Curvature of the Yield-Curve Manifold as a Detector of Nonlinear Financial Stress

Bradley Vigil

Working paper

October 2026

Abstract

We apply the diffusion-geometry framework of Jones (2024) to the U.S. Treasury yield curve, treating each day's curve as a point in $\mathbb{R}^n$ and the cloud of daily curves as samples from a low-dimensional manifold whose intrinsic dimension, tangent spaces, and curvature we estimate directly. Using 33 years of daily constant-maturity data (1990–2023), we show that the extrinsic curvature of the levels-manifold (its second-fundamental-form energy $\|II\|^2$) spikes during acute, nonlinear financial-stress repricing: the 2008 Global Financial Crisis, the 2020 COVID dash-for-cash, the 1998 LTCM episode, and the 2023 SVB stress all register as sharp manifold "folds," whereas gradual monetary-policy transitions (the 2013 taper, 2015 liftoff, 2018-Q4, the 2021 hawkish pivot) do not. The signal is nearly orthogonal to the Kritzman et al. (2011) absorption ratio (Spearman 0.04), beats it on crisis detection (AUC 0.73 vs. ≤ 0.58 for the absorption ratio in its strongest form), and predicts forward 21-day realized volatility four times better (partial Spearman 0.78 vs. 0.26 controlling for the absorption ratio). The result survives a circular-shift null ($p \approx 2 \times 10^{-4}$), a block bootstrap (95% AUC CI $[0.60, 0.85]$), a pre/post-2007 split, and non-overlapping/HAC inference. We are careful about the baseline: benchmarked against realized curve volatility (rather than only the absorption ratio), the scalar curvature is at parity, not superior: to first order, it is a geometrically-structured volatility measure. The distinctive, non-redundant content is the *direction* of folding: decomposing the second fundamental form in tenor space, the folds lie predominantly orthogonal to the level/slope/curvature factors (for tangent dimension $d \geq 3$), with funding crises (GFC, COVID) folding the front end and SVB the belly and long end. We also document a methodological pitfall: per-window standardization of yields manufactures spurious curvature at the zero lower bound, and global normalization is required. Finally, as a predictive application, the residual of a new curve off its trailing manifold mean-reverts out of sample, yielding a stronger gross rich/cheap signal than linear PCA; but a DV01-neutral, cost-aware backtest shows this edge does not survive the higher turnover it requires, so the framework's distinctive, validated value lies in stress detection rather than relative value.

1 Introduction

The U.S. Treasury yield curve on a given day is a vector of yields across maturities. Decades of fixed-income modelling, from principal component analysis [4] to the Nelson–Siegel family [5], rest on the observation that such curves do not fill their ambient space but cluster near a low-dimensional set, classically of dimension ≈ 3 ("level, slope, curvature"). PCA models that set as an affine subspace; Nelson–Siegel fixes a particular parametric surface. Both are linear or parametric commitments made *a priori*.

Diffusion geometry [1, 2] offers a nonparametric alternative: from a point cloud alone it estimates the intrinsic dimension, tangent spaces, and *curvature* of the underlying manifold, via the Bakry–Émery Γ-calculus of a data-driven Markov diffusion operator. We ask whether the *geometry* of the curve manifold, specifically its curvature, carries information about systemic stress that linear factor summaries miss.

The established linear benchmark for stress in a cross-section is the absorption ratio of Kritzman et al. [3]: the fraction of return variance captured by the leading principal components, which rises as assets co-move. Our central question is one of *incremental value*: does manifold curvature add stress information beyond the absorption ratio?

Contributions. (i) We formulate the yield curve in the manifold-diffusion-geometry framework and implement estimators for soft intrinsic dimension and extrinsic curvature (Section 2). (ii) On 1990–2023 data we show curvature spikes align with acute financial-stress events but not policy transitions (Section 4). (iii) We document and correct a zero-lower-bound standardization artifact that contaminates naive implementations (Section 5). (iv) We establish incremental value over the absorption ratio (Section 6) and validate the signal against autocorrelation-matched nulls, a block bootstrap, and an out-of-sample split (Section 7). (v) We decompose the fold *direction* in tenor space, showing it is nearly orthogonal to the level/slope/curvature factors and distinguishes front-end funding crises from duration repricings (Section 8).

