so the standard reading of crisis comovement is incomplete for the assets that matter most for diversification. The topological data analysis literature has argued that the geometry of state space carries early warning content that moment-based methods miss [36,37]; our results support the geometric premise while differing in instrument, since the varifold lift is differentiable, channel-resolved, and, as the controls show, not reducible to the persistence summaries those methods use. Finally, the recent demonstration that visibility geometry can be monetized as a trading indicator [9] and the fractal market account of horizon collapse in crises [40] together frame our result as the structural, cross-asset counterpart to a single-series, decision-oriented use of the same raw construction. What the study confirms, then, is that the visibility graph carries economically meaningful information beyond its fractal skeleton, and that this information is most visible exactly where the conventional toolkit is blind.

7.1. Mechanism

The interaction structure, the channel anatomy, and the placebo benchmark jointly support a flight-to-safety reading with a geometric signature. Commodity shocks impose a characteristic visibility geometry, long, steep, directionally concentrated edges, on the windows of commodity prices. Terms of trade exposure transmits that geometry to the currencies of commodity exporters, whose alignment with the commodity block rises far above the independence benchmark precisely when global risk appetite collapses. The yen,

the funding leg of the carry trade [25,26], does not import the shock geometry. Its alignment stays near the level attained by a time-shifted surrogate, and this is the operational meaning of safe haven neutrality in our framework. The franc behaves as a managed currency over a large part of the sample and contributes only through the contrast form. Gold, sitting in both the commodity and the safe haven worlds [27,28], dilutes the commodity block and sharpens the effect when removed. None of this is visible to linear or rank correlation on the decoupling side, in line with the synthetic demonstration that sign-free, volatility-timed dependence is exactly the blind spot of second moments.

7.2. Practical Interpretation

The contrast has a direct reading for market analysis and risk monitoring, which we state without overstating it. A rising contrast is a real-time signal that the currencies of commodity exporting economies are importing the geometry of the commodity shock, their price trajectories taking on the same shape as the commodity complex, while the yen retains a shape of its own. For a risk monitor, this is a structural early picture of flight to safety that does not require the move to have already shown up in returns or in linear correlations. For portfolio interpretation, it says when the geometric diversification between commodity currencies and the commodity complex is collapsing, the two groups becoming shape aligned, even as the conventional correlation between them may still look unexceptional; the yen, by contrast, continues to supply genuine geometric diversification in exactly those windows. Because the contrast is sign-free and shape-based, it flags this regime through the channel that second-moment risk measures are constructed to miss.

We are deliberately careful about what this does and does not deliver. The signal is diagnostic, contemporaneous to slightly lagging relative to implied volatility, and our own hedging exhibit returned a null, so we do not present the contrast as a trading rule. This is the complement of the recent demonstration that visibility geometry can be turned into a profitable indicator [9]: that work builds a decision rule on a single series, whereas our contribution is a cross-asset structural diagnostic of how the geometry of a market system reorganizes under stress. The two uses are compatible, and connecting a structural diagnostic of this kind to an executable rule is a natural next step that we leave open.

7.3. Relation to Classical Network Statistics and Alternative Descriptors

A reader grounded in the classical visibility graph literature will ask how the layers compare in the standard network statistics of that tradition, each computed for a single series in isolation [1,4,5]. Table 6 reports, for the full sample natural visibility graph of every layer, the mean degree, the exponent of a power law fit to the degree distribution, the degree assortativity, and the global clustering coefficient. The statistics separate the blocks by type. Commodities have the heaviest tailed degree distributions, with fitted exponents near 1.74, and the lowest clustering, near 0.23. The currencies are more homogeneous and more clustered, with exponents near 1.86 and clustering near 0.28, a range that agrees with the exchange rate values reported by [4]. Same-type series sit close together. The two safe havens share almost identical exponents, 1.883 for the yen against 1.880 for the franc, and identical clustering, 0.267 against 0.266, and the five commodity currencies occupy a tight region of the statistic space. One within-type contrast is informative for our narrative. The franc has by far the lowest assortativity of any layer, 0.131 against 0.316 for the yen, together with the single largest hub degree. This is the network signature of the long flat stretches that the managed regime floor imposed on the franc, and it corroborates, from the combinatorial side, the franc neutrality discussed in Section 6.1.

