WORKING PAPER

Geometry over Taxonomy in Portfolio Construction

The End of Asset Classes as Primitives

Clear Fog Capital

May 2026

ABSTRACT

Portfolio construction is conventionally organized around asset class taxonomy: equities, fixed income, commodities, and their subcategories. This paper argues that taxonomy is not a structural primitive — it is a naming convention. The return-generating process is governed by statistical relationships between instruments, not by the labels assigned to them. Those relationships are time-varying, regime-dependent, and empirically measurable. We replace taxonomy with geometry: specifically, the eigenstructure of the correlation matrix and the metrics derived from it. The central quantity is the effective number of independent bets, N_{eff} — a geometric measure of diversification that adapts to correlation regimes. A taxonomy-based framework cannot adapt. The empirical record on a representative sixteen-asset institutional universe reveals that regime shifts are not uniform: in some episodes, cross-asset dispersion expands and label-balance happens to align with geometric balance; in others, the correlation structure compresses and label-balance ceases to correspond to diversification. The label framework cannot distinguish the two. The geometric framework can. The implication is direct: allocation should be constructed across independent directions in return space, not across institutional categories.

1. Introduction

The standard architecture of portfolio construction begins with a taxonomy. Equities. Fixed income. Commodities. Alternatives. These categories are treated as primitives — as if the label itself carried structural content, as if naming a thing were the same as understanding its behavior.

It is not.

Asset class labels are descriptive artifacts. They reflect how markets are reported, how mandates are written, how institutions are organized. They do not reflect how returns are generated. The return-generating process does not consult the label. It operates through the statistical co-movement of instruments — a structure that is continuous, time-varying, and indifferent to category.

The practical consequences of this confusion are not minor. A portfolio constructed to hold "equities, bonds, and commodities" achieves categorical balance, not geometric diversification. Whether the two coincide is a regime-dependent question, not a structural guarantee. Sometimes the labels happen to track the geometry. Sometimes they do not. A label-based framework cannot tell which case it is in.

This paper develops the alternative. We replace taxonomy with geometry. The correlation matrix defines a metric space over assets. Diversification is a property of that space. It can be measured, monitored, and constructed — independently of any label system.

The argument proceeds as follows. Section 2 establishes the formal framework. Section 3 develops the geometric quantities. Section 4 documents the empirical record on a representative institutional universe and identifies the regimes in which taxonomy and geometry diverge. Section 5 draws the implications for portfolio construction. Section 6 synthesizes the position. Section 7 concludes.

---

2. Formal Framework

Let $r_t \in \mathbb{R}^N$ denote the vector of asset returns at time t, where N is the number of instruments. Define the covariance matrix:

$$\Sigma = \mathbb{E}\left[(r_t - \mu)(r_t - \mu)^\top\right], \qquad \mu = \mathbb{E}[r_t]$$

and the correlation matrix:

$$C_{ij} = \frac{\Sigma_{ij}}{\sqrt{\Sigma_{ii}\,\Sigma_{jj}}}$$

All diversification-relevant structure is contained in C. The covariance matrix Σ encodes both magnitude (volatility) and direction (co-movement). For the purposes of structural diversification, direction is what matters: C is sufficient.

Asset class labels do not appear in this formulation. This is not an omission. It is the point.

In practice, Σ and C are unobservable and must be estimated. We use a rolling estimator over a window of length T:

$$\hat{\Sigma}_t = \frac{1}{T-1} \sum_{s=t-T}^{t-1} (r_s - \hat{\mu}_t)(r_s - \hat{\mu}_t)^\top$$

The choice of T governs the regime sensitivity of the estimate. Short windows respond quickly but introduce noise. Long windows are stable but lag. This tension is not resolved by taxonomy. It is a property of the estimation problem, and it applies regardless of how assets are labelled.

In practice we use the exponentially weighted variant:

$$\hat{\Sigma}_t = \lambda\, \hat{\Sigma}_{t-1} + (1-\lambda)\, r_t r_t^\top, \qquad \lambda = \left(\tfrac{1}{2}\right)^{1/h}$$

where h is the half-life in trading days. EWMA decays smoothly rather than truncating at a hard window edge, removing the artifacts produced when a single observation drops out of an equal-weighted window. The institutional convention (RiskMetrics, DCC) is built on the same construction. Throughout, the primary estimator uses $h = 60$ days; the empirical record in Section 4 is shown to be qualitatively unchanged for $h \in [30, 120]$.