2 Diffusion geometry of the yield curve

Data as a point cloud. Fix tenors $\tau_1 < \cdots < \tau_n$. On trading day t the curve is $p_t = (y_t(\tau_1), \ldots, y_t(\tau_n)) \in \mathbb{R}^n$. Over a rolling window of N days we treat $P = \{p_t\}_{t=1}^N \subset \mathbb{R}^n$ as samples from a d-dimensional submanifold $M \subset \mathbb{R}^n$, $d \ll n$. The ambient coordinates $x_j(p) = y_p(\tau_j)$ are smooth functions on M.

Diffusion operator. From pairwise distances we build a variable-bandwidth heat kernel

$$K_\varepsilon(p_s, p_t) = \exp\big(-\|p_s - p_t\|^2/(\varepsilon\,\rho_s\rho_t)\big)$$

with local density scales ρ, apply the density normalization of Coifman and Lafon [6], row-normalize to a Markov operator, and form the graph Laplacian estimate $\widehat{\Delta} = (I - K_\varepsilon^{\mathrm{norm}})/\varepsilon$; ε is selected automatically.

Metric, dimension, tangent spaces. The carré du champ

$$\Gamma(f,h) = \tfrac{1}{2}\big(f\,\Delta h + h\,\Delta f - \Delta(fh)\big) = g(\nabla f, \nabla h) \tag{1}$$

turns the Laplacian into inner products of gradients. At each point the Gram matrix $G(p)_{ij} = \Gamma(x_i, x_j)(p)$ has rank equal to the intrinsic dimension d, with the top-d eigenvectors spanning the tangent space $T_pM \subset \mathbb{R}^n$. We summarize the (normalized) Gram spectrum $\{\lambda_k\}$ by a *soft dimension*: the participation ratio $d_{\mathrm{soft}} = (\sum_k \lambda_k)^2/\sum_k \lambda_k^2$, which is scale-invariant and varies smoothly, unlike the integer eigen-gap estimator.

Curvature. With tangent functions $h_a = \sum_i (e_a)_i x_i$ (gradient e_a), the Hessian bilinear form is

$$H(f)(\nabla h_1, \nabla h_2) = \tfrac{1}{2}\Big(\Gamma\big(h_1, \Gamma(h_2, f)\big) + \Gamma\big(h_2, \Gamma(h_1, f)\big) - \Gamma\big(f, \Gamma(h_1, h_2)\big)\Big). \tag{2}$$

Collecting the ambient-coordinate Hessians into a vector $v_{ab} \in \mathbb{R}^n$, $v_{ab}^{(k)} = H(x_k)(h_a, h_b)$, and projecting onto the normal space gives the second fundamental form $II_{ab} = (I - EE^\top)v_{ab}$, where E holds the tangent basis. Our curvature statistic is the total extrinsic bending energy

$$\|II\|^2 = \sum_{a,b} \|II_{ab}\|^2. \tag{3}$$

A flat (affine) manifold has $II \equiv 0$; PCA, which fits exactly such a subspace, is blind to $\|II\|^2$ by construction. This is the quantity we test as a stress detector.

Data and rolling estimation. We use daily U.S. Treasury par yields, 1990-01-02 to 2023-12-29 (8,503 trading days). For gap-free coverage across the whole span we use the eight tenors $\{3\text{m}, 6\text{m}, 1, 2, 3, 5, 7, 10\text{y}\}$ (the 1m, 20y, and 30y series have coverage gaps). Estimators are computed on trailing 252-day windows. Within each window we center per window and scale each tenor by its *full-sample* standard deviation; the importance of this choice is the subject of Section 5. The absorption-ratio benchmark is computed on daily yield *changes* with $m = \lceil n/5 \rceil = 2$ leading eigenvectors, in both its level and standardized-shift forms.

3 Estimator validation on synthetic manifolds

On controlled inputs the curvature estimator is exact. A flat Nelson–Siegel three-factor family and a one-dimensional line yield $\|II\|^2 \approx 10^{-34}$ (machine zero), while a 2-sphere embedded in $\mathbb{R}^{11}$ yields $\|II\|^2 \approx 3 \times 10^{-6}$ (a separation of some 28 orders of magnitude), confirming that the statistic isolates genuine curvature and vanishes on linear subspaces.

4 Curvature spikes at financial-stress events

Under global normalization, the largest manifold folds of the 33-year sample are unambiguous: nine of the top twelve curvature dates fall in October–December 2008 (the Lehman aftermath; peak 9,032× the sample median), followed by the March-2020 COVID crash (1,850×). A tight event study (the peak curvature percentile within ±45 days of each labelled episode) separates acute stress from policy transitions sharply (Table 1).