These per-series statistics also clarify why a cross-type comparison, of a currency with a commodity, is well posed in our framework even though it would be meaningless for

raw price levels. We never compare magnitudes. The alignment (6) is a normalized inner product of geometric measures, invariant by construction to the type, scale, and units of each series, so the shape of a yen window and the shape of a copper window lie in the same Hilbert space and their alignment is defined exactly as it is for two currencies. More to the point, the classical statistics of Table 6 describe each series in isolation and are structurally unable to report whether two series move with a common geometry. The economically central question, whether commodity shocks reach the currencies of commodity exporters through the terms of trade [23,24], is irreducibly a cross-type question. The varifold lift exists precisely to make that question measurable, and the within-type comparisons above are reported alongside it for completeness, not as substitutes.

Table 6. Classical network statistics of the full sample natural visibility graph of each layer (6652 nodes): mean degree, exponent γ of a power law fit to the degree distribution, degree assortativity r , and global clustering coefficient C . Same-type series occupy similar regions of the table; the managed regime franc is the assortativity outlier.

Layer	Mean Degree	γ	Assortativity r	Clustering C
gold	19.45	1.559	0.328	0.183
silver	18.21	1.636	0.153	0.204
copper	16.61	1.732	0.275	0.235
WTI crude	16.16	1.792	0.327	0.249
natural gas	15.31	1.858	0.264	0.267
wheat	14.85	1.840	0.170	0.252
AUD	15.20	1.837	0.316	0.280
CAD	14.80	1.859	0.318	0.273
NOK	15.08	1.847	0.392	0.271
NZD	13.89	1.938	0.371	0.310
ZAR	15.49	1.808	0.408	0.255
CHF	14.60	1.880	0.131	0.266
JPY	14.77	1.883	0.316	0.267

The lift is one of several mathematical languages for the geometry of a time series, and we use the others as benchmarks rather than ignore them. The combinatorial degree channel of the classical literature [1,2,4] enters the analysis directly in Section 6.4, where a degree sequence alignment captures part of the effect while loading on fractality. Topological persistence [36,37] enters as a control in Section 6.4. Beyond these, the natural alternatives include recurrence networks [6], ordinal and permutation entropy representations [7], correlation filtered networks such as minimum spanning trees and planar maximally filtered graphs [29,32], and graph kernels that compare networks combinatorially [46]. Each captures a different facet of dynamical structure. The varifold lift is distinguished among them by retaining the embedded first-order geometry, position and direction, that the others discard or never form, and by supplying a single positive definite cross-series alignment. A systematic comparison of these descriptors on the same cross-type-alignment task is a natural direction for future work.

7.4. Negative Results

Three null findings shape the claims. First, supervised learning on the lift does not transfer across market eras at this sample size: a learnable kernel embedding decayed from a test AUC of 0.68 at initialization to 0.52 after training, a frozen embedding behaved identically (0.69 to 0.49), and a linear varifold machine in the spirit of [22] produced era-inverted out-of-sample classifications. The failure mode is informative rather than mysterious: the effective sample is roughly 55 non-overlapping windows spanning a

