

Unified Boundary Field Theory of Interest Rates: An Empirical Study of Robin Short-Rate Boundary Coordinates and Low-Rank Covariance in the Yield Curve

Working Paper

July 4, 2026

Abstract

This paper provides an empirical test of the Robin short-rate boundary-coordinate layer in Unified Boundary Field Theory (UBFT). The central claim is not merely that Treasury yield curves are low-dimensional. Rather, the paper asks whether a single short-rate boundary geometry can organise several empirical objects that are usually modelled separately: yield-curve shape, maturity-dependent volatility, local boundary anchoring, covariance geometry, and phase-state behaviour.

Using daily fitted U.S. Treasury zero-coupon yield curves from the Federal Reserve/Svensson data set over 1985–2026, I estimate a three-coordinate Robin short-rate boundary layer consisting of an amplitude coordinate, a local short-end slope, and a long-boundary anchor. The same maturity-loading geometry that represents the curve also governs variance transport across maturities through the interaction of the maturity loading vector with the boundary-coordinate covariance matrix. At the innovation level, the empirical correlation matrix of these boundary coordinates lies close to a signed rank-two manifold. Across broad regimes, blind five-year windows, rolling windows, holistic phase windows, human-selected windows, corrfine windows, random partitions, coordinate placebo tests, and targeted identification checks, the implied correlations reproduce observed correlations with high explanatory power and persistent branch selection.

The evidence supports a segmentation-robust boundary-layer interpretation: the Robin short-rate boundary coordinates are not only a curve-fitting device, but a common empirical geometry linking shape, volatility transport, boundary anchoring, and innovation covariance. Placebo tests using Nelson–Siegel coordinates, Svensson subsets, endpoint–slope decompositions, and random rotations show that alternative representations can reproduce individual pieces of the evidence, but do not provide the same unified boundary interpretation. Additional identification checks show that the result is not explained by replacing the Robin amplitude with the raw fitted short-boundary level, and that the stable residual wedge is not materially generated by nonlinear transformation bias in the correlation estimator. The contribution is therefore a structural organisation result rather than a claim of exact rank reduction.

1 Introduction

Yield-curve dynamics are usually described by statistical factors, macro-finance state variables, or no-arbitrage volatility restrictions. This paper asks a different empirical question: whether several objects that are normally modelled separately become jointly organised when the curve is represented through boundary coordinates. It is the fourth paper in the Unified Boundary Field Theory (UBFT) sequence. Paper 1, *An Axiomatically Consistent Single-Factor Term Structure Model with Admissible Humped Volatility*, establishes the admissible single-factor boundary-field setting. Paper 2, *Unified Boundary Field Theory of Interest Rates II: Monetary Policy as a Short-End Boundary Condition*, develops the short-end boundary classification. Paper 3, *Unified Boundary Field Theory of Interest Rates III: The Stochastic Long Boundary: Volatility-Aligned Geometry, Affine Admissibility, and Two-Channel MSNR for Instantaneous Rolling Forward Rates*, completes the two-boundary geometry used here. The present paper turns those theoretical restrictions into empirical diagnostics.

The word “Robin” is used here in the boundary-value sense, but the empirical object is specifically a short-rate boundary geometry. UBFT applies standard boundary-condition terminology to the short end of the yield curve. A Neumann short-rate boundary is the limiting case in which the short boundary is locally free, so the short-end slope is the primitive boundary object. A Dirichlet short-rate boundary is the opposite limiting case, in which the short boundary is pinned to a fixed level, denoted x_D . The Robin short-rate boundary is the mixed case: the short-end slope responds to the distance between the fitted short-boundary level and the pinned level x_D . In the empirical specification below this response coefficient is denoted by χ ; it measures local anchoring strength. The Robin case is the main object of the paper because it allows the short boundary to be neither completely free nor completely fixed, which is the economically realistic intermediate case between a fully floating short end and a fully pinned short end.

The Robin short-rate boundary representation describes the curve through a long-boundary state X_t , a fitted short-boundary amplitude A_t , and a fitted short-boundary slope s_t . These variables are not treated as new tradable instruments. They are local coordinates for the shape and boundary behaviour of the fitted yield curve. The central empirical object is therefore the Robin short-rate boundary layer

$$Z_t^R = (A_t, s_t, X_t)^\top. \quad (1.1)$$

The paper asks whether curve shape, maturity-volatility structure, local anchoring, covariance geometry, and phase behaviour appear as different projections of this same layer.

1.1 The unification thesis

Standard term-structure models usually begin with a statistical decomposition of the curve,

$$\text{yield curve} \longrightarrow \text{factors},$$

and then ask how those factors perform in separate applications such as curve fitting, volatility modelling, risk aggregation, or regime classification. The Robin short-rate boundary construction reverses this logic. It begins with a boundary-value geometry and asks whether the same coordinate layer organises several empirical dimensions of the yield curve at once:

boundary geometry $\longrightarrow$ curve shape
 $\longrightarrow$ volatility transport
 $\longrightarrow$ innovation covariance
 $\longrightarrow$ local boundary states.

This is the sense in which the framework is unified. The object of interest is not any single diagnostic, such as a low-rank covariance matrix or a volatility fit. The object is the boundary-coordinate layer Z_t^R , and the empirical question is whether the same coordinates organise multiple features of the Treasury curve that are usually estimated and interpreted separately.

This distinction matters for comparison with standard low-dimensional factor models. Principal components and Nelson–Siegel factors can explain substantial variation in Treasury yields, and simple endpoint–slope decompositions can sometimes generate mechanical low-rank covariance. The Robin claim is not that no alternative model can match any one of these features. Rather, the claim is that a single theoretically motivated boundary-coordinate system gives a coherent account of the whole bundle: cross-sectional shape, volatility transport, anchoring behaviour, signed covariance geometry, and boundary-state evolution.

