

Nested Robin Boundary States and the Geometry of Yield-Curve Regimes

Companion Note to UBFT Paper IV

S. Oza

Working Paper, July 2026

Abstract

Paper IV of the Unified Boundary Field Theory of Interest Rates [4] studies US Treasury curve *innovations*: the day-to-day, conditional movements of the fitted curve around its current Robin boundary representation, rather than the level of the boundary layer itself. In that setting, Robin boundary coordinates organise a stable rank-two covariance geometry for curve innovations. This companion note asks a different question. Motivated by apparent sub-period structure among the fitted short-boundary slope, amplitude, and long-boundary anchor, it asks whether the fitted Robin boundary layer also contains persistent state structure in levels and locations. Starting from broad slope-amplitude regions in (s, A) , each region is subdivided using the slope-long-boundary projection (s, X) . The resulting nested states are persistent, macro-related but not macro-redundant, and economically meaningful. They exhibit a one-day stay probability of 0.983 compared with a memoryless benchmark of 0.107, and a mean dwell time of roughly 60 trading days compared with about one day under the same benchmark. Their normalised mutual information with broad macro regimes is 0.469, far above random-label baselines, while the states also resolve substantial within-macro variation. Most importantly, the nested states condition the observed volatility term structure: state-conditioned volatility curves separate into short/belly-hump, intermediate/balanced, and long-boundary maximum or tilt families. The long-boundary cases are not interpreted as literal 30-year interior humps; rather, they are weak-anchoring states in which realised volatility is transmitted toward the long boundary. The states also condition portfolio risk and add substantial explanatory power for curve shape after macro-regime dummies are included. The evidence supports the interpretation that the Robin layer is a coupled boundary regime with nested structure: broad slope-amplitude regions contain finer short-long boundary substates that persist over economically meaningful horizons and condition the observable curve and risk environment.

Keywords: yield curve, Robin boundary condition, boundary states, volatility term structure, covariance geometry, regime classification, Treasury curve.

JEL codes: E43, G12, C58, E52, G17, C32.

Acknowledgements

The author acknowledges the use of OpenAI’s ChatGPT as a research and drafting assistant for code generation, empirical checks, exploratory analysis, figure preparation, and manuscript drafting support. Responsibility for the theory, interpretation, empirical claims, and final manuscript remains with the author.

1 Introduction

Paper IV of the Unified Boundary Field Theory of Interest Rates [4] shows that Robin boundary coordinates organise US Treasury curve innovations into a robust low-rank covariance geometry. In that paper, an *innovation* means a conditional movement of the curve around its fitted boundary representation: roughly, the daily change or residual displacement of the curve after the prevailing Robin boundary location has already been fixed. The object of interest is therefore not the level of the curve itself, but how the curve moves locally around the boundary layer.

This note studies the complementary object. It asks whether the fitted boundary layer itself contains meaningful *location states*. In other words, once the curve has been represented by Robin boundary coordinates, does the level-location geometry of the boundary layer contain persistent regimes that matter for observed curve shape and risk?

The answer appears to be yes. The evidence developed below suggests that the Robin layer behaves as a coupled boundary regime with nested structure. Broad regions in slope-amplitude space (s, A) contain finer substates in short-long boundary space (s, X) . These substates persist for economically meaningful intervals, are related to macro history without being reducible to macro dates, and condition the maturity location of realised yield-curve volatility.

This distinction is central. The states studied here are not proposed as replacements for the rank-two covariance geometry in Paper IV. They are not covariance-homogeneous blocks, nor are they shown to be standalone forecasting regimes. Their role is different: they describe where the Robin boundary layer sits. Paper IV describes how innovations move around that boundary geometry. The two objects are complementary:

boundary-location states and innovation covariance geometry.

This note is intended as a standalone companion working paper rather than an appendix to, or replacement for, Paper IV. For that reason the state construction and diagnostic checks are described explicitly below, while the full daily labels and code are included in the replication package.