More sophisticated estimators address noise and persistence directly. Shrinkage methods (Ledoit & Wolf, 2004) reduce estimation error in the sample covariance toward a structured target. Dynamic conditional correlation models (Engle, 2002) parameterize the time variation explicitly. These refinements operate within the estimation problem; they do not change which object is being estimated. The geometry is in C, regardless of how it is recovered.

---

3. *Geometry of Returns*

The correlation matrix C induces a geometry over the asset universe. Diversification, properly understood, is a geometric property of that induced space — not a consequence of holding instruments from different categories.

3.1 Average Correlation

The simplest geometric summary is the average pairwise correlation:

$$\bar{\rho} = \frac{2}{N(N-1)} \sum_{i<j} C_{ij}$$

$\bar{\rho}$ provides a first-order measure of global dependence. Low $\bar{\rho}$ implies a dispersed structure; high $\bar{\rho}$ implies a compressed one. A portfolio constructed over a high-$\bar{\rho}$ universe holds many instruments but few independent bets.

3.2 Eigenstructure

The spectral decomposition of C provides the complete geometric picture. Let $\lambda_1 \geq \lambda_2 \geq \cdots \geq \lambda_N$ denote the eigenvalues of C. The fraction of variance explained by the dominant eigenvector is:

$$\rho_{\text{PCA}} = \frac{\lambda_1}{\sum_{i=1}^{N} \lambda_i}$$

When ρ_{PCA} is high, one factor dominates the return space. All assets move together. Adding instruments adds exposure — not independence. When ρ_{PCA} is low, variance is distributed across many eigenvectors. Independent structure exists.

This is the geometry of crowding. A market in which ρ_{PCA} is rising is a market in which the available diversification is being compressed — irrespective of how many different labels are represented.

The empirical eigenstructure of financial correlation matrices is well-studied. Most eigenvalues are statistically indistinguishable from those of a random matrix; only a small number — typically the largest — carry meaningful structural information (Laloux et al., 1999). The same conclusion appears in portfolio applications: the apparent diversification benefits of holding many instruments derive from a small number of dominant directions (Pafka & Kondor, 2004).

3.3 Effective Number of Independent Bets

The full eigenvalue spectrum can be collapsed into a single, interpretable quantity. Following Meucci (2009), define:

$$N_{\text{eff}} = \frac{\left(\sum_{i=1}^{N} \lambda_i\right)^2}{\sum_{i=1}^{N} \lambda_i^2}$$

This is the Herfindahl-style effective count applied to the eigenvalue distribution. It measures how many independent directions exist in return space — not how many assets are held. It is a property of the asset universe — the eigenstructure of C — not of any particular portfolio's weights.

N_{eff} ranges from 1 (single dominant factor, no diversification) to N (all eigenvectors equally weighted, maximum diversification). A portfolio of 20 assets with $N_{\text{eff}} = 3$ is not a diversified portfolio. It is three bets with 17 labels attached.

A related geometric measure, the diversification ratio of Choueifaty and Coignard (2008), captures the same insight through the gap between portfolio volatility and the weighted average volatility of its components.

This distinction is the operational heart of the geometric approach.

3.4 Distance and Clustering

The correlation matrix induces a natural distance metric between assets, due to Mantegna (1999):

$$d_{ij} = \sqrt{2(1 - C_{ij})}$$

This satisfies the axioms of a metric — non-negativity, symmetry, and the triangle inequality — and maps correlation structure directly into a geometric space. When $C_{ij} = 1$, assets are coincident: $d_{ij} = 0$. When $C_{ij} = -1$, they are maximally distant: $d_{ij} = 2$. Independent assets ($C_{ij} = 0$) are at distance $\sqrt{2}$.

Clustering methods applied under this metric — hierarchical linkage, minimum spanning trees, graph-theoretic methods — recover the empirical structure of the asset universe. The resulting clusters represent genuine behavioral similarity. They do not respect asset class boundaries. Equities cluster with credit. Certain commodities cluster with growth equities. Government bonds may cluster with nothing — or with everything, depending on regime.

Hierarchical risk parity (López de Prado, 2016) operationalizes this directly: portfolio construction proceeds from the cluster structure rather than from category labels.