Diagnostic	What it establishes or rules out
Curve-shape fit	Robin coordinates are not introduced only for covariance; they first organise the cross-section of the fitted curve.
Volatility transport	The same loading geometry maps boundary-coordinate covariance into the maturity term structure of realised volatility.
Boundary anchoring	The coordinates describe measurable local boundary states rather than only a statistical factor rotation.
Rank-two triangle	The covariance footprint is quantitative, signed, and testable rather than visually suggestive.
Bidirectional triangle	The result is not created by choosing one convenient dependent correlation.
Window and random-partition tests	The geometry is segmentation-robust and not induced by hand-selected phase boundaries.
Tangent-direction test	Rotating phase means and planar innovation covariance are distinct empirical objects.
Coordinate placebo tests	The result is not simply generic PCA, Nelson–Siegel, or endpoint-slope low-dimensionality.
Identification checks	The result is not explained by the raw short-boundary coordinate, bootstrap nonlinear-transformation bias, OLS reversal asymmetry, or small-window leverage.

Table 1: Empirical role of the main diagnostics. The tests are intended as a bounded robustness suite rather than an open-ended search for additional regularities.

The empirical design is deliberately sequential. The paper first estimates the maturity-decay parameter a^2 from curve shape. It then asks whether the same loading geometry organises the term structure of realised volatility. Only after constructing these boundary-coordinate objects does the paper turn to the covariance matrix of innovations in (A_t, s_t, X_t) . This ordering is important because the low-rank covariance result should be earned rather than engineered. Within the UBFT sequence, this paper therefore plays the role of an empirical boundary-coordinate test: it asks whether the Robin coordinates derived theoretically in Paper 3 leave a measurable covariance-space footprint in U.S. Treasury curve data.

The main finding is that the Robin short-rate boundary layer exhibits a robust near-two-dimensional correlation geometry. Define the innovation correlations

$$\rho_{AX} = \text{Corr}(\Delta A_t, \Delta X_t), \quad \rho_{sX} = \text{Corr}(\Delta s_t, \Delta X_t), \quad \rho_{sA} = \text{Corr}(\Delta s_t, \Delta A_t). \quad (1.1.1)$$

If the 3×3 correlation matrix of $(\Delta A_t, \Delta s_t, \Delta X_t)$ is rank two, the three correlations cannot vary independently. For example,

$$\rho_{sA} = \rho_{AX}\rho_{sX} \pm \sqrt{(1 - \rho_{AX}^2)(1 - \rho_{sX}^2)}. \quad (1.1.2)$$

Empirically, the same orientation branch is selected in all broad Robin subperiods, all blind non-overlapping five-year windows, all rolling five-year windows, and all main holistic boundary-phase windows. In the extended phase-specific inverse test there is one short corrfine exception, but the minus branch remains selected in all other windows. This branch stability is as important as the fit itself: it suggests a persistent orientation in the covariance geometry rather than a mechanical low-dimensional fitting artefact.

The paper is cautious about interpretation. The evidence does not say that the covariance matrix is exactly rank two at every date. Nor does it say that directly measured drift ratios are maturity-invariant in daily data. The stronger, more defensible claim is that Robin boundary innovations cluster near a stable two-dimensional correlation manifold, with a small positive residual wedge. The new robustness tests sharpen this interpretation: careful phase identification is useful for economic description, but the rank-two covariance footprint is already present under blind and random partitions. This covariance-space footprint is the empirical counterpart to the affine return-to-risk discipline developed in Paper 3.

2 Data and fitted boundary coordinates

The data are daily fitted zero-coupon yield curves from the Federal Reserve/Svensson data set over 1985–2026. Let $x_{t,t+\tau}$ denote the fitted curve value at date t and maturity τ . The empirical long boundary X_t is proxied by the long end of the fitted curve. The maturity grid is treated as a set of observed cross-sectional curve points.

The three short-rate boundary cases therefore have a simple empirical interpretation. Dirichlet

treats the short boundary as pinned at x_D ; Neumann treats it as locally free; Robin estimates an intermediate anchoring law with finite χ . The two-boundary Neumann benchmark is

$$x_{t,t+\tau}^N = e^{-a^2\tau} x_{t,t}^N + (1 - e^{-a^2\tau}) X_t. \quad (2.1)$$

The parameter a^2 controls the rate at which maturity loadings move away from the short boundary and towards the long boundary.

The Robin representation adds a short-boundary slope channel:

$$x_{t,t+\tau}^R = X_t + [A_t + (s_t + a_R^2 A_t)\tau] e^{-a_R^2\tau}. \quad (2.2)$$

Equivalently,

$$x_{t,t+\tau}^R = \ell_R(\tau)^\top Z_t^R, \quad Z_t^R = (A_t, s_t, X_t)^\top, \quad (2.3)$$

where

$$\ell_R(\tau) = \begin{pmatrix} (1 + a_R^2\tau)e^{-a_R^2\tau} \\ \tau e^{-a_R^2\tau} \\ 1 \end{pmatrix}. \quad (2.4)$$

Here $A_t = x_{t,t}^R - X_t$ is the fitted amplitude of the short boundary relative to the long boundary, and $s_t = \partial_\tau x_{t,t+\tau}^R|_{\tau=0}$ is the fitted short-end slope.

3 Maturity decay and volatility loading

The first empirical object is the maturity-decay parameter a^2 . Estimating the Neumann loading in (2.1) over the full sample gives

$$\hat{a}^2 \simeq 0.3254, \quad \hat{a} \simeq 0.5704. \quad (3.1)$$

The paper does not require a^2 to be literally constant. The claim is more modest: maturity-decay geometry is substantially more stable than the Robin anchoring parameters introduced below.