The principal empirical payoff is the volatility-shape result. State-conditioned volatility curves do not share a fixed maturity profile. Instead they separate into short/belly-hump states, intermediate or balanced states, and long-boundary maximum or tilt states. For visual presentation the main figures show the 1Y-20Y segment, while the 30Y endpoint is retained in the classification tables. This prevents a terminal 30Y maximum from being visually overstated as an interior “hump.” The long-end cases are better interpreted as weak-anchoring, long-boundary-dominated states.

The note proceeds as follows. Section 2 defines the Robin coordinates and the nested state construction. Section 3 documents persistence and transition structure. Section 4 shows that the states are macro-related but not macro-redundant. Section 5 presents the volatility-shape result. Section 6 shows portfolio-risk consequences. Section 7 reports curve-shape and covariance conditioning, including macro-adjusted explanatory power. Section 8 documents the ordered volatility-shape transition constraint. Section 9 explains what the states are not. Section 10 concludes.

2 Robin boundary coordinates and nested state construction

Let the fitted Robin boundary representation of the curve be summarised by three daily coordinates,

$$(A_t, s_t, X_t),$$

where s_t is the local short-boundary slope or short-end pressure, X_t is the long-boundary anchor, and A_t is the amplitude displacement relative to that anchor. A useful derived quantity is the short-boundary Robin coordinate

$$x_t^R = A_t + X_t.$$

Local Robin anchoring can then be summarised by an empirical regression of the form

$$s_t = -\chi x_t^R + \chi x_D + \varepsilon_t, \tag{1}$$

where χ measures the strength of local short-boundary anchoring and x_D is the implied Dirichlet reference location.

The nested-state construction is deliberately sequential. First, the sample is partitioned into broad regions in (s, A) space. These are slope-amplitude regions: they identify where the short-boundary slope and amplitude displacement place the curve in a broad boundary-location manifold. Second, within each broad (s, A) region, the observations are subdivided using the short-long boundary projection (s, X) . The resulting double labels are denoted $i.j$, where i indexes the broad slope-amplitude region and j indexes the short-long substate.

The first-stage (s, A) labels are taken as fixed inputs from the empirical Robin segmentation used in Paper IV; they are not re-estimated in this companion note. Conditional on those labels, the second-stage subdivision is algorithmic: within each first-stage segment, (s, X) is globally standardized and a full-covariance Gaussian mixture is fit using $K = 1, 2, 3$ candidate components, BIC selection, three initializations, fixed random seed 2026, and a 3% minimum component-share screen. The amplitude coordinate A is not used in the second-stage split. Components are relabeled by increasing mean X and then mean s for stable display labels. Persistence, volatility shape, macro overlap, and transition ordering are evaluated only after these labels have been fixed.

This construction is related to, but distinct from, standard low-dimensional yield-curve representations. Level-slope-curvature factors, Nelson-Siegel-type coordinates, principal components, Markov-switching models, and curve-clustering methods all provide ways to reduce or classify yield-curve shapes. The claim here is not that a two-dimensional projection of slope against amplitude or slope against a long-end coordinate is unique to UBFT. The contribution is narrower: the coordinates are fitted Robin boundary coordinates, the subdivision is explicitly nested, and the resulting labels are tested as boundary-location states with persistence, macro non-redundancy, volatility-shape consequences, and ordered transition topology.

The construction is therefore not ordinary clustering on a generic factor vector. It reflects the coupled boundary interpretation. The broad (s, A) region locates slope-amplitude geometry, while the within-region (s, X) subdivision identifies how the short boundary and long anchor are configured inside that broad location. The double label therefore represents a nested boundary state rather than a black-box cluster.

Figure 1 shows the double states projected into both the (s, A) and (s, X) planes. The key visual point is that the double labels remain coherent across both projections. They are learned through

the nested construction, but they do not collapse when viewed in the alternative boundary plane.