handful of stress regimes, and the sign structure of feature-to-outcome maps is regime-dependent, so decision rules fitted in one era invert in another, while representation-level statistics remain stable. The information is in the representation, not in fitted decision rules, which is why every claim in this paper rests on pre-specified statistics with dependence-robust inference rather than on trained models. Second, a hedging exhibit failed: the contrast does not forecast the next-month effectiveness of a minimum variance safe haven hedge of a commodity basket, whose average effectiveness is itself indistinguishable from zero; the contrast is a measurement of structure, contemporaneous to slightly lagging relative to the VIX, not a trading signal. The contribution is therefore diagnostic rather than prescriptive: the instrument identifies a mechanism, the channel-resolved geometry of flight to safety, and we leave decision-relevant applications as an open question rather than claiming them. Third, the raw decoupling ratio that motivated the study is partly mechanical, as our own surrogate analysis revealed; we regard the resulting reformulation, from a raw level to a benchmark-referenced contrast, as a template for placebo discipline in network measures of market structure, complementary in spirit to the topological literature [36,37].

7.5. Limitations

The binding constraint on every claim in this paper is the event structure of the sample: five canonical crises and roughly 55 non-overlapping windows bound what any analysis of these data can establish, and we accordingly read the evidence as a strong but event-limited signal rather than as settled fact. The leave-one-crisis-out and dose-response checks of Table 4 show that the signal is not the property of any single episode or of any regime threshold, which mitigates the boundary without removing it. The effect is also conditional and tail-like: it requires the coincidence of commodity turbulence and global risk-off, is weak at mild stress thresholds in raw form, and the early subsample is marginal on its own. The aspect normalization, although swept over a factor of four and replaced by a z-score variant without changing conclusions, remains a modeling choice, as does the window length. Sample sizes at the episode level are small, five canonical crises, and the permutation event study should be read accordingly. The franc result is confounded by a managed exchange rate regime for part of the sample. Commodity prices are front month continuous series without roll adjustment; the weekly frequency robustness and the detrended variant bound the plausible impact of roll effects, but a fully roll-adjusted reconstruction would be a useful extension. The equity sector replication establishes that the coupling response is not specific to one asset class and that the contrast is, by design, specific to genuine safe haven assets; replications on bond and credit panels, where flight to safety operates through yet another mechanism, are a natural next step. Finally, our fractal instruments operate at window length 120, where DFA and MF-DFA estimates are noisy; they are used as controls and descriptive separators, not as precise exponent estimates.

8. Conclusions

Visibility graphs are known to encode fractality; this paper measured what they encode beyond it. The varifold lift retains the embedded geometry that the adjacency matrix discards, embeds graphs of unequal size in one Hilbert space through a provably positive definite multiscale kernel, and supports a placebo-disciplined, channel-resolved measure of cross-market geometric alignment. On a quarter century of commodity and currency data, the resulting relative alignment contrast isolates a flight-to-safety signature: commodity currencies import commodity shock geometry while the yen remains at geometric independence. The signature is invisible to correlation on its defining channel, orthogonal to moment and fractal controls, stable across sixteen specifications and every single crisis

exclusion, and zero under a time-shift placebo. The accompanying decomposition, directional dispersion tracking the Hurst exponent while the varifold scaling exponent and the alignment contrast do not, gives precise content to the claim that visibility graph geometry contains fractal-driven and fractal-free components. The economically informative part of the stress response resides in the fractal-free component, with the mechanical share of raw alignment levels explicitly separated and discarded by the placebo benchmark.

We close by placing the contribution against the classical reading of the visibility graph. Standard per-series network statistics, reported here for every layer, separate currencies from commodities and place same-type series close together, but by construction, they describe one series at a time and cannot measure whether two series share a geometry. The varifold lift adds exactly that missing cross-series object, and the safe haven signature it isolates is the kind of finding that the classical statistics, however refined, are not built to express.

The substantive findings about commodities and safe haven currencies can be stated plainly. In stress, the currencies of commodity exporters take on the visibility geometry of the commodity complex, while the yen does not; gold and the yen are also the only two assets in the panel with positive average returns on high volatility days, confirming their safe haven role on the level of returns as well as of shape. The decoupling is scale-dependent: it lives at fine to intermediate kernel scales and fades at the coarsest scale, which is the geometric expression of the horizon collapse that the fractal market hypothesis associates with crises. We are careful about the portfolio reading. The geometric contrast identifies, in real time, when commodity currencies stop diversifying the commodity complex while the yen still does, which is information about the structure of portfolio risk; but a direct minimum variance hedging exhibit returned a null, so we document the scale-dependent structural role without claiming a realized hedging gain, and we treat the construction of an executable hedge from the signal as open.