The same loading geometry should also organise realised volatility. Under the two-boundary benchmark,

$$\Delta x_{t,t+\tau} = e^{-a^2\tau} \Delta x_{t,t} + (1 - e^{-a^2\tau}) \Delta X_t, \quad (3.2)$$

so the maturity- τ innovation variance is

$$v_N(\tau) = e^{-2a^2\tau} g_x^2 + (1 - e^{-a^2\tau})^2 g_X^2 + 2\rho e^{-a^2\tau} (1 - e^{-a^2\tau}) g_x g_X. \quad (3.3)$$

In Robin coordinates the corresponding projection is

$$v_R(\tau) = \ell_R(\tau)^\top \Sigma_R \ell_R(\tau), \quad \Sigma_R = \text{Cov}(\Delta A_t, \Delta s_t, \Delta X_t). \quad (3.4)$$

Thus the volatility section is not the paper’s final result. Its role is to establish that boundary loadings organise maturity risk before the paper studies the covariance matrix of the boundary coordinates themselves.

Model	a^2	Curve RMSE	Curve R^2	Vol NRMSE	Vol shape corr.
Dirichlet	0.147	0.231	0.874	0.225	0.547
Neumann	0.127	0.387	0.689	0.290	0.382
Robin	0.113	0.174	0.924	0.161	0.679

Table 2: Average curve and volatility fit across subperiods by boundary model.




Figure 1: Term-structure-of-volatility fit by boundary model. The Robin projection uses the fitted boundary state and its covariance matrix to explain realised maturity-volatility shape.

4 Robin boundary anchoring

The Robin boundary condition motivates a local relation between the fitted short-boundary slope and the fitted short-boundary level:

$$s_t = -\chi(x_{t,t}^R - x_D) + \varepsilon_t. \quad (4.1)$$

The parameter χ measures local anchoring strength: larger values indicate that the short-end slope responds more strongly when the fitted short-boundary level moves away from the pinned level. The parameter x_D is the effective pinned boundary level implied by the local regression, not an externally imposed policy-rate target. The regression R^2 measures the strength of level anchoring.

The broad-regime estimates show that x_D and χ are far more regime-sensitive than maturity decay. This distinction is central to the empirical narrative: a^2 controls maturity loading, while x_D and χ control local boundary anchoring.

The Post-2021 tightening period is the clear exception to stable level anchoring in Table 3. Its low anchoring R^2 is consistent with a period in which the short boundary was being actively repriced rather than pinned to a stable local boundary level. This is economically informative rather than anomalous: it separates the movement of boundary-state means during a policy transition from

the innovation-covariance geometry studied below.

Regime	a_R^2	a_R	χ	x_D	R^2
Full sample	0.128	0.357	0.094	10.523	0.244
Pre-GFC	0.122	0.349	0.145	9.409	0.628
GFC-ZLB-QE	0.145	0.381	0.387	2.902	0.521
Liftoff	0.040	0.200	0.077	3.477	0.819
COVID	0.169	0.412	0.380	0.886	0.479
Post-2021	0.088	0.297	0.069	6.687	0.076

Table 3: Broad-regime Robin boundary estimates.

Note: The full-sample x_D is an era-wide fixed point from a pooled regression, not a weighted average of the subperiod pinned levels. Because the pooled anchoring slope χ is small and the sample mixes distinct policy environments, the extrapolated crossing level can lie above every individual subperiod estimate.




Figure 2: Estimated Robin anchoring strength χ by broad regime. Boundary anchoring is substantially more regime-sensitive than maturity-decay geometry.

5 From gross regimes to holistic boundary phases

Broad policy regimes are economically interpretable, but they need not coincide with locally homogeneous Robin boundary phases. The fitted coordinates often show substructure within a gross regime. In particular, scatter plots of s_t against A_t and s_t against X_t suggest that the boundary layer moves through local trend phases in the three-dimensional space (A_t, s_t, X_t) .

To make this observation usable without making it the primary source of identification, I treat the hand-guided windows as diagnostic boundary phases. For each phase j , I estimate

$$s_t = c_j + \beta_{A,j} A_t + \varepsilon_t, \quad (5.1)$$

$$s_t = d_j + \beta_{X,j} X_t + \eta_t, \quad (5.2)$$

and the joint local trend

$$\begin{pmatrix} A_t \\ s_t \\ X_t \end{pmatrix} = z_{0,j} + v_j \tau_t + u_t. \quad (5.3)$$

This separates slow boundary-layer movement from the innovation covariance geometry studied later.

The holistic windows are not used to replace the broad subperiods, and they are not the primary source of the rank-two result. They provide an interpretable economic segmentation of the boundary layer. The stronger covariance claim below is that the rank-two footprint survives across broad windows, blind calendar windows, holistic windows, human-selected windows, corrfine windows, and random partition benchmarks. Table 4 gives six core windows where the level geometry and covariance geometry are both strong.

H	Source	Dates	n	$\text{Corr}(A, s)$	$\text{Corr}(s, X)$	$R^2(s \sim A + X)$	$\lambda_3/3$
1	1a	1985-11-25–1986-07-16	160	-0.956	0.877	0.915	0.035
18	14ba	2003-08-26–2007-05-07	925	-0.942	0.660	0.942	0.018
26	18bb	2012-01-04–2015-12-15	989	-0.931	-0.511	0.966	0.040
28	20	2020-02-25–2021-05-28	318	-0.925	0.658	0.889	0.023
29	21	2021-06-01–2021-12-31	148	-0.943	0.768	0.891	0.041
33	24ba	2018-10-10–2019-05-24	155	-0.980	0.829	0.966	0.031

Table 4: Core holistic boundary-phase windows. These windows have coherent local A – s and s – X behaviour and small residual third-dimensional covariance mass.

6 Tangent directions and covariance planarity

A separate diagnostic asks whether adjacent boundary phases move in similar directions in the level space (A, s, X) . For each window I estimate a local trend direction and then compare adjacent directions. This tangent-direction test addresses a different question from the rank-two covariance triangle. Tangent angles describe the movement of local phase means; the triangle describes the covariance geometry of innovations around those local movements.

Table 5 shows that adjacent phase directions are not generally persistent. Median adjacent tangent angles are close to orthogonal for the broad Robin and corrfine windows, and obtuse for the holistic and human-selected windows. This is not evidence against the covariance result. It means that local boundary phases can turn or reverse while innovations continue to lie near a common low-dimensional covariance surface. The within-window regression $R^2(s \sim A + X)$ remains high, with medians from 0.816 to 0.895, supporting the distinction between rotating phase means and persistent covariance planarity.