Figure 1: Nested Robin boundary states. Broad amplitude-slope regions are subdivided using short-long boundary coordinates. The resulting labels remain coherent when projected back into both (s, A) and (s, X) space.

The double states explain a large share of coordinate variation relative to random labels: approximately 0.609 of the variation in A , 0.791 of the variation in s , and 0.801 of the variation in X . This confirms that the states are location states in Robin boundary space. It does not, by itself, imply that they are innovation covariance states; that distinction is tested below.

3 Persistence and transition structure

The strongest first check is temporal persistence. If the double states were merely noisy daily partitions, transition behavior should resemble a random relabeling with the same marginal state frequencies. It does not.

Table 1 compares observed transition statistics with a memoryless benchmark preserving marginal state frequencies. The observed one-day stay probability is 0.983, compared with a benchmark value of 0.107. Mean dwell time is approximately 60 trading days, compared with roughly 1.12 days under the memoryless benchmark. The maximum observed dwell is 1419 trading days, versus a random median of 5 days.

Table 1: Persistence of nested Robin boundary states versus a memoryless benchmark

Metric	Observed	Benchmark	5% sim.	95% sim.
One-day stay probability	0.983	0.107	0.102	0.112
Mean dwell time	59.923	1.120	1.114	1.126
Median dwell time	5.000	1.000	1.000	1.000
Maximum dwell time	1419.000	5.000	5.000	7.000
Number of runs	169.000	9041.000	8990.000	9094.000
Transition HHI	0.104	0.012	0.011	0.012
Top transition share	0.201	0.042	0.038	0.046

Figure 2 plots the transition matrix. The diagonal dominance is the visual counterpart to the stay-probability result. These states are not high-frequency labels attached to daily noise. They are persistent boundary-location episodes.




Figure 2: Transition matrix for nested Robin boundary states. The diagonal dominance reflects substantial state persistence.

The persistence result is central to the interpretation. The double states are not simply a way to draw a denser scatterplot. They identify regions of the boundary layer in which the curve remains for extended periods and from which it transitions in a structured manner.

4 Macro relation without macro redundancy

The nested states are not macro-neutral. They are meaningfully related to broad macro periods, but they are not equivalent to those periods. The normalised mutual information between nested states and macro-regime labels is approximately 0.469 using the geometric normalisation. A random-label benchmark preserving state sizes has median normalized mutual information close to zero and an empirical probability of matching the observed value of approximately 0.0002.

Metric	Observed	Random median	Random 5%	Random 95%	$p(\text{random} \geq \text{obs})$
Mutual information, nats	0.814	0.002	0.001	0.003	0.0002
NMI, arithmetic	0.448	0.001	0.001	0.002	0.0002
NMI, geometric	0.469	0.001	0.001	0.002	0.0002
NMI, minimum	0.639	0.002	0.001	0.002	0.0002

Table 2: Information overlap between nested states and macro regimes

This result cuts two ways. On one hand, the states are not arbitrary: they overlap with recognisable macro history. On the other hand, the overlap is incomplete. The states resolve within-regime boundary variation. In the empirical sample, the pre-GFC period contains nine distinct double states and the post-2021 tightening period contains six. That resolution is economically meaningful because broad macro dating alone cannot identify how the Robin boundary layer is configured inside those periods.




Figure 3: Overlap between nested Robin boundary states and broad macro regimes. The states are macro-related but not one-to-one macro labels.

The key conclusion is therefore not that macro history is irrelevant, it is that macro history is insufficient. Nested Robin states capture a boundary-layer geometry that is partly aligned with macro periods but not reducible to them.

5 Volatility-shape consequence: from short/belly humps to long-boundary tilt

The main economic consequence of the nested states appears in the maturity shape of realised volatility. For each state, realised daily yield-change volatility is computed across the observed maturity grid. The state-level volatility maximum and the ratio of 2Y to 30Y volatility are then used to classify the state-conditioned volatility term structure.