Theoretically, the paper shows that the visibility graph carries economically meaningful information beyond the fractality its skeleton encodes, and it gives that statement an operational decomposition into fractal driven and fractal free components, aligning the multiscale design with the horizon-based logic of the fractal market hypothesis. Practically, the contrast offers risk monitors and analysts a structural diagnostic of flight to safety that activates through a channel second-moment measures are built to miss. The limitations are equally clear: the evidence rests on five canonical crises and roughly fifty-five non-overlapping windows, the franc is confounded by a managed regime, the fractal instruments are noisy at the chosen horizon, and the futures are front month series without roll adjustment. Future work should test the lift on bond and credit panels where flight to safety operates differently, pursue richer multifractal and pointwise regularity controls on longer windows, and connect the diagnostic to an executable rule in the spirit of recent visibility graph indicators [9].

Author Contributions: Conceptualization, M.A.B. and Ö.A.; methodology, M.A.B., Ö.A. and D.R.E.; software, M.A.B. and D.R.E.; validation, Ö.A. and L.G.; formal analysis, M.A.B., D.R.E. and Ö.A.; investigation, D.R.E.; data curation, D.R.E. and L.G.; writing, original draft preparation, M.A.B. and Ö.A.; writing, review and editing, L.G.; visualization, M.A.B.; supervision, Ö.A. and L.G. All authors have read and agreed to the published version of the manuscript.

Funding: This research received no external funding.

Institutional Review Board Statement: Not applicable.

Informed Consent Statement: Not applicable.

Data Availability Statement: All data are from public sources: front month futures, sector exchange-traded funds, and the VIX from public market data services, and exchange rates from the European Central Bank data portal. The code used in this study is available at <https://github.com/ADUSlab/vg-varifold> (accessed on 8 July 2026).