Figure 3: Holistic boundary-phase scatter diagnostics. The purpose of the finer windows is to isolate local movement in (A, s, X) , not to maximise any single correlation mechanically.




Figure 4: Holistic phase scores combining local level coherence and innovation covariance diagnostics.

Window set	Windows	Median adjacent angle	Median within-half angle	Median $R^2(s \sim A + X)$
Broad Robin	20	87.01	65.53	0.816
Corrfine	30	87.75	102.23	0.895
Holistic	34	121.85	48.09	0.890
Human-selected	22	116.57	90.36	0.875

Table 5: Tangent-direction stability and within-phase level geometry. Angles are in degrees.

7 Boundary innovations and the Robin correlation matrix

The stochastic object of interest is the covariance matrix of boundary innovations:

$$\Delta Z_t^R = (\Delta A_t, \Delta s_t, \Delta X_t)^\top, \quad \Sigma_{R,j} = \text{Cov}_j(\Delta A_t, \Delta s_t, \Delta X_t). \quad (7.1)$$

The corresponding correlation matrix is

$$R_R = \begin{pmatrix} 1 & \rho_{sA} & \rho_{AX} \\ \rho_{sA} & 1 & \rho_{sX} \\ \rho_{AX} & \rho_{sX} & 1 \end{pmatrix}. \quad (7.2)$$

A global covariance-homogeneity test across the broad Robin subperiods rejects equality of covariance matrices, with corrected statistic approximately 46005.7 on 114 degrees of freedom and permutation p -value about 0.005. Thus the subperiods correspond not only to different level fits but to distinct local covariance geometries.

The Paper 3 MSNR condition motivates the search for low-dimensional covariance geometry, but this paper does not treat daily realised drift ratios as direct evidence for maturity-invariant MSNR. The stable empirical object is the correlation matrix (7.2).

8 Rank-two covariance triangle

If R_R is locally rank two, then

$$\det R_R = 0. \quad (8.1)$$

Solving for one off-diagonal correlation gives the rank-two triangle

$$\rho_{sA} = \rho_{AX}\rho_{sX} \pm \sqrt{(1 - \rho_{AX}^2)(1 - \rho_{sX}^2)}. \quad (8.2)$$

Geometrically, the three standardised boundary innovations behave like three unit vectors in a plane. The sign selects the orientation branch.

Empirically, the minus branch is always selected:

$$\rho_{sA}^{(-)} = \rho_{AX}\rho_{sX} - \sqrt{(1 - \rho_{AX}^2)(1 - \rho_{sX}^2)}. \quad (8.3)$$

Table 6 summarises the original test. Across broad Robin subperiods, the implied ρ_{sA} explains 95.2 per cent of the observed variation. Across blind non-overlapping and rolling five-year windows, it explains more than 99 per cent. The stable positive intercept indicates that the exact rank-two identity under-predicts ρ_{sA} by a small and regime-stable residual amount.

The small-window-count summaries should be read as diagnostics rather than standalone inference. In particular, the non-overlapping five-year and core-holistic checks contain few windows, so their high R^2 values are most informative when viewed alongside the larger rolling, phase-specific, random-partition, placebo, and identification tests.

Window set	Windows	Intercept	Slope	R^2	Corr.	MAE
Robin subperiods	20	0.157	0.981	0.952	0.976	0.166
Blind non-overlap 5y	9	0.158	0.847	0.992	0.996	0.179
Blind rolling 5y quarterly	143	0.162	0.852	0.992	0.996	0.185

Table 6: Original rank-two test: observed ρ_{sA} regressed on implied $\rho_{sA}^{(-)}$ from (ρ_{AX}, ρ_{sX}) .




Figure 5: Observed versus implied ρ_{sA} under the rank-two covariance triangle.

9 Why this is not ordinary yield-curve PCA

A natural objection is that Treasury yield changes are already known to be low-dimensional. Standard yield-curve principal-component analysis usually finds that a small number of statistical factors, often labelled level, slope, and curvature, explains most cross-sectional yield variation. The claim in this paper is different. A principal-component decomposition concerns variance concentration in the original panel of yield changes. The Robin test concerns the admissible correlation geometry among the estimated boundary coordinates (A_t, s_t, X_t) .

This distinction matters because a generic two- or three-factor yield curve does not imply the signed triangle tested here. Ordinary PCA can say that most curve movements lie in a low-dimensional span; it does not predict that the three innovation correlations $(\rho_{AX}, \rho_{sX}, \rho_{sA})$ should cluster near the determinant-zero surface

$$\rho_{sA} = \rho_{AX}\rho_{sX} \pm \sqrt{(1 - \rho_{AX}^2)(1 - \rho_{sX}^2)}, \quad (9.1)$$

nor does it determine which orientation branch should be selected. The persistent minus-branch selection is therefore not a restatement of ordinary factor concentration. It is a statement about the handedness and stability of the Robin boundary-correlation triangle.

The empirical novelty is consequently not that Treasury curves are low-dimensional in the familiar PCA sense. The novelty is that, after the curve is represented in Robin boundary coordinates, the innovations in those coordinates occupy a stable signed correlation manifold across broad regimes, blind calendar windows, holistic windows, human-selected windows, corrfine windows, and random partitions. The boundary map supplies the coordinate language; the segmentation-robust branch stability is the empirical fact.

10 Coordinate-system placebo tests

The PCA comparison raises a sharper empirical question: is the signed near-planar manifold specific to Robin coordinates, or does it appear automatically in any reasonable three-coordinate representation of the curve? I therefore run a fifth set of placebo tests. The same triangle diagnostics are applied to Nelson–Siegel coordinates, three-coordinate subsets of the Svensson parameterisation, simple endpoint–slope decompositions, and random orthogonal rotations of standardized Robin coordinates.