The structure is empirical. The labels are not.

3.5 Time Variation

All geometric quantities are functions of time:

$$C = C(t), \qquad \lambda_i = \lambda_i(t), \qquad N_{\text{eff}} = N_{\text{eff}}(t)$$

Correlation regimes are not stationary. The geometry breathes with the market. The direction in which it breathes — toward compression or toward dispersion — is not given by theory. It is a regime-dependent empirical question, addressed directly in Section 4.

Taxonomy, by contrast, does not breathe at all. A fixed label system encodes a static picture of relationships that do not remain static. The model error introduced at the first step — treating labels as structural primitives — propagates forward through every subsequent construction decision, regardless of which direction the geometry happens to move.

---

4. The Empirical Record

The argument against taxonomy is not merely theoretical. It has a consistent empirical record. The relevant question is not whether the geometry shifts under stress — it does — but how, and whether label-balance and geometric balance remain aligned across those shifts.

4.1 Convergence — when, and when not

The conventional intuition is that diversification "fails under stress" because all assets begin to move together. The pattern is documented at the level of extreme tail co-movement (Longin & Solnik, 2001) and in the practitioner literature on diversification failure (Page & Panariello, 2018). The intuition is half right. Convergence does occur. It does not occur uniformly.

A label assigned at portfolio construction time reflects the statistical relationships of the estimation period. Those relationships are not guaranteed to persist. Under stress, the geometry changes. The label does not. But the *direction* of the change is regime-specific. Examined across the modern record, stress episodes sort cleanly into two distinct geometric modes. They are not variations of a single failure pattern. They are structurally different events.

We label them Regime A and Regime B.

4.2 Regime A — Dispersion

Regime A is the flight-to-quality regime. The trigger is a deflationary, credit, or demand shock. Equities and credit decline; long duration, gold, and reserve currencies rally. Different asset categories move in different directions and on different timescales. The correlation matrix, far from compressing, *expands*: average pairwise correlation falls toward zero, and variance distributes across multiple eigenvectors.

In Regime A:

$$\bar{\rho}\downarrow, \qquad N_{\text{eff}} \approx \text{baseline or higher}$$

The geometry, in this mode, is more diversified during stress than it is in calm. This is not a paradox. It reflects the fact that systemic deflationary shocks activate the very negative cross-asset correlations that taxonomy assumes are always present. Bonds rally precisely because equities fall.

The institutional consequence is that label-balance and geometric balance happen to align. A categorically diversified portfolio — equities offset by duration, credit offset by gold — receives the diversification its construction implies. Taxonomy works. Not because the framework is correct, but because the regime is the one for which the framework was implicitly calibrated.

4.3 Regime B — Compression

Regime B is the correlated risk-off regime. The trigger is an inflationary, monetary, or sovereign shock. The "diversifying" asset is the source of the stress, not the offset to it. Stocks and bonds sell off together. Gold, commodities, and credit move on the same factor. The cross-asset structure does not disperse; it concentrates.

The mechanism is empirically established. Campbell, Sunderam, and Viceira (2017) document that the sign of the stock-bond covariance is determined by the inflation regime: negative under deflationary conditions, positive under inflationary ones. The structural relationship between the diversifier and the risk asset is regime-conditional, not fixed. The deflationary case is the one institutional memory encodes as "normal."

In Regime B:

$$\bar{\rho}\uparrow \text{ or sign-flipped}, \qquad \rho_{\text{PCA}}\uparrow, \qquad N_{\text{eff}}\downarrow$$

This is the regime in which the conventional "diversification-fails-under-stress" intuition holds — and in which the taxonomy framework breaks. A portfolio constructed for label balance discovers that the labels were tracking a regime A correlation structure that no longer applies. The diversifier and the risk asset are now exposures to the same factor. The hedge is the position.

The institutional consequence is that label-balance ceases to correspond to geometric balance. The portfolio realizes higher correlation — and therefore higher portfolio variance — than its construction implied. This is not a tail event. It is the standard behavior of the system in Regime B.

4.4 Empirical Exhibit

To make these regimes concrete, we compute realized N_{eff} on a representative sixteen-asset institutional universe approximating a typical 60/30/10 allocation, broadened along two axes that label-based aggregates conceal: equity sub-factors (a growth/duration tilt and an energy/commodity tilt) and the rate/credit ladder. The estimation uses the EWMA correlation matrix described in Section 2 with a 60-day half-life, daily, from 2008 forward. All tickers use their native price histories.