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

References

1. Lacasa, L.; Luque, B.; Ballesteros, F.; Luque, J.; Nuño, J.C. From time series to complex networks: The visibility graph. *Proc. Natl. Acad. Sci. USA* **2008**, *105*, 4972–4975. [CrossRef] [PubMed]
2. Lacasa, L.; Luque, B.; Luque, J.; Nuño, J.C. The visibility graph: A new method for estimating the Hurst exponent of fractional Brownian motion. *EPL* **2009**, *86*, 30001. [CrossRef]
3. Luque, B.; Lacasa, L.; Ballesteros, F.; Luque, J. Horizontal visibility graphs: Exact results for random time series. *Phys. Rev. E* **2009**, *80*, 046103. [CrossRef] [PubMed]
4. Yang, Y.; Wang, J.; Yang, H.; Mang, J. Visibility graph approach to exchange rate series. *Physica A* **2009**, *388*, 4431–4437. [CrossRef]
5. Zou, Y.; Donner, R.V.; Marwan, N.; Donges, J.F.; Kurths, J. Complex network approaches to nonlinear time series analysis. *Phys. Rep.* **2019**, *787*, 1–97. [CrossRef]
6. Donner, R.V.; Zou, Y.; Donges, J.F.; Marwan, N.; Kurths, J. Recurrence networks: A novel paradigm for nonlinear time series analysis. *New J. Phys.* **2010**, *12*, 033025. [CrossRef]
7. Bandt, C.; Pompe, B. Permutation entropy: A natural complexity measure for time series. *Phys. Rev. Lett.* **2002**, *88*, 174102. [CrossRef] [PubMed]
8. Vamvakaris, M.D.; Pantelous, A.A.; Zuev, K.M. Time series analysis of S&P 500 index: A horizontal visibility graph approach. *Physica A* **2018**, *497*, 41–51. [CrossRef]
9. Rak, R. Visibility graphs can make money in financial markets. *arXiv* **2026**, arXiv:2605.01300. [CrossRef]
10. Mandelbrot, B. The variation of certain speculative prices. *J. Bus.* **1963**, *36*, 394–419. [CrossRef]
11. Cont, R. Empirical properties of asset returns: Stylized facts and statistical issues. *Quant. Financ.* **2001**, *1*, 223–236. [CrossRef]
12. Di Matteo, T. Multi-scaling in finance. *Quant. Financ.* **2007**, *7*, 21–36. [CrossRef]
13. 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. [CrossRef] [PubMed]
14. Kantelhardt, J.W.; Zschiegner, S.A.; Koscielny-Bunde, E.; Havlin, S.; Bunde, A.; Stanley, H.E. Multifractal detrended fluctuation analysis of nonstationary time series. *Physica A* **2002**, *316*, 87–114. [CrossRef]
15. Kristoufek, L.; Vosvrda, M. Commodity futures and market efficiency. *Energy Econ.* **2014**, *42*, 50–57. [CrossRef]
16. Allard, W.K. On the first variation of a varifold. *Ann. Math.* **1972**, *95*, 417–491. [CrossRef]
17. Federer, H. *Geometric Measure Theory*; Springer: Berlin/Heidelberg, Germany, 1996. [CrossRef]
18. Vaillant, M.; Glaunès, J. Surface matching via currents. In *Information Processing in Medical Imaging*; Lecture Notes in Computer Science; Springer: Berlin/Heidelberg, Germany, 2005; Volume 3565, pp. 381–392. [CrossRef] [PubMed]
19. Charon, N.; Trounev, A. The varifold representation of nonoriented shapes for diffeomorphic registration. *SIAM J. Imaging Sci.* **2013**, *6*, 2547–2580. [CrossRef]
20. Kaltenmark, I.; Charlier, B.; Charon, N. A general framework for curve and surface comparison and registration with oriented varifolds. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR), Honolulu, HI, USA, 21–26 July 2017*; IEEE: New York, NY, USA, 2017; pp. 3346–3355. [CrossRef]
21. Lee, J.; Cai, X.; Schönlieb, C.-B.; Masnou, S. Neural varifolds: An aggregate representation for quantifying the geometry of point clouds. *arXiv* **2024**, arXiv:2407.04844. [CrossRef]
22. Hartman, E.; Charon, N. SVarM: Linear support varifold machines for classification and regression on geometric data. *arXiv* **2025**, arXiv:2506.01189. [CrossRef]
23. Rinaldo, A.; Söderlind, P. Safe haven currencies. *Rev. Financ.* **2010**, *14*, 385–407. [CrossRef]
24. Habib, M.M.; Stracca, L. Getting beyond carry trade: What makes a safe haven currency? *J. Int. Econ.* **2012**, *87*, 50–64. [CrossRef]
25. De Bock, R.; de Carvalho Filho, I. The behavior of currencies during risk-off episodes. *J. Int. Money Financ.* **2015**, *53*, 218–234. [CrossRef]
26. Brunnermeier, M.K.; Nagel, S.; Pedersen, L.H. Carry trades and currency crashes. *NBER Macroecon. Annu.* **2008**, *23*, 313–348. [CrossRef] [PubMed]
27. Baur, D.G.; Lucey, B.M. Is gold a hedge or a safe haven? An analysis of stocks, bonds and gold. *Financ. Rev.* **2010**, *45*, 217–229. [CrossRef]
28. Baur, D.G.; McDermott, T.K. Is gold a safe haven? International evidence. *J. Bank. Financ.* **2010**, *34*, 1886–1898. [CrossRef]
29. Mantegna, R.N. Hierarchical structure in financial markets. *Eur. Phys. J. B* **1999**, *11*, 193–197. [CrossRef]