The purpose of these tests is not to require Robin to dominate every alternative on every isolated number. That would be the wrong benchmark. Nelson–Siegel coordinates are designed to describe curve shape, principal components are designed to capture variance concentration, and endpoint–slope triples can create exact algebraic rank reduction when one coordinate is a linear combination of the other two. The harder benchmark is the joint package: high triangle fit, low residual third-eigenvalue mass, low direct error, persistent branch behaviour, and a boundary interpretation that also connects to volatility and anchoring diagnostics.

Coordinate family	Window set	R^2	MAE	Median $\lambda_3/3$	Dominant branch	Interpretation
Nelson–Siegel $(\beta_0, \beta_1, \beta_2)$	Holistic	0.493	0.013	0.004	1.000	Low residual rank can appear, but triangle fit is not Robin-like across windows.
Nelson–Siegel $(\beta_0, \beta_1, \beta_2)$	Human	0.463	0.016	0.005	1.000	Good factor compression does not by itself reproduce the Robin signed geometry.
Svensson $(\beta_0, \beta_2, \beta_3)$	Holistic	0.895	0.021	0.001	0.647	Some subsets are strongly compressed, but branch behaviour is less stable.
Endpoint tautology $(Y_1, Y_{30}, Y_{30} - Y_1)$	Holistic	1.000	0.000	0.000	1.000	Exact rank two by construction; not economically informative.
Endpoint segmented slopes	Holistic	0.584	0.598	0.134	0.706	Non-tautological endpoint decompositions perform materially worse.
Random Robin rotations	Holistic	0.909	–	0.008	0.971	Rotation preserves some compression, but loses the boundary-coordinate interpretation.

Table 7: Coordinate-system placebo tests. Values are representative summaries from the Test 5 control families. Endpoint triples that include a coordinate such as $Y_{30} - Y_1$ are tautologically rank two and are not interpreted as economic evidence.

The placebo evidence supports a limited and defensible conclusion. The rank-two footprint is not unique in the trivial sense that no other coordinate system can exhibit low residual eigenvalue mass. Other curve coordinates can inherit ordinary factor compression, and some endpoint triples are mechanically rank two. What remains distinctive is the combination of signed branch stability, low direct prediction error, robustness across several window systems, and the fact that the same Robin coordinates also organise curve shape, volatility transport, and boundary-state behaviour. The comparison therefore strengthens the unification claim rather than a narrow uniqueness claim: alternatives can explain fragments of the empirical picture, while the Robin layer gives a coherent account of the bundle.

11 Bidirectional tests of the rank-two triangle

A natural concern is that the result might depend on choosing ρ_{sA} as the dependent correlation. The rank-two condition is symmetric, so any one off-diagonal correlation can be implied from the other two. In particular,

$$\rho_{sX} = \rho_{sA}\rho_{AX} \pm \sqrt{(1 - \rho_{sA}^2)(1 - \rho_{AX}^2)}. \quad (11.1)$$

I therefore repeat the test in the inverse direction, again selecting the branch with lower absolute error in each window.

Table 8 reports the initial inverse result. Across broad Robin subperiods, the implied ρ_{sX} explains 91.7 per cent of the observed variation. Across all holistic boundary-phase windows the fit is 95.8 per cent, and across the six core holistic windows it is 99.3 per cent. This should not be read as evidence that the holistic boundaries create the geometry. Rather, the finer windows are a descriptive segmentation on top of a correlation-space regularity that is already visible under broad and blind windows.

The inverse fit is not one-for-one. In this direction the slope is greater than unity and the intercept is positive. This should not be interpreted as a separate geometric law. Both the observed and triangle-implied correlations are estimated with noise, so reversing the dependent variable

Window set	Windows	Intercept	Slope	R^2	Corr.	MAE	Minus branch
Broad Robin subperiods	20	0.252	1.149	0.917	0.958	0.151	20/20
Holistic all windows	34	0.303	1.232	0.958	0.979	0.151	34/34
Core holistic windows	6	0.241	1.175	0.993	0.997	0.173	6/6

Table 8: Inverse rank-two test: observed ρ_{sX} regressed on implied $\rho_{sX}^{(-)}$ from (ρ_{sA}, ρ_{AX}) .

can create errors-in-variables and regression-to-the-mean asymmetry in the OLS summaries. This interpretation is checked directly below by bootstrapping daily innovations within each window and rerunning the forward and inverse regressions. The relevant evidence is therefore the joint pattern of high explanatory power, low direct error, stable branch selection, and small third-eigenvalue mass, not exact equality of the forward and inverse regression coefficients.

Table 9 reports the extended phase-specific covariance footprint. For the ρ_{sA} direction the R^2 range is 0.967–0.983 across all four window sets, and the minus branch is selected in every window. For the inverse ρ_{sX} direction the R^2 range is 0.930–0.977, with one plus-branch exception in the corrfine windows. The median residual third-eigenvalue share, $\lambda_3/3$, remains small at 0.024–0.032.

Window set	Direction	Windows	Intercept	Slope	R^2	Median abs. error	Minus branch
Broad Robin	ρ_{sA}	20	0.134	0.945	0.983	0.138	20/20
Broad Robin	ρ_{sX}	20	0.293	1.237	0.977	0.122	20/20
Corrfine	ρ_{sA}	30	0.155	0.940	0.967	0.136	30/30
Corrfine	ρ_{sX}	30	0.191	1.079	0.930	0.107	29/30
Holistic	ρ_{sA}	34	0.138	0.927	0.982	0.144	34/34
Holistic	ρ_{sX}	34	0.303	1.232	0.958	0.111	34/34
Human-selected	ρ_{sA}	22	0.132	0.928	0.983	0.137	22/22
Human-selected	ρ_{sX}	22	0.323	1.263	0.974	0.094	22/22

Table 9: Extended phase-specific rank-two covariance footprint. Each row regresses the observed correlation on the branch-selected triangle-implied correlation for the indicated direction.

The important point is that the planar identity works in both directions and the orientation branch remains essentially unchanged.




Figure 6: Observed versus implied ρ_{sX} under the inverse rank-two covariance triangle.

