



Econosysmographie

UNIVERSE RISK FRAMEWORK

EDUCATIONAL NOTE · INSTITUTIONAL ASSET MANAGERS & RISK OFFICERS · READING TIME ~7 MIN

Your Portfolio Lives in a Space You Never See

A journey through market states, diversification and the geometry of risk

Core insight. A portfolio does not exist in isolation: it lives inside an evolving market state. Same weights, same securities, same conviction — and yet, from a risk perspective, it may no longer be the same portfolio at all. This note asks how to rigorously measure the journey from one state to another, and introduces the Eratosthenes Test as one answer.

Your portfolio can remain completely unchanged from one day to the next. The same securities. The same weights. The same fundamental conviction. And yet, from a risk perspective, it may no longer be the same portfolio at all.

What changed is the geometry of the market around it. When assets begin to move in unison, diversification evaporates without a single trade being executed. A portfolio holding fifty distinct securities can suddenly collapse into just a handful of effective bets.

Core insight

A portfolio does not exist in isolation. It lives inside an evolving **market state**. If so, **how do we rigorously measure the journey from one state to another?**

This question leads directly to the geometry of correlation matrices, and to the **Eratosthenes Test**.

1. A portfolio is more than its weights

Modern portfolio analysis typically starts with the allocation vector:

$$\mathbf{w} = (\mathbf{w}_1, \dots, \mathbf{w}_n)^\top$$

Weights alone, however, reveal nothing about real-time diversification. Let $\mathbf{r}_t$ be the vector of asset returns and $\Sigma_t = \text{Cov}(\mathbf{r}_t)$ its covariance matrix at time t , decomposed as:

$$\Sigma_t = \mathbf{D}_t \mathbf{C}_t \mathbf{D}_t$$

where $\mathbf{D}_t$ is the diagonal matrix of individual asset volatilities and $\mathbf{C}_t$ is the correlation matrix. Portfolio variance is then:

$$\sigma^2_{p,t} = \mathbf{w}^\top \Sigma_t \mathbf{w} = \mathbf{w}^\top \mathbf{D}_t \mathbf{C}_t \mathbf{D}_t \mathbf{w}$$

Even if $\mathbf{w}$ is fixed, portfolio risk shifts whenever individual volatilities ($\mathbf{D}_t$) or the dependence structure ($\mathbf{C}_t$) change. The matrix $\mathbf{C}_t$ captures which assets move together, and therefore how much true diversification remains.

$$\text{Portfolio at time } t = (\mathbf{w}, \mathbf{C}_t)$$

2. The market moves when you stand still

Consider a diversified portfolio holding tech equities, government bonds, gold and commodities. In calm markets these assets occupy relatively orthogonal directions. When macro stress strikes, correlations surge: previously uncorrelated assets align under systemic liquidity pressure. Your weights did not change, but the state did:

$$\mathbf{C}_{t-1} \rightarrow \mathbf{C}_t, \quad \mathbf{C}_{t-1} \neq \mathbf{C}_t$$

- A **market regime** is an ex-post classification: calm, stressed, inflationary, risk-on, risk-off.
- A **market state**, in this framework, is the raw dependence structure itself: $\mathbf{C}_t$.

Treating $\mathbf{C}_t$ as the state allows us to monitor the structural path continuously:

- Where is the portfolio living today?
- Is the dependence structure contracting or expanding?

- How many effective structural directions remain?
- Is today's state historically anomalous relative to its own baseline?

3. Correlation states and the choice of ruler

A valid full-rank correlation matrix is symmetric, positive definite, and has unit diagonal. It therefore belongs to the open ellipsope, a constrained subset of $\text{SPD}(n)$, the open convex cone of symmetric positive-definite matrices:

$$\text{SPD}(n) = \{ C \in \mathbb{R}^{n \times n} : C = C^T, \mathbf{x}^T C \mathbf{x} > 0 \ \forall \mathbf{x} \neq \mathbf{0} \}$$

In this note, $\text{SPD}(n)$ is used as a diagnostic space for comparing observed correlation states. Several metrics can be placed on it, flat or curved, and when market dependencies move from A to B, the chosen metric determines the measured length of the trajectory. We compare two of them.