A terminology point is important. When the maximum occurs at 30Y, this is not interpreted as a literal interior 30Y hump. Since 30Y is the endpoint of the grid, the correct interpretation is a long-boundary maximum or long-side tilt. The visual figures therefore display the 1Y-20Y segment, while the 30Y endpoint is retained in the classification tables.

The state-conditioned volatility curves separate into three families:

- (i) short/belly-hump states, where volatility is concentrated around the 3Y-5Y region;
- (ii) intermediate or balanced states, where the term structure is flatter or transitional;
- (iii) long-boundary maximum or tilt states, where volatility is transmitted toward the long side or long boundary.

Table 3: Volatility-curve families and Robin anchoring

Family	States	Median χ	Median anchor R^2	Max-maturity range	Median 2Y/30Y	Median max vol
Short/belly hump	6	0.099	0.338	3Y-5Y	1.10	105.9
Intermediate/balanced	2	0.019	0.116	5Y-7Y	0.92	91.1
Long-boundary max/tilt	4	0.026	0.051	14Y-30Y	0.62	99.0

Table 3 shows that the long-boundary family has weak median short-boundary anchoring: median χ is approximately 0.026 and median anchoring R^2 is approximately 0.051. The short/belly family has stronger median anchoring. This supports the interpretation that long-side maximum states are weak-anchoring states, not ordinary short-end-dominated volatility regimes.




Figure 4: Observed state-conditioned volatility curves, displayed over the 1Y-20Y window. The 30Y endpoint is retained in classification tables but omitted from the figure to avoid visually overstating a terminal maximum as an interior hump.




Figure 5: Normalised average volatility shapes by family over the 1Y-20Y window. The short/belly family exhibits an interior concentration, while other families flatten or tilt toward the long side.

The most careful statement is therefore not that “the volatility hump migrates to 30Y.” Rather:

The state-conditioned volatility maximum migrates across shape families. Some states exhibit an interior short/belly hump, while others are better described as weak-anchoring long-boundary maximum or long-side tilt states.

This is the most theoretically load-bearing result in the note. It links the nested boundary-state classification to volatility transport along the maturity axis. The boundary state is not merely a label attached to the curve. It conditions where realized risk appears along the curve.

6 Portfolio-risk consequence

The volatility-shape result has direct implications for risk measurement. If state-conditioned volatility profiles differ, the same stylised exposure should face different realised risk depending on the boundary state. That is exactly what the portfolio-risk summaries show.

Across nested states, the ratio of maximum to minimum annualized realized risk is approximately 3.10 for a 2Y rate-change exposure, 2.14 for a 10Y rate-change exposure, and 4.21 for a 5s10s30s butterfly exposure. These are not small differences. They imply that a portfolio with the same static exposure can occupy very different risk environments depending on the Robin boundary state.

Exposure	Minimum ann. vol	Maximum ann. vol	Max/min ratio
2Y rate-change exposure	42.7	132.6	3.10
10Y rate-change exposure	65.3	139.4	2.14
5s10s30s butterfly exposure	30.0	126.1	4.21

Table 4: Representative state-conditioned portfolio-risk dispersion

Table 4 reports the ratio summary rather than all state-level entries. The practical message is simple: nested Robin boundary states classify risk environments, not just curve shapes. The visual heatmap in Figure 6 displays the state-by-exposure risk structure.



Figure 6: State-conditioned realized risk across stylized rate, slope, and butterfly exposures. The same exposure can face materially different realized risk across nested Robin boundary states.

For a practitioner, this is the most direct consequence of the boundary-state construction. A state label identifies not only where the curve sits in Robin boundary space but also the kind of maturity-risk environment in which curve exposures are being carried.