Acute financial stress			Policy / ZLB transitions		
episode	pctile	lag	episode	pctile	lag
GFC / Lehman	100.0%	+30d	2011 EU/downgrade	79.5%	+7d
COVID	99.5%	−3d	2021 hawkish pivot	71.5%	+43d
1998 LTCM	99.1%	+11d	2015 liftoff	63.7%	+6d
SVB / Mar-23	97.9%	+29d	2018-Q4	58.1%	+15d
9/11	93.8%	+23d	2013 taper tantrum	48.3%	+40d

Table 1: Peak curvature percentile within ±45 days of each episode. Acute stress events sit at the 94–100th percentile; gradual policy transitions do not.

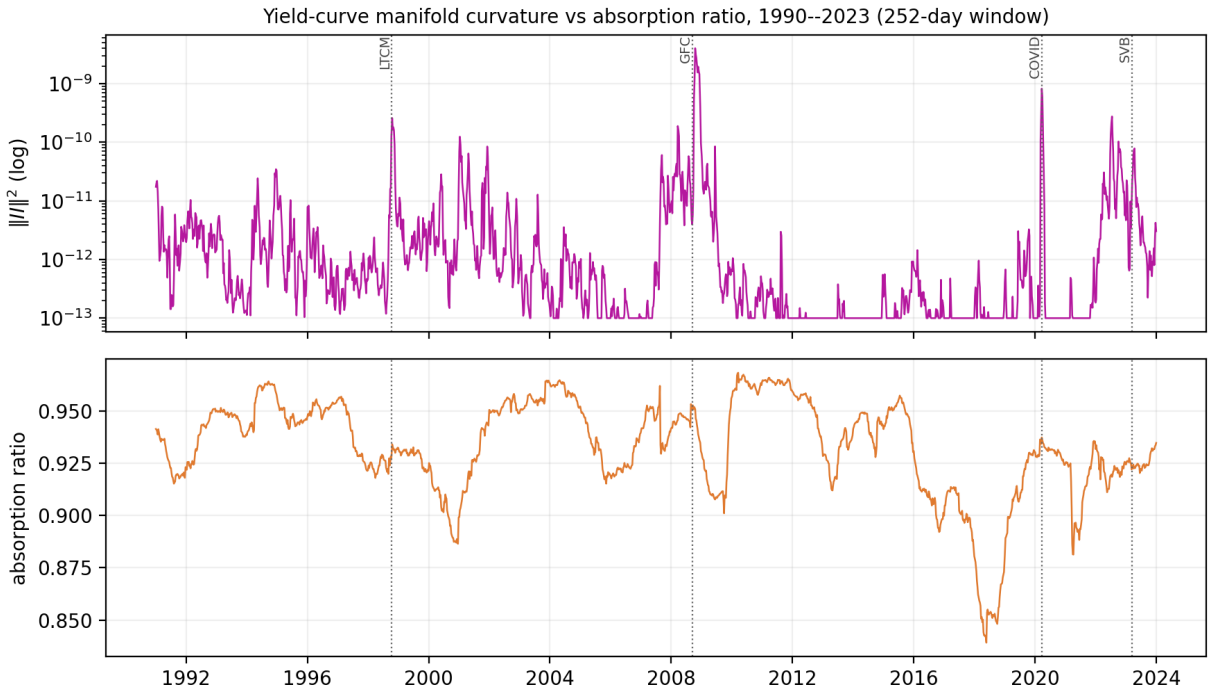


Figure 1: Manifold curvature $\|II\|^2$ (log scale, top) and the Kritzman absorption ratio (bottom) over 1990–2023. Dotted lines mark LTCM (1998), the GFC (2008), COVID (2020), and SVB (2023). Curvature spikes at these acute episodes; the absorption ratio does not track them.