The universe figure characterizes the eigenstructure of the asset space — what diversification is *available*. To make the institutional implication concrete we compute, alongside it, a portfolio-level N_{eff} for an illustrative 60/30/10 allocation: VTI 30%, EFA 12%, EEM 8%, IWM 5%, QQQ 3%, XLE 2% in equities; AGG 12%, IEF 5%, TLT 4%, LQD 3%, TIP 3%, HYG 2%, EMB 1% in fixed income; GLD 4%, VNQ 3%, DBC 3% in real assets. The portfolio measure is the Herfindahl on variance contributions across the principal components of the realized covariance — how many independent directions a *specific* portfolio loads on (Meucci, 2009), bounded above by the universe figure.

Plotted alongside universe N_{eff} in Figure 1, the result is unambiguous. Portfolio N_{eff} holds near 1.15 in calm and dips to 1.04 in stress. The gap to the universe figure is not a stress phenomenon — it is the standing arrangement. Equity volatility ($\approx$ 16% annualized) is several times bond volatility ($\approx$ 4–5%), so the equity sleeve dominates portfolio variance even at 60% capital weight; the bond and real-asset sleeves reduce *absolute* portfolio variance modestly but do not add independent directions in the variance decomposition. The picture is robust to weight perturbation: portfolio N_{eff} in calm holds at 1.15 across VTI weights from 25 to 35 percent, with the equity sleeve rebalanced accordingly.

Table 1. *Asset universe. Sixteen liquid US-listed ETFs spanning the conventional institutional taxonomy, broadened to include equity sub-factors and a fuller rate/credit ladder.*

Ticker	Role	Conventional label
VTI	US total market equity	Equities
IWM	US small-cap equity	Equities
QQQ	Nasdaq-100 (growth/tech)	Equities
XLE	US energy sector	Equities
EFA	Developed ex-US equity	Equities
EEM	Emerging markets equity	Equities
IEF	7–10y US Treasuries	Fixed income (rates)
TLT	Long-duration US Treasuries	Fixed income (rates)
TIP	US TIPS (inflation-linked)	Fixed income (rates)
AGG	Broad US fixed income aggregate	Fixed income
LQD	US investment-grade corporate	Fixed income (credit)
HYG	US high yield credit	Fixed income (credit)
EMB	EM sovereign debt (USD)	Fixed income (credit)
GLD	Gold	Commodities / Alternatives
DBC	Broad commodities	Commodities
VNQ	US listed REITs	Alternatives / Real assets

The taxonomy baseline is defined as the time-average of N_{eff} over the calm window 2012–2019: a value of approximately 3.54, against a maximum possible of 16. This is the implicit assumption embedded in a label-balanced framework — that approximately 3.5 independent directions are reliably available. The realized series records what was actually present. That sixteen labelled positions collapse to roughly three-and-a-half independent directions in calm is itself the first empirical observation: the labels overcount diversification by more than four-to-one. The additional tickers — sectors, sub-categories, geographies — load onto the existing eigenstructure rather than expanding it.

The baseline is not a threshold for whether taxonomy "works." It is a quantitative description of what a label-balanced framework implicitly assumes is reliably available. The load-bearing argument is not that N_{eff} varies — it is that the variation is invisible to the labels.