The flat Euclidean (Frobenius) distance:

$$d_F(A, B) = \|B - A\|_F = \sqrt{(\sum_{i,j} (B_{ij} - A_{ij})^2)}$$

The affine-invariant Riemannian metric (AIRM):

$$d_R(A, B) = \| \log(A^{-1/2} B A^{-1/2}) \|_F = \sqrt{(\sum_{i=1 \dots n} \log^2(\lambda_i))}$$

where λ_i are the generalized eigenvalues of the pair (A, B) , i.e. the eigenvalues of $A^{-1}B$.

Two different rulers measuring the exact same market transition can return radically different lengths.

4. Why the ruler matters near systemic stress

The discrepancy is most visible in a simple two-asset system, with a single correlation ρ between the two assets. For a correlation move from ρ_1 to ρ_2 :

$$\text{Flat distance: } d_F = \sqrt{2} \cdot |\rho_2 - \rho_1|$$

$$\text{Affine-invariant distance: } d_R = \sqrt{(\log^2((1 + \rho_2) / (1 + \rho_1)) + \log^2((1 - \rho_2) / (1 - \rho_1)))}$$

The stress paradox

A shift from $\rho = 0.10$ to 0.30 and a shift from $\rho = 0.70$ to 0.90 are both $\Delta\rho = 0.20$ on a flat ruler: $d_F = 0.28$ in both cases. On the affine-invariant ruler, the first move measures $d_R \approx 0.30$; the second measures $d_R \approx 1.10$, almost four times longer. As $\rho \rightarrow 1$, the term $(1 - \rho) \rightarrow 0$ and the logarithm blows up.

Moving from moderate to high correlation degrades diversification; moving from high correlation to near-singularity destroys an entire dimension of diversification. A ruler that measures differences in correlation treats both changes as identical; a ruler that measures ratios, on a logarithmic scale, registers the collapse. In this two-asset case the effect comes from that logarithmic scale alone: both states share the same principal directions.

5. Tracking effective dimensionality

The nominal count of positions in a portfolio is rarely its effective dimension. Two complementary measures apply.

Effective positions by weight (a property of the allocation w):

$$N_{\text{eff,pos}} = 1 / \sum_i w_i^2$$

Effective structural directions (a property of the market state C_t , via the entropy of its spectrum):

$$N_{\text{eff,dir}} = \exp(- \sum_i \tilde{\lambda}_i \log \tilde{\lambda}_i)$$

where $\tilde{\lambda}_i = \lambda_i / \sum_j \lambda_j$ is the normalised eigenspectrum of C_t .

A portfolio holding 50 securities may show 35 effective positions by weight, while the universe it invests in offers only 10 to 12 effective market directions under stress. Monitoring $N_{\text{eff,dir}}$ reveals silent structural crowding that nominal position counts may obscure.

6. The Eratosthenes Test: measuring ruler disagreement

The **Eratosthenes Test** quantifies the relative gap between the flat Euclidean measurement and the affine-invariant Riemannian distance of an observed transition $A \rightarrow B$:

$$\kappa = (d_R(A, B) - d_F(A, B)) / d_F(A, B)$$

Over an observed sequence of states, the cumulative geometry gap is:

$$U = 1 - \sum_t d_F(C_{t-1}, C_t) / \sum_t d_R(C_{t-1}, C_t)$$

κ reads as “how much longer the affine-invariant measurement is than the flat one” for a single transition; U reads as the share of the total affine-invariant length on which the two rulers disagree over a whole period.