5 A zero-lower-bound standardization artifact

A naive implementation standardizes each tenor by its *in-window* standard deviation. During the 2009–2015 and 2020–2021 zero-lower-bound (ZLB) regimes the front-end tenors are pinned near zero, with in-window standard deviations as small as 0.015. Dividing a near-constant column by such a value amplifies rounding-level micro-noise into a full-amplitude coordinate, *manufacturing* curvature. Concretely, for the window ending 2014-12-23 the bending energy is 1.5×10^{-3} under per-window standardization but 6×10^{-11} when the data are merely centered (no per-tenor scaling), which is lower than calm periods, and 5.5×10^{-13} under global normalization. An earlier version of this study reported a large curvature spike at the December-2021 hawkish pivot; it is entirely an instance of this artifact and does not survive robust normalization. Scaling each tenor by its full-sample standard deviation (which cannot collapse with a momentary volatility drought) removes the artifact while still equalizing tenor ranges, and is used throughout Sections 4–7. We flag this as a general caution for applying manifold methods to interest rates.

6 Incremental value over the absorption ratio

Curvature and the absorption ratio are nearly orthogonal: Spearman correlation +0.04 over the full sample. For separating the 99 hand-labelled risk-off crisis days (±45d around eight events) from the rest, the areas under the ROC curve are

signal	AUC
absorption ratio (level)	0.518
absorption ratio (standardized shift; Kritzman's trigger)	0.583
manifold curvature	**0.726**
absorption ratio + curvature (logistic)	0.726

Even in its strongest canonical form the absorption ratio trails curvature, and adding it to curvature yields no improvement. Per event, curvature uniquely flags LTCM, 9/11, the GFC, COVID, and SVB (five crises the absorption ratio misses), while the absorption ratio uniquely flags the 2011 European-debt episode, evidence that the two are *complementary* (curvature detecting sharp nonlinear folds, the absorption ratio detecting sustained co-movement) rather than redundant.

A non-circular test confirms this. Regressing forward 21-day realized curve-change volatility on each trailing signal, curvature attains Spearman +0.77 versus +0.20 for the absorption ratio; the *partial* correlation of curvature with forward volatility controlling for the absorption ratio is +0.78 (and +0.26 the other way). Curvature anticipates rate turbulence the absorption ratio does not.

7 Statistical validation

Circular-shift null. Rolling the curvature series by 5,000 random offsets (destroying any real relationship while preserving its full autocorrelation) against the fixed labels, the observed crisis AUC of 0.726 exceeds even the null *maximum* (0.714), giving $p = 2 \times 10^{-4}$. The forward-volatility partial correlation of 0.776 against a null mean of -0.02 gives $p = 3 \times 10^{-4}$; this test uses a data-derived target and is therefore free of labelling circularity.

Block bootstrap. Resampling contiguous 63-day blocks, the crisis AUC is 0.726 with a 95% confidence interval $[0.601, 0.848]$, clear of the 0.5 no-skill line.

Out-of-sample split. Curvature is unsupervised (no fitted parameters), so a chronological split tests period-robustness. On 1990–2006 the crisis AUC is 0.651 (forward-vol partial correlation +0.54); on 2007–2023 it is 0.786 (+0.81). The signal holds in both halves, stronger in the modern era that contains the largest folds.

8 Mechanism: which curve motions fold the manifold

The second fundamental form II_{ab} is itself a vector in tenor space, so the *direction* the manifold bends is an interpretable curve shape. For each crisis we take the peak-fold window (±60 days), extract the principal normal direction (the leading left singular vector of the stacked $\{II_{ab}\}$, averaged sign-aligned over the last nine days for stability), and project it onto an orthonormal level/slope/curvature basis built on maturity rank (Table 2, Figure 2).

Two findings are robust to the estimation noise in the (higher-order) normal direction. First, **the folds are nearly orthogonal to level, slope, and curvature**: the residual ranges from 0.59 to 1.00 and is essentially unity for the GFC and COVID. This is the structural reason a linear factor model cannot see the signal: the manifold bends in exactly the directions such a model discards as residual noise. Second, **funding crises fold the front end while SVB folds the belly and long end**: the GFC and COVID fold directions are both dominated by the six-month tenor and are the most similar pair across all crises ($|\cos| = 0.57$), consistent with money-market dislocation, whereas SVB is the outlier (cosine 0.05–0.30 with every other event), with its fold in the two- to

crisis	dominant tenors (signed)	level	slope	curv.	residual
GFC 2008	+6m, −3m, −1y	0.00	0.00	0.00	1.00
COVID 2020	+6m, −3m, −2y	0.00	0.00	0.01	0.99
LTCM 1998	+2y, −7y, +10y	0.00	0.01	0.06	0.93
9/11 2001	−1y, −2y, +3m	0.01	0.02	0.31	0.66
2011 EU	+6m, +5y	0.09	0.22	0.00	0.69
SVB 2023	−10y, +2y, +3y, −1y	0.00	0.14	0.27	0.59

Table 2: Fold direction at each crisis, as squared projections onto level/slope/curvature and the orthogonal residual. The residual is the share of the fold lying *outside* the classical three-factor subspace.