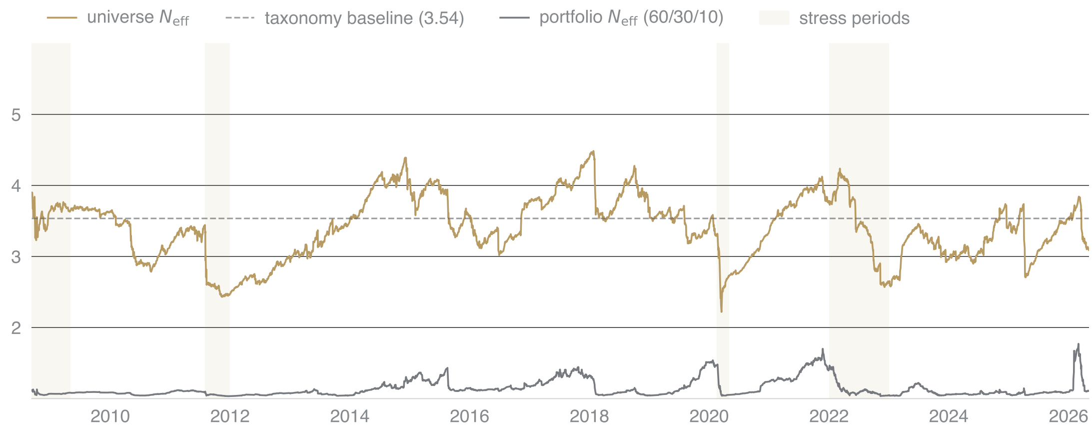

Figure 1. *Realized N_{eff}, 2008–2026. Sixty-day half-life EWMA. Solid gold: universe N_{eff} on the asset space (Table 1). Solid muted: portfolio N_{eff} for the illustrative 60/30/10 in §4.4. Dashed: taxonomy baseline (mean 2012–2019). Shaded bands: GFC (Sep 2008 – Apr 2009), Eurozone sovereign (Aug – Dec 2011), COVID (Feb – Apr 2020), 2022 inflation/rates (full year).*

Table 2. *Selected stress troughs.*

Episode	Regime	Universe N_{eff}	Portfolio N_{eff}	$\bar{\rho}$ at trough
GFC 2008–2009	A	3.23	1.05	0.14
Lehman vertex (Sep 2008)	A	3.55	1.10	0.07
Eurozone sovereign 2011	B	2.43	1.03	0.18
COVID March 2020	A → mixed	2.22	1.04	0.38
2022 inflation/rates	B	2.57	1.04	0.54

Calm baselines: universe $N_{\text{eff}} \approx 3.54$, portfolio $N_{\text{eff}} \approx 1.15$. The picture is robust to estimator choice: the 2022 universe trough ranges from 2.27 ($h = 30$) to 2.93 ($h = 120$); the regime ordering and the GFC/2022 contrast are preserved across the range.

Four observations follow from Figure 1 and Table 2.

First, the GFC — the canonical "diversification fails" event in the institutional memory — was a Regime A episode in the geometry. At the Lehman vertex itself, average pairwise correlation fell to 0.07 and universe N_{eff} sat essentially at the calm baseline (3.55 against 3.54). The broader GFC episode showed only a modest dip (trough $N_{\text{eff}} = 3.23$) — a small disturbance of the eigenstructure compared with the deep compressions later seen in Regime B. The portfolio constructed on label-balance was not penalized — but it could not have known why. The regime did the diversification work that the labels were assumed to be doing.

Second, the events at which universe N_{eff} collapsed substantially below baseline — Eurozone 2011, the 2022 inflation regime — were not the canonical crises. They were regimes in which the taxonomy's central diversifier (duration) became correlated with the risk asset it was meant to offset. In 2022, with $\bar{\rho} = 0.54$ and $N_{\text{eff}} = 2.57$, sixteen labelled positions held 2.6 independent directions against an implicit 3.5 — a 27 percent shortfall in realized diversification at the universe level.

Third, the structure of the realized universe series is visually obvious. The 2022 episode is the most sustained Regime B compression in the eighteen-year record — average pairwise correlation reached 0.54, far above the 0.38 at the COVID trough or the 0.18 in the Eurozone crisis. The

COVID compression was briefly deeper in absolute N_{eff} (2.22) but decayed within months; the 2022 reading held below baseline for the better part of a year. The GFC, despite its institutional memory, traced only a shallow dip. The series end — May 2026 — shows renewed compression: N_{eff} at approximately 3.1, ρ_{PCA} above 0.52, against a still-rising equity market. The framework identifies this regime; the labels do not.

Fourth, and most institutionally consequential, the portfolio figure barely tracks any of this. Portfolio N_{eff} sits near 1.1 throughout the eighteen-year record, dipping only to 1.04 in the worst stress periods. The series shows brief upticks within the COVID and 2022 bands — to roughly 1.4 and 1.5 respectively — when surging bond volatility momentarily raises the FI sleeve's share of portfolio variance; both episodes revert within months. The universe-level dynamics the geometry exposes are dynamics a typical 60/30/10 does not capture: the bond and real-asset sleeves reduce absolute portfolio variance modestly but do not deliver independent directions in the variance decomposition. The compression of 2022 was, for this portfolio, the closing of opportunities it was already not using. The "diversification" the label-balanced framework promises is a capital-allocation property, not a risk-allocation one.