WHAT THE ERATOSTHENES GAP (κ , U) IS

WHAT IT IS NOT

A measure of how much two rulers disagree

Not a VaR correction factor

An indicator of metric sensitivity under stress

Not an expected-loss calculation

A model-risk test for stress scenarios

Not a directional market-timing signal

7. Distance vs. direction: detecting structural contraction

Because distance is symmetric ($d(A, B) = d(B, A)$), length alone cannot tell whether dependence contracted or expanded. Direction is recovered from the logarithm of the step, taken in the frame of the starting state:

$$V = \log(A^{-1/2} B A^{-1/2})$$

Its trace reveals the global volumetric change of the dependence ellipsoid:

$$\text{tr}(V) = \log \det(B) - \log \det(A)$$

- $\text{tr}(V) < 0$: contraction of the dependence space (diversification collapsing).

- $\text{tr}(V) > 0$: re-expansion of effective structural directions.

8. Empirical evidence across regimes

Calibration on a demonstration universe of 54 US-listed exchange-traded funds spanning US and world sectors, developed- and emerging-market countries, government bonds, credit, gold and commodities, daily data from 2020-01-03 to 2026-09-14, with regimes labelled by quartiles of the equal-weight book's own realised volatility (calm = lowest quartile, stressed = highest), gives the following medians of κ :

MARKET REGIME	21-SESSION STATE WINDOW (MEDIAN κ)	63-SESSION CONTROL WINDOW (MEDIAN κ)
Calm	41%	23%
Normal	46%	25%
Stressed	55%	44%

Absolute values of κ depend on the estimation window and on matrix conditioning, but the ordering is invariant across calibrations:

$$\kappa_{\text{calm}} < \kappa_{\text{normal}} < \kappa_{\text{stressed}}$$

Across both calibrations, the disagreement between the flat and the affine-invariant rulers is larger in stressed states than in calm states, indicating greater metric sensitivity precisely when dependence structures are more compressed.

9. Scenario admissibility and model risk

When risk teams construct hypothetical stress tests, linear shocks are standard:

$$C_{\text{flat}}(t) = A + t(B - A)$$

Positive definiteness is preserved for interpolation ($0 \leq t \leq 1$), but extrapolating beyond recent experience ($t > 1$) frequently produces matrices with negative eigenvalues, which are not valid correlation matrices.

Conversely, an unconstrained Riemannian geodesic on $\text{SPD}(n)$ preserves positive definiteness but lets the diagonal entries drift away from 1, leaving the manifold of valid correlation matrices (the open ellipsope). Geometries built directly on that manifold (Thanwerdas & Pennec, 2022) are the rigorous answer.

Model-risk takeaway

Before debating the loss generated by an extreme scenario, risk committees should ask a prior question: **is the underlying correlation state geometrically admissible?**

Diagnostic summary: seeing the state you live in

TRADITIONAL RISK FOCUS

MARKET-STATE INTELLIGENCE LAYER

What are our nominal exposures?

What market state does the portfolio occupy today?

What is our parametric or historical VaR?

Is the underlying dependence structure contracting?

How did the portfolio perform in 2008 / 2020?

How many effective risk directions actually remain?

Linear factor stress testing

Are scenario correlation matrices admissible and positive definite?

Request an Eratosthenes Market-State Review

The Eratosthenes diagnostic can benchmark the market state of an investment universe without ingesting internal positions. For a portfolio-specific Market-State Review, controlled portfolio weights can be incorporated to quantify allocation and structural diversification measures.

Book a 30-minute Discovery Session at econosystemographe.com · contact@econosystemographe.eu

Selected references

- Münnix, M. C. et al. (2012). Identifying States of a Financial Market, *Scientific Reports* 2, 644.
- Thanwerdas, Y. & Pennec, X. (2022). Theoretically and computationally convenient geometries on full-rank correlation matrices, *SIAM Journal on Matrix Analysis and Applications* 43(4).
- Bhatia, R. (2007). *Positive Definite Matrices*, Princeton University Press.
- Higham, N. J. (2002). Computing the nearest correlation matrix — a problem from finance, *IMA Journal of Numerical Analysis* 22(3), 329–343.

This note is for educational purposes only and does not constitute investment advice.