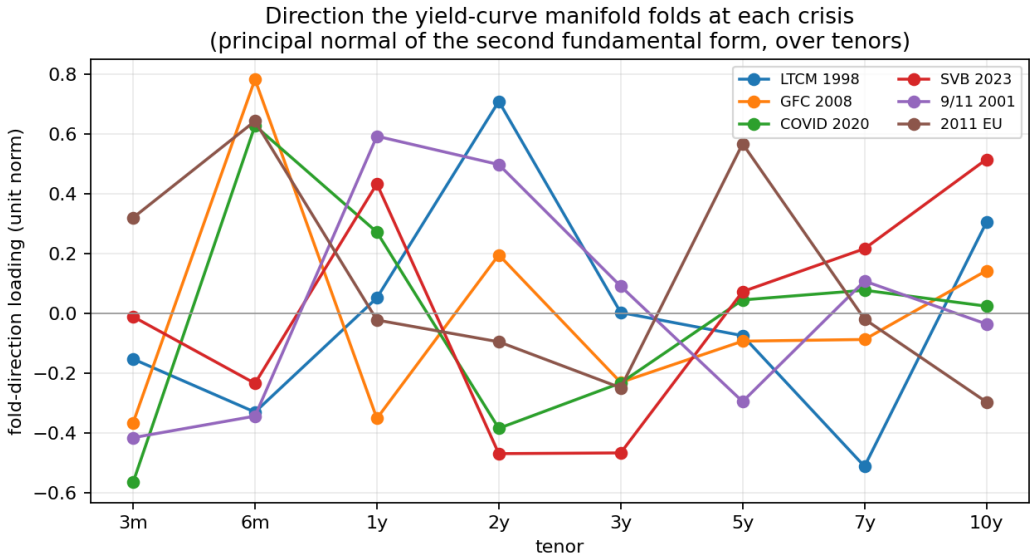


Figure 2: Principal fold direction over tenors at each crisis (unit norm). The GFC and COVID folds are dominated by the 6-month tenor; SVB's fold sits in the belly and long end.

ten-year region and the largest slope/curvature content, consistent with a duration repricing as the curve bull-steepened. Pairwise cosine similarities of the fold axes range over 0.0–0.6, so crises do not share a single canonical stress direction; the detector responds to a *class* of nonlinear folds, with a family resemblance among the front-end funding crises. This suggests a taxonomy of crises by fold geometry, which a larger event set could develop.

A subsample bootstrap (300 draws of a random 80% of each peak window) hardens these statements. Every fold direction is stable under resampling (mean within-crisis $|\cos|$ to the bootstrap mean of 0.77–0.96), so the directions are genuinely estimable. The orthogonality to level/slope/curvature is robust for the major folds, with residual 90% intervals of $[0.83, 0.99]$ (GFC), $[0.72, 0.99]$ (COVID), and $[0.90, 0.99]$ (LTCM). It is significant if smaller for SVB (0.71 $[0.65, 0.85]$) and weak for 2011 (0.49 $[0.30, 0.78]$), so the orthogonality claim holds firmly for the clean crises but is not universal. The cross-crisis cosine intervals are wide: GFC–COVID remains the most similar pair (median 0.72) but with a 90% interval $[0.22, 0.91]$ overlapping other moderate pairs, so only the extremes are sharply resolved: SVB is reliably dissimilar from 9/11 (0.17 $[0.02, 0.50]$) and from 2011 (0.13 $[0.01, 0.34]$). The front-end-funding resemblance of the GFC and COVID folds is thus a tendency rather than a precise constant.