12 Random partition benchmark

The random partition benchmark is a direct guard against over-interpreting the phase boundaries. For each candidate window set I compare the actual windows with random partitions of the same length profile. The benchmark does not support the claim that the chosen phase boundaries significantly outperform random partitions on the own/best NMSE ratio. The corresponding p -values are 0.215 for the human-selected windows, 0.446 for the holistic windows, and 0.198 for the corrfine windows. The score statistic is more favourable but remains marginal rather than conventionally significant.

This is an important clarification, not a failure of the covariance result. It means that the rank-two geometry is not created by the window boundaries. Instead, the near-planar correlation structure is a global feature of the Robin coordinate system that remains visible under arbitrary partitioning. The residual third-dimensional mass is mildly favourable to the actual windows: actual median $\lambda_3/3$ values are 0.024–0.025, compared with random medians of 0.029–0.031. The corresponding permutation p -values are 0.124 for the human-selected windows, 0.074 for the holistic windows, and 0.264 for the corrfine windows, so this evidence should be described as directionally supportive rather than a decisive rejection of random partitions.

A second possible concern is mechanical dependence. The coordinates A_t , s_t , and X_t are all estimated from the same fitted yield curve, and A_t is defined relative to the long boundary. Since $A_t = x_{t,t}^R - X_t$, innovations satisfy $\Delta A_t = \Delta x_{t,t}^R - \Delta X_t$, which mechanically contributes a negative term to the covariance between ΔA_t and ΔX_t . The earlier robustness tests show that the rank-two footprint is not phase-selection-induced, but this definitional channel requires a more direct check.

Window set	Windows	Actual ratio	Random median	p_{ratio}	$\lambda_3/3$ actual/random	p_λ
Human-selected	22	0.323	0.393	0.215	0.024 / 0.031	0.124
Holistic	34	0.418	0.430	0.446	0.024 / 0.031	0.074
Corrfine	30	0.349	0.400	0.198	0.025 / 0.029	0.264

Table 10: Random partition benchmark for phase-specific prediction. The own/best NMSE ratio is lower when own-window prediction is better relative to the best alternative.

13 Identification checks: raw short boundary and bootstrap bias

Two targeted identification checks address the main remaining mechanical objections. First, I repeat the rank-two triangle using the raw fitted short-boundary level $x_{t,t}^R$ instead of the Robin amplitude $A_t = x_{t,t}^R - X_t$. This test asks whether the same result appears in $(x_{t,t}^R, s_t, X_t)$ without differentially subtracting the long boundary from one coordinate. Second, I bootstrap the nonlinear triangle map within each window to test whether the positive residual wedge is generated by finite-sample transformation bias in the square-root correlation formula.

The raw-short-boundary placebo does not reproduce the Robin-amplitude result. In the forward direction, the R^2 values fall to 0.343–0.517, compared with 0.967–0.983 for the corresponding Robin-




Figure 7: Random partition benchmark for phase-specific prediction. The actual windows do not significantly outperform random partitions on the own/best NMSE ratio, so the rank-two covariance footprint should be interpreted as segmentation-robust rather than segmentation-induced.

amplitude phase-specific tests. The third-eigenvalue share is also higher for the raw short-boundary coordinate, especially in broad Robin windows. The inverse raw-coordinate direction fits better, but its branch selection is no longer stable. This distinction is important. A high R^2 alone is not the identifying feature of the Robin result, because generic low-dimensional coordinates can sometimes generate strong pairwise prediction. The discriminating feature is the joint package of high fit, low direct error, low third-eigenvalue mass, and persistent signed branch selection. The raw short-boundary coordinate only reproduces part of this package. This is useful evidence: the near-planar signed triangle is not a generic consequence of estimating three variables from the same curve. It depends on the Robin amplitude coordinate, which is the boundary-relative object predicted by the theory.

The bootstrap check gives the opposite result for the residual wedge. Across the four phase-specific window sets, the median observed wedge is about 0.136–0.144, while the median absolute bootstrap transformation bias is only about 0.003–0.006. Thus the stable positive wedge is not materially explained by Jensen or delta-method bias in the nonlinear correlation map. This does not prove that the wedge is a structural friction term, but it rules out the simplest estimator-bias explanation.

14 Additional sensitivity: OLS asymmetry and small-window leverage

Two further checks close smaller but important inference issues. First, the difference between the forward and inverse OLS slopes could be mistaken for a new geometric fact. I therefore bootstrap daily innovations within each window and rerun both regressions across windows. The simulated

Window set	Raw x^R R^2	Raw x^R median error	Raw x^R $\lambda_3/3$	Median wedge	Median boot bias
Broad Robin	0.517	0.159	0.047	0.138	0.004
Corrfine	0.343	0.147	0.035	0.136	0.006
Holistic	0.494	0.143	0.033	0.144	0.003
Human-selected	0.419	0.135	0.031	0.137	0.005

Table 11: Test 6 identification checks. The raw x^R columns report the forward triangle test for $(x_{t,t}^R, s_t, X_t)$. The bootstrap columns report the original Robin-amplitude wedge and the median absolute finite-sample transformation bias from resampling daily innovations within each window.

slope gaps closely match the empirical gaps: the empirical inverse-minus-forward slope gap lies inside the bootstrap 90 per cent interval in all four window systems. This is a confirmatory diagnostic rather than a unique identification result, but it supports the interpretation that the forward/inverse slope asymmetry is consistent with errors-in-variables regression asymmetry and should not be read as a separate structural law.

Second, the smallest headline window sets are checked by leave-one-window-out regressions. The blind non-overlapping five-year result is highly stable: dropping any one of the nine windows leaves R^2 between 0.990 and 0.994. The six-window core holistic inverse test is also stable, with leave-one-out R^2 between 0.986 and 0.995. The core holistic forward test is somewhat more leverage-sensitive, with one leave-one-out value falling to 0.942, but the median leave-one-out R^2 remains 0.997. These results justify keeping the small-window summaries as diagnostics while maintaining the caution that they are not standalone inference.