30. Onnela, J.-P.; Chakraborti, A.; Kaski, K.; Kertész, J.; Kanto, A. Dynamics of market correlations: Taxonomy and portfolio analysis. *Phys. Rev. E* **2003**, *68*, 056110. [[CrossRef](#)] [[PubMed](#)]
31. Bonanno, G.; Caldarelli, G.; Lillo, F.; Mantegna, R.N. Topology of correlation-based minimal spanning trees in real and model markets. *Phys. Rev. E* **2003**, *68*, 046130. [[CrossRef](#)] [[PubMed](#)]
32. Tumminello, M.; Aste, T.; Di Matteo, T.; Mantegna, R.N. A tool for filtering information in complex systems. *Proc. Natl. Acad. Sci. USA* **2005**, *102*, 10421–10426. [[CrossRef](#)] [[PubMed](#)]
33. Kwapień, J.; Drożdż, S. Physical approach to complex systems. *Phys. Rep.* **2012**, *515*, 115–226. [[CrossRef](#)]
34. Diebold, F.X.; Yilmaz, K. Better to give than to receive: Predictive directional measurement of volatility spillovers. *Int. J. Forecast.* **2012**, *28*, 57–66. [[CrossRef](#)]
35. Adekoya, O.B.; Oliyide, J.A. How COVID-19 drives connectedness among commodity and financial markets: Evidence from TVP-VAR and causality-in-quantiles techniques. *Resour. Policy* **2021**, *70*, 101898. [[CrossRef](#)] [[PubMed](#)]
36. Gidea, M.; Katz, Y. Topological data analysis of financial time series: Landscapes of crashes. *Physica A* **2018**, *491*, 820–834. [[CrossRef](#)]
37. Guritanu, E.; Barbierato, E.; Gatti, A. Topological machine learning for financial crisis detection: Early warning signals from persistent homology. *Computers* **2025**, *14*, 408. [[CrossRef](#)]
38. Lacasa, L.; Nicosia, V.; Latora, V. Network structure of multivariate time series. *Sci. Rep.* **2015**, *5*, 15508. [[CrossRef](#)] [[PubMed](#)]
39. Peters, E.E. *Fractal Market Analysis: Applying Chaos Theory to Investment and Economics*; John Wiley & Sons: New York, NY, USA, 1994.
40. Kristoufek, L. Fractal markets hypothesis and the global financial crisis: Scaling, investment horizons and liquidity. *Adv. Complex Syst.* **2012**, *15*, 1250065. [[CrossRef](#)]
41. Bekaert, G.; Hoerova, M. The VIX, the variance premium and stock market volatility. *J. Econom.* **2014**, *183*, 181–192. [[CrossRef](#)]
42. Newey, W.K.; West, K.D. A simple, positive semi-definite, heteroskedasticity and autocorrelation consistent covariance matrix. *Econometrica* **1987**, *55*, 703–708. [[CrossRef](#)]
43. Mann, H.B.; Whitney, D.R. On a test of whether one of two random variables is stochastically larger than the other. *Ann. Math. Stat.* **1947**, *18*, 50–60. [[CrossRef](#)]
44. Politis, D.N.; Romano, J.P. The stationary bootstrap. *J. Am. Stat. Assoc.* **1994**, *89*, 1303–1313. [[CrossRef](#)]
45. Engle, R. Dynamic conditional correlation: A simple class of multivariate generalized autoregressive conditional heteroskedasticity models. *J. Bus. Econ. Stat.* **2002**, *20*, 339–350. [[CrossRef](#)]
46. Shervashidze, N.; Schweitzer, P.; van Leeuwen, E.J.; Mehlhorn, K.; Borgwardt, K.M. Weisfeiler-Lehman graph kernels. *J. Mach. Learn. Res.* **2011**, *12*, 2539–2561.

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.