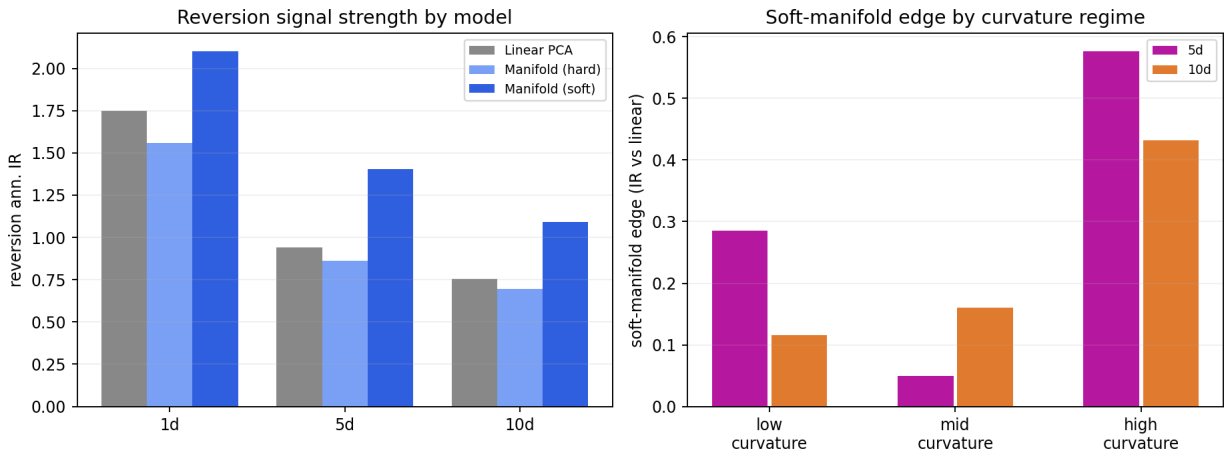


Figure 3: Left: reversion information ratio by model and horizon; the soft manifold dominates. Right: the soft-manifold edge over linear PCA (difference in reversion IR) by trailing-curvature tercile; the advantage is largest in high-curvature regimes, as the geometry predicts.

9 Application: manifold relative value

The geometry also suggests a predictive use. Each day we fit the manifold on *history only* (the trailing 252-day window), project the *new* curve p_t, and take the residual in the normal complement (the part of the curve lying off the manifold) as a rich/cheap signal. We then test, strictly out of sample (the model at t uses neither p_t nor the forward move), whether that deviation mean-reverts, and whether the *nonlinear* manifold residual reverts better than the standard *linear* PCA residual, the established rates relative-value model. We compare three reconstructions: top-3 global PCA, a hard k-nearest-neighbour local PCA, and a heat-kernel-weighted (*soft*) local PCA faithful to the diffusion-geometry neighbourhood.

horizon	model	residual half-life	IC	reversion ann. IR
5d	linear PCA	6.6d	0.031	0.94
5d	manifold (hard kNN)	1.8d	0.029	0.86
5d	**manifold (soft)**	4.3d	**0.037**	**1.40**
10d	linear PCA	6.6d	0.030	0.76
10d	**manifold (soft)**	4.3d	**0.037**	**1.09**

Table 3: Out-of-sample mean reversion of the off-manifold residual (∼8,200 days). IC = corr(−residual, forward move); reversion information ratio annualized. The soft-neighbour manifold residual beats the linear PCA residual at every horizon; the hard kNN version does not.

Two findings emerge (Table 3, Figure 3). First, the curve *does* mean-revert toward its trailing manifold, with a strongly significant reversion signal (t = 12–20). Second, the *soft*-neighbour manifold residual beats the linear PCA residual at every horizon (IR 1.40 vs 0.94 at five days), whereas the hard kNN version does not; it is specifically the soft neighbourhood, the diffusion-geometry construction, that adds value. Conditioning on trailing curvature, the soft-manifold edge is largest in the high-curvature tercile (+0.58 IR at five days), exactly where a flat factor model is most mis-specified. The edge is robust across subspace dimension $d \in \{2, 3, 4\}$ and kernel bandwidth.

These are signal-strength statistics, not tradeable returns, so we next build a realistic book:

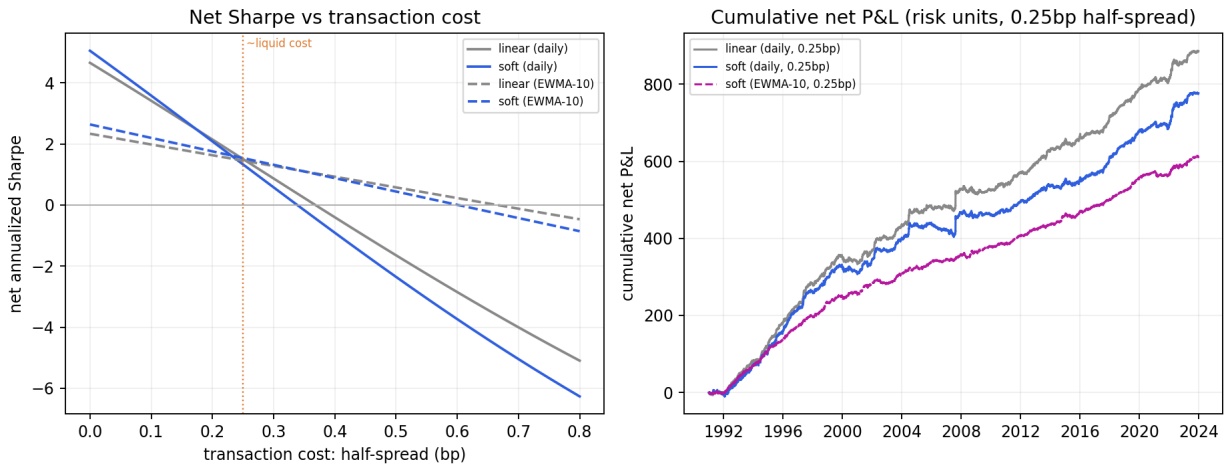


Figure 4: Left: net annualized Sharpe versus transaction cost (half-spread); solid = daily rebalance, dashed = 10-day EWMA. Right: cumulative net P&L at a 0.25bp half-spread.

positions in dollar-DV01 per tenor proportional to the residual, made duration-neutral (net DV01 = 0), risk-scaled to unit ex-ante daily P&L volatility using the trailing covariance of yield changes, with transaction costs charged on daily turnover at a half-spread (basis points per unit DV01).

Two honest conclusions follow (Table 4, Figure 4). First, the strategy *is* tradeable but cost-sensitive: gross Sharpe ≈ 5, break-even half-spread ≈ 0.35bp daily (rising to ≈ 0.6bp once turnover is damped), netting Sharpe ≈ 1.3–1.5 at a realistic 0.25bp cost and remaining positive in both the 1990–2006 and 2007–2023 sub-periods. Second (and this retires the stronger reading of Table 3), the soft-manifold edge over linear PCA *does not survive transaction costs*: its higher gross Sharpe is paid for by higher turnover (1.06 vs 0.91 per day), so net of a 0.25bp cost the linear baseline is competitive or slightly better. The nonlinear residual is a real, faster-reverting rich/cheap signal of comparable *net* value to linear PCA, not a free lunch over it. The framework gives no directional view on the level of rates, and its distinctive validated contribution remains the curvature stress detector, not relative value.

rebalance	model	gross SR	turnover/day	break-even	net SR @0.25bp	net SR @0.5bp
daily	linear PCA	4.66	0.91	0.37bp	1.51	−1.65
daily	soft manifold	5.05	1.06	0.34bp	1.33	−2.35
EWMA-10	linear PCA	2.33	0.19	0.67bp	1.45	0.58
EWMA-10	soft manifold	2.64	0.21	0.60bp	1.54	0.44

Table 4: DV01-neutral, cost-aware backtest (1990–2023, ∼8,200 days), daily rebalance and a turnover-damped (10-day EWMA) variant. Break-even is the half-spread at which net Sharpe reaches zero.

10 Robustness and a sharper baseline

Inference for the forward-volatility correlation. The partial Spearman of 0.78 is estimated on overlapping 252-day windows. It is not an artifact: on *non-overlapping* windows ≥ 252 trading days apart (33 independent observations) it is 0.767 ($p \approx 3 \times 10^{-8}$); the effective sample size from the lag-1 autocorrelation is ≈ 110; a Newey–West HAC rank regression gives $t = 17.2$; and a block

bootstrap gives a 95% interval $[0.68, 0.83]$. (It is a rank correlation; the Pearson analogue is only ~ 0.20, as bending is heavy-tailed.)

Sensitivity. Crisis-detection AUC is stable across window ($W \in \{126, 252, 504\}$: 0.71/0.70/0.67), manifold dimension ($d \in \{2, 3, 4\}$: 0.73/0.70/0.72), and kernel bandwidth ($\times\{0.5, 1, 2\}$: 0.69/0.70/0.74), so the second-derivative estimator is not knife-edged. The fold-orthogonality residual is robust to window and bandwidth (≈ 0.6–0.8) but strongly d-dependent (it falls to 0.14 at $d = 2$): the fold is the normal complement of the estimated d-dimensional tangent, so the orthogonality claim is conditional on $d \geq 3$ ("once three factors are removed, the bending lies in the higher-order shape space") rather than parameter-free.