Window set	Check fwd slope	Check inv. slope	Check gap	Sim. gap med.	Sim. gap 90% interval	Inside?
Broad Robin	0.945	1.237	0.292	0.274	[0.225, 0.332]	Yes
Corrfine	0.940	1.200	0.260	0.245	[0.201, 0.299]	Yes
Holistic	0.927	1.232	0.305	0.275	[0.220, 0.331]	Yes
Human-selected	0.928	1.263	0.335	0.310	[0.244, 0.386]	Yes

Table 12: Test 7A: bootstrap check for forward/inverse OLS asymmetry. The check slopes are the empirical forward and inverse slopes used in the bootstrap asymmetry diagnostic; they should not be read as a restatement of every slope reported in the earlier phase-specific summary tables, which aggregate different diagnostics. The gap is the inverse slope minus the forward slope. The simulation resamples daily innovations within each window and reruns both triangle regressions across windows.

Small-window summary	Windows	Base R^2	Leave-one-out R^2 range	Median LOO R^2	LOO slope range
Blind non-overlap 5y ρ_{sA}	9	0.992	0.990–0.994	0.991	0.838–0.864
Core holistic ρ_{sX}	6	0.993	0.986–0.995	0.995	1.153–1.208
Core holistic ρ_{sA}	6	0.997	0.942–0.999	0.997	0.927–1.006

Table 13: Test 7B: leave-one-window-out sensitivity for the smallest headline window sets. The results support the use of these rows as diagnostics, while the broader evidence remains anchored in larger rolling, phase-specific, random-partition, placebo, and identification tests.

15 Eigenvalue and residual-wedge interpretation

The constrained rank-two matrix has zero determinant by construction. Therefore the informative eigenvalue diagnostic is computed on the raw observed matrix (7.2). Let

$$\lambda_{1,j}^{\text{raw}} \geq \lambda_{2,j}^{\text{raw}} \geq \lambda_{3,j}^{\text{raw}} \quad (15.1)$$

be the eigenvalues of the observed correlation matrix in window j . Because $\text{tr}(R_R) = 3$, the residual third-dimensional mass is

$$\omega_{3,j} = \frac{\lambda_{3,j}^{\text{raw}}}{3}. \quad (15.2)$$

Small values of $\omega_{3,j}$ indicate that the observed correlation matrix lies close to a two-shock covariance system. Larger values indicate local rank-three behaviour, transition-window mixing, or misspecification of the boundary window.

The empirical evidence is consistent with a near-planar geometry rather than exact rank reduction. In the extended phase-specific tests, median $\lambda_3/3$ is about 0.032 for broad Robin windows, 0.025 for corrfine windows, and 0.024 for holistic and human-selected windows. These values are small but non-zero. They are below the corresponding random partition medians, although the permutation evidence is only directionally supportive rather than uniformly significant. They match the interpretation from the regression intercepts: the rank-two triangle is an organising surface, not an exact identity.

The residual intercept wedge should not be read as a failure of the rank-two geometry. Rather, it indicates that the correlation triangle describes the admissible covariance surface, while observed Treasury dynamics sit slightly and systematically away from the exact rank-two boundary. In this sense, the Robin coordinates behave like a near-planar structural sheet with a small positive residual thickness. The bootstrap check in Table 11 shows that nonlinear transformation bias in the estimated correlations is too small to explain this wedge. The consistently suppressed values of $\lambda_3/3$ show that the third covariance direction is economically minor even when the OLS regressions exhibit nonzero intercepts and non-unit slopes. The wedge may reflect higher-rank contamination, time-varying anchoring, liquidity or microstructure effects, or departures from the local affine/MSNR idealisation, but it does not alter the main conclusion: the dominant covariance geometry is two-dimensional.

16 Interpretation and contribution

The updated evidence changes the paper’s empirical contribution in an important way. The result is no longer simply that one correlation can be predicted from two others, nor that a particular hand-guided segmentation discovers a hidden relationship. The stronger claim is that the Robin boundary-coordinate layer organises several empirical objects that are usually modelled separately. The rank-two triangle is the sharpest diagnostic of this layer, but it is not the whole layer.

The economic interpretation is that Treasury dynamics contain both a short-boundary adjustment channel and a long-boundary anchoring channel. The amplitude coordinate A_t captures the scale of deviation from the long boundary, the short-slope coordinate s_t captures near-end adjustment pressure, and X_t captures the long-boundary anchor. A near-planar covariance manifold means that daily innovations in these three objects are not free to vary independently. Once the long-boundary co-movement and one short-boundary co-movement are known, the remaining correlation is largely pinned down by the triangle geometry. In financial terms, this is consistent with a market in which policy-rate expectations, term-premium adjustment, and long-anchor repricing are jointly constrained by a boundary condition rather than moving as three independent risk channels.

This interpretation is deliberately weaker than a full equilibrium derivation. The paper does not identify one Robin coordinate with a single institutional mechanism such as monetary policy, liquidity, or term premia. Instead, it shows that the reduced-form covariance structure behaves as if Treasury innovations are constrained by two dominant boundary-risk directions.

The Robin coordinates also provide a quantitative alternative to purely qualitative regime labels. Standard macro-finance discussions often describe policy environments using discrete categories such as tightening, easing, zero-lower-bound policy, quantitative easing, or inflationary stress. Such labels are economically useful, but they are not themselves state variables. In the Robin representation, changes in the policy environment appear through measurable boundary quantities, including anchoring strength χ , the pinned level x_D , and the local covariance structure of (A_t, s_t, X_t) . This does not imply that macroeconomic regimes disappear. Rather, it translates regime narratives into boundary behaviour.

The tangent-direction diagnostic clarifies a second distinction. Adjacent phase directions can be nearly orthogonal or even obtuse in (A, s, X) space. This says that phase means rotate. It does not contradict covariance planarity, because the rank-two triangle concerns the innovations around those local movements. The macro economy drives the coordinate means along wandering, non-persistent pathways, but the daily structural innovations around those pathways are governed by a rigid, near-planar boundary law. The empirical picture is therefore one of rotating local phase paths embedded in a persistent near-planar innovation covariance geometry.