The interpretation is structural, not retrospective. The taxonomy framework appears to have worked across the GFC because the correlation regime aligned with its assumptions at the universe level. It appears to have failed in 2022 because the correlation regime did not. At the portfolio level, the universe's apparent successes were already hollow: the realized variance decomposition shows one effective direction, regardless of which universe regime is active. In neither case did the labels themselves carry information about what the portfolio was actually doing. The geometry did.

4.5 False Independence and False Redundancy

Beyond regime-level effects, taxonomy produces two systematic errors at the cross-section.

The first is false independence: two assets in different categories that occupy the same cluster, load on the same eigenvector, and behave identically under the relevant stress scenarios. Holding both provides no additional independent direction. The label suggests diversification; the geometry refutes it.

The second is false redundancy: two assets within the same category that are structurally independent — with distinct factor loadings and low correlation across all regimes. A taxonomy-based portfolio that limits exposure to a category discards a genuine independent direction, reducing N_{eff} unnecessarily.

Both errors are invisible to a label-based framework. Both are directly observable in the geometry.

---

5. Implications for Portfolio Construction

5.1 Taxonomy Is Not a Primitive

Asset classes are institutional artifacts. They are naming conventions developed for reporting, regulatory, and mandate purposes. Used as structural inputs to portfolio construction, they introduce model error at the first step.

This is not a small technical objection. It is a foundational one. A framework built on the wrong primitive will produce systematically incorrect outputs, regardless of the sophistication applied in subsequent steps. The error is not in the optimization — it is in what is being optimized over.

5.2 The Allocation Space

Portfolio construction, under the geometric framework, becomes allocation across independent directions in return space. The relevant questions are not "how much equities, how much bonds" but:

- What is the current N_{eff} of the opportunity set?
- How is variance distributed across eigenvectors?
- Which clusters represent genuinely independent directions?
- How stable is the cluster structure across regimes?
- Which regime — dispersion or compression — is currently active?

These questions have quantitative answers. They can be monitored in real time. They respond to the actual structure of the market, not to an institutional classification scheme.

5.3 Risk Allocation as the Unit

The natural consequence of the geometric framework is a shift in the unit of allocation. Capital is not the relevant quantity. Risk is.

A portfolio that allocates equal capital across categories does not allocate equal risk — and does not achieve equal diversification — unless the geometry happens to be uniformly distributed. It rarely is. The correct unit is risk contribution per independent direction: a quantity that is defined geometrically and measured through the eigenstructure of the realized covariance.

This reframes the construction problem entirely. Equal risk contribution across independent eigenvectors is a well-defined objective. Equal capital contribution across label categories is not.

The risk-allocation literature is developed (Roncalli, 2013), as are Bayesian frameworks that combine equilibrium with active views (Meucci, 2010). Both typically operate over a predefined asset universe with assumed structural properties. The geometric approach makes those properties the object of estimation rather than assumption.

5.4 Dynamic Monitoring and Regime Identification

Because the geometry is time-varying, construction is not a one-time event. A portfolio that achieves high N_{eff} at construction time may see that N_{eff} collapse as correlation regimes shift. Monitoring must therefore be continuous: tracking $\bar{\rho}$, ρ_{PCA}, cluster stability, and N_{eff} against the portfolio's intended construction.

The two-regime structure documented in Section 4 has direct operational content. Regime-switching frameworks for asset allocation have a developed literature (Ang & Bekaert, 2002); the contribution of the geometric approach is to define the regimes over the eigenstructure rather than over labelled categories. Regime A and Regime B are distinguishable in real time through the joint behavior of $\bar{\rho}$ and N_{eff}, supplemented by the sign of the equity-duration correlation. A drawdown accompanied by falling $\bar{\rho}$ and stable N_{eff} is a Regime A event: the taxonomy is, coincidentally, providing the diversification it claims. A drawdown accompanied by rising $\bar{\rho}$ and falling N_{eff} is a

Regime B event: the taxonomy is failing, and position sizing must respond to the actual geometry rather than the labelled construction.