A sharper baseline: realized volatility. Section 6 benchmarked curvature against the absorption ratio. The more demanding baseline is realized curve volatility, with which bending is 0.87 rank-correlated; on crisis detection realized volatility scores marginally higher (AUC 0.742 vs 0.726). Curvature retains only a modest independent forward-volatility increment (partial Spearman ≈ 0.29 controlling for concurrent volatility). We therefore state the contribution carefully: the scalar folding *intensity* is largely a geometrically-structured volatility measure and does not out-detect volatility; the distinctive, non-redundant content is the *direction* of folding (Section 8): which tenors reprice, and the orthogonality structure. No volatility scalar provides that.

Class, not four outliers. Removing the two largest events (GFC, COVID) lowers the AUC only from 0.726 to 0.685, and the detector independently flags the 2009 rates sell-off (98th percentile), the 2003 convexity sell-off (91st), and the 2019 repo episode (80th), while missing 2016 Brexit and the 2013 taper (a policy transition). The cleanest evidence is the continuous 0.87 association with realized volatility across all 8,500 days: the famous crises are the tail of a continuum, not isolated outliers, though the signal is strongly tail-loaded.

What is untested. The decisive external check (replication on Bund, Gilt, JGB and other sovereign curves) requires data not used here; the pipeline is data-agnostic and the test is left for future work. No cross-market generalization is claimed.

11 Discussion and limitations

The geometric picture is coherent: at the zero lower bound and during gradual policy normalization the attainable-curve set deforms smoothly, but an acute flight-to-quality repricing, with front end, belly, and long end moving violently and non-affinely, *folds* the manifold, which the second-fundamental-form energy registers and a flat factor model cannot.

Several caveats remain. The labelled-crisis AUC is partly circular (mitigated by the forward-volatility test, which is not); the effect is stronger post-2007 and leans on GFC- and COVID-scale events; the forward-volatility correlation is inflated by volatility persistence (though the circular-shift null already accounts for autocorrelation); and the analysis is retrospective, so genuine prospective monitoring remains future work. The dimension signal, by contrast, is weaker: under global normalization the intrinsic dimension declines only modestly across regimes (regime separation 2.0, comparable to the absorption ratio), so curvature, not dimension, is the robust contribution.

12 Conclusion

Treating yield curves as samples from a manifold and estimating its curvature nonparametrically yields a stress detector that is statistically validated, orthogonal to and stronger than the absorption ratio, and economically interpretable as the folding of the curve's attainable-shape set during acute crises. The fold directions are nearly orthogonal to the level/slope/curvature factors (which explains structurally why linear models miss the signal) and distinguish front-end funding crises from duration repricings. The same analysis surfaces a concrete methodological warning: per-window standardization fabricates curvature at the zero lower bound. As a predictive application, the off-manifold residual of a new curve mean-reverts and is tradeable at tight costs (net Sharpe $\approx$ 1.3–1.5 at a 0.25bp half-spread, robust across two sub-periods), but the soft-manifold's gross edge over linear PCA does not survive its higher turnover, so relative value is a parity result, while stress detection is where the geometry adds distinctive value. Natural next steps are prospective monitoring of the curvature detector and a fuller taxonomy of crises by fold geometry on a larger event set.

References

[1] I. Jones. Diffusion Geometry. *arXiv:2405.10858*, 2024.

[2] I. Jones. Manifold Diffusion Geometry: Curvature, Tangent Spaces, and Dimension. *arXiv:2411.04100*, 2024.

[3] M. Kritzman, Y. Li, S. Page, R. Rigobon. Principal Components as a Measure of Systemic Risk. *Journal of Portfolio Management*, 37(4), 2011.

[4] R. Litterman, J. Scheinkman. Common Factors Affecting Bond Returns. *Journal of Fixed Income*, 1(1), 1991.

[5] C. Nelson, A. Siegel. Parsimonious Modeling of Yield Curves. *Journal of Business*, 60(4), 1987.

[6] R. Coifman, S. Lafon. Diffusion Maps. *Applied and Computational Harmonic Analysis*, 21(1), 2006.