The significance of the Robin framework is not that it dominates every alternative representation on every isolated metric. Nelson–Siegel coordinates can fit curve shape; principal components can compress variance; endpoint–slope decompositions can create mechanical rank reduction in special cases. The distinctive feature of the Robin representation is that the same boundary-coordinate object,

$$Z_t^R = (A_t, s_t, X_t),$$

appears across several empirical layers that are usually modelled separately. First, the Robin loadings provide a cross-sectional representation of the curve. Second, the same loading geometry transports boundary-coordinate covariance into the maturity term structure of realised volatility. Third, the local anchoring parameters describe how the short boundary is tied to the long boundary across monetary environments. Fourth, the innovation covariance of the boundary coordinates lies

close to a signed rank-two manifold with stable branch selection across segmentations and placebo tests.

The empirical contribution is therefore not simply another low-dimensional factor representation of Treasury yields. It is evidence that yield-curve shape, volatility transport, boundary anchoring, and innovation covariance share a common coordinate geometry. This is the empirical content of the word “unified” in Unified Boundary Field Theory.

17 Conclusion

The main result of this paper is not that Treasury yields are approximately low-dimensional. That fact is already familiar from standard factor models. The stronger finding is that a single Robin boundary-coordinate layer organises several objects that are normally treated as separate: curve shape, maturity-dependent volatility, boundary anchoring, innovation covariance, and local phase behaviour. The near rank-two covariance triangle is the sharpest diagnostic of this structure, but it is not the whole structure. It is one projection of the same boundary geometry that also governs the volatility and anchoring results.

The robustness evidence is important because it limits the interpretation in the right way. The covariance matrix is not exactly rank two at every date, the phase boundaries do not significantly outperform random partitions on every prediction criterion, and alternative coordinate systems can reproduce fragments of the evidence. These qualifications do not weaken the central claim. They clarify it. The empirical object is not a hand-selected segmentation or a single low-rank statistic. It is a segmentation-robust signed boundary-correlation manifold embedded in a broader Robin coordinate geometry.

The placebo evidence also clarifies what the paper is not claiming. Nelson–Siegel coordinates describe curve shape; principal components explain variance concentration; endpoint–slope triples can generate mechanical low rank. But these alternatives do not by themselves provide the same joint account of maturity geometry, volatility transport, boundary anchoring, signed covariance geometry, and boundary-state evolution. The relevant benchmark is therefore the whole empirical bundle, not any single metric in isolation.

The language of the yield curve is, to a significant degree, the language of its boundaries. In that sense, Paper 4 supports the central UBFT proposition: the same boundary-coordinate geometry that organises yield-curve shape also organises volatility transport, boundary-layer covariance, and local regime behaviour. The evidence should be read as a structural organisation result, not as a claim of exact rank reduction or a completed pricing theory.

Appendix A Branch Stability

Across the original ρ_{sA} test, the minus branch is selected in all broad Robin subperiods, all non-overlapping five-year windows, and all rolling five-year windows. In the initial inverse ρ_{sX} test, the minus branch is again selected in all broad subperiods, all holistic windows, and all core holistic windows. The extended phase-specific tests preserve this conclusion: all broad Robin, holistic, and human-selected windows select the minus branch in both directions, and all corrfine windows select the minus branch for ρ_{sA} . There is one corrfine exception in the inverse ρ_{sX} direction: window 3, 1987-04-22 to 1987-09-29, with $n = 112$, selects the plus branch, with absolute error 0.115. This short late-1987 window is the only branch break in the extended tests. Branch stability remains one of the most important empirical features of the result.

Window set	Windows	Plus best	Minus best	MAE minus
Robin subperiods	20	0	20	0.166
Non-overlap 5y	9	0	9	0.179
Rolling 5y quarterly	143	0	143	0.185

Table 14: Branch stability in the original ρ_{sA} rank-two test.

Appendix B Implementation Notes

All reported correlations use innovations in fitted Robin boundary coordinates. The broad Robin subperiods are estimated from local slope-level anchoring relations. The holistic boundary-phase windows are diagnostic windows based on joint movement in (A, s, X) and are used to describe locally coherent boundary phases. The rank-two identities are evaluated using observed local correlation matrices, and eigenvalue diagnostics are computed on the raw matrices rather than on constrained matrices. Random partition benchmarks use partitions with the same length profile as the corresponding actual window set, so they test whether the observed covariance geometry depends on the chosen phase boundaries.

References

- Unified Boundary Field Theory, Paper 1. (2026). *An Axiomatically Consistent Single-Factor Term Structure Model with Admissible Humped Volatility*. Zenodo. [doi:10.5281/zenodo.20943574](https://doi.org/10.5281/zenodo.20943574).
- Unified Boundary Field Theory, Paper 2. (2026). *Unified Boundary Field Theory of Interest Rates II: Monetary Policy as a Short-End Boundary Condition*. Zenodo. [doi:10.5281/zenodo.20963161](https://doi.org/10.5281/zenodo.20963161).
- Unified Boundary Field Theory, Paper 3. (2026). *Unified Boundary Field Theory of Interest Rates III: The Stochastic Long Boundary: Volatility-Aligned Geometry, Affine Admissibility, and Two-Channel MSNR for Instantaneous Rolling Forward Rates*. Zenodo. [doi:10.5281/zenodo.21169397](https://doi.org/10.5281/zenodo.21169397).

- Bianchi, F., Ludvigson, S. C., and Ma, S. (2022). Monetary-Based Asset Pricing: A Mixed-Frequency Structural Approach. Working paper.
- Ritchken, P. and Chuang, I. (2000). Interest rate option pricing with volatility humps. *Review of Derivatives Research*, 3, 237–262.
- Svensson, L. E. O. (1994). Estimating and interpreting forward interest rates: Sweden 1992–1994. *IMF Working Paper*.