When the geometry compresses, the appropriate response is not to rebalance toward label targets. It is to recognize that the available independent directions have temporarily contracted, adjust position sizing accordingly, and preserve capital for the period when the geometry re-expands. This is not a trading decision. It is a structural one. It follows directly from the geometry.

---

6. Synthesis

The argument can be stated without qualification:

Asset class labels do not constitute structural primitives in the return-generating process. The correspondence between label-balance and geometric balance is regime-dependent. In Regime A, the two coincide and taxonomy appears to function. In Regime B, they diverge and taxonomy fails. The framework cannot distinguish these cases from inside its own categories. Portfolio construction based on taxonomy is structurally fragile not because labels are always wrong, but because they are sometimes right for reasons unrelated to the framework that produced them.

Therefore: taxonomy is abandoned as a primitive. Geometry is adopted as the basis for diversification.

The practical commitments that follow are:

- Measure diversification through N_{eff}, not through count of labels
- Monitor $\bar{\rho}$, ρ_{PCA}, and the equity-duration correlation jointly to identify the active regime
- Construct portfolios across independent geometric directions, not across categories
- Size positions by risk contribution per independent direction, not by capital allocation per label
- Respond to compression episodes by reducing gross exposure, not by rebalancing toward label targets
- Monitor continuously; the geometry does not stay where it was placed

---

7. Conclusion

The organization of markets is not categorical. It is geometric.

Labels describe. Geometry determines.

The correlation matrix does not care which sector an instrument belongs to, how it is classified in an index, or what mandate category a portfolio manager assigns it to. It records co-movement — which is the only quantity relevant to diversification.

A portfolio constructed on labels is constructed on the map, not the territory. The map may be a reasonable approximation in some regimes, and the empirical record contains episodes — including the most institutionally memorable one — in which it has performed acceptably. Those episodes are not vindication of the framework. They are coincidences of regime. When the territory shifts and the regime changes, the map fails. The failure is not accidental. It is built into the framework at the level of its first assumption.

The geometric approach does not guarantee superior returns. No framework does. It guarantees that the structure being optimized reflects the actual statistical properties of the opportunity set — that diversification, when claimed, exists; that regime shifts, when they occur, are observable and identifiable; that the unit of construction is risk, not label.

That is the correct basis for portfolio construction. The rest is execution.

References

Ang, A., & Bekaert, G. (2002). International asset allocation with regime shifts. *The Review of Financial Studies*, 15(4), 1137–1187.

Campbell, J. Y., Sunderam, A., & Viceira, L. M. (2017). Inflation bets or deflation hedges? The changing risks of nominal bonds. *Critical Finance Review*, 6(2), 263–301.

Choueifaty, Y., & Coignard, Y. (2008). Toward maximum diversification. *Journal of Portfolio Management*, 35(1), 40–51.

Engle, R. F. (2002). Dynamic conditional correlation: A simple class of multivariate generalized autoregressive conditional heteroskedasticity models. *Journal of Business & Economic Statistics*, 20(3), 339–350.

Laloux, L., Cizeau, P., Bouchaud, J.-P., & Potters, M. (1999). Noise dressing of financial correlation matrices. *Physical Review Letters*, 83(7), 1467–1470.

Ledoit, O., & Wolf, M. (2004). Honey, I shrunk the sample covariance matrix. *Journal of Portfolio Management*, 30(4), 110–119.

Longin, F., & Solnik, B. (2001). Extreme correlation of international equity markets. *The Journal of Finance*, 56(2), 649–676.

López de Prado, M. (2016). Building diversified portfolios that outperform out of sample. *Journal of Portfolio Management*, 42(4), 59–69.

Mantegna, R. N. (1999). Hierarchical structure in financial markets. *The European Physical Journal B*, 11(1), 193–197.

Meucci, A. (2009). Managing diversification. *Risk*, 22(5), 74–79.

Meucci, A. (2010). The Black-Litterman approach: Original model and extensions. *The Encyclopedia of Quantitative Finance*. Wiley.

Pafka, S., & Kondor, I. (2004). Estimated correlation matrices and portfolio optimization. *Physica A: Statistical Mechanics and its Applications*, 343, 623–634.

Page, S., & Panariello, R. A. (2018). When diversification fails. *Financial Analysts Journal*, 74(3), 19–32.

Roncalli, T. (2013). *Introduction to Risk Parity and Budgeting*. Chapman & Hall/CRC Financial Mathematics Series.