7 Curve-shape and covariance conditioning

The nested states also condition observed curve shape. A natural objection is that this may merely reflect macro-period overlap. To address this, curve-shape regressions compare macro-regime dummies, state dummies, and macro-plus-state dummies. The key statistic is incremental state explanatory power after macro dummies have already been included.

Target	Macro R^2	State R^2	Macro+state R^2	Incremental state R^2
10Y level	0.698	0.744	0.907	0.209
2s10s slope	0.411	0.610	0.678	0.267
5s30s slope	0.377	0.603	0.700	0.323

Table 5: Macro-adjusted curve-shape explanatory power

These incremental values are substantial. They show that nested boundary states are not merely refined macro-calendar labels. They add within-macro-regime information about curve level and slope structure. Because the macro-plus-state specification adds eleven state dummy variables to the macro-only regression, Table 5 should be read as a descriptive contemporaneous decomposition rather than as a parsimonious forecasting exercise. The replication package therefore also reports adjusted R^2 and nested-model F -tests for the incremental state block; for the three headline curve-shape targets, the adjusted incremental R^2 values remain essentially unchanged and the incremental state-block F -tests are overwhelmingly significant.




Figure 7: Incremental explanatory power of nested state dummies after macro-regime dummies. The strongest macro-adjusted effects occur in curve-shape targets.

The covariance-surface evidence is more subtle. In yield space, state-conditioned covariance surfaces differ: first eigenvalue shares, first-two-eigenvalue shares, and long-short correlations vary across states. This is consistent with the idea that different boundary locations project into different maturity-space risk surfaces.

However, these states do not materially sharpen the rank-two covariance triangle in Robin coordinates relative to random partitions. That is not a failure of the state construction; it clarifies what the states are. The double states are boundary-location episodes. The Paper IV covariance triangle is an innovation geometry: it describes the covariance of local curve movements around the

fitted boundary state, not the unconditional distribution of boundary locations. Location states can shift the projected maturity-risk surface without being locally covariance-homogeneous in the Robin innovation coordinates.




Figure 8: State-conditioned covariance eigenvalue shares in yield space. Boundary-location states shift projected maturity-space risk surfaces even though they do not replace the Robin-coordinate covariance geometry of Paper IV.

This distinction is important enough to state explicitly:

The nested states describe where the boundary layer sits; the rank-two covariance geometry describes how local curve innovations move around that boundary layer.

8 Volatility-shape transitions and ordered boundary transport

The preceding sections establish that the nested states are persistent, macro-related but not macro-redundant, and economically meaningful for volatility shape, portfolio risk, and curve geometry. A stronger question is whether the three volatility-shape families also obey a sequential transition structure. To test this, daily family labels are compressed into family-level episodes and transitions are counted between consecutive episodes.

Across 63 family-level episode transitions, the volatility-shape families obey a strict adjacency constraint. The system never transitions directly between short/belly-hump and long-boundary-tilt environments. All observed movement between these endpoints passes through the intermediate/balanced family. The result is exact in the sample: the short/belly-to-long-boundary and long-boundary-to-short/belly cells of the episode transition matrix are both zero (1985–2026).

From / To	Short/belly	Intermediate	Long-boundary
Short/belly	0	9	0
Intermediate	9	0	22
Long-boundary	0	23	0

Table 6: Family-level episode transition counts

This adjacency constraint is not an artefact of the marginal distribution of family episodes. In random spell re-orderings preserving the number of spells in each family, the median number of adjacent transitions is 32, compared with 63 observed, while the median number of direct short-long jumps is 6, compared with 0 observed. The observed transition graph is therefore not merely persistent or unbalanced; it is sequentially constrained.

Metric	Observed	Random median	Random 5%	Random 95%
Adjacent transitions	63	32	26	39
Direct short-long jumps	0	6	3	10
Adjacent share	1.000	0.508	0.413	0.619
Direct-jump share	0.000	0.095	0.048	0.159

Table 7: Ordered-transition random baseline

The coordinate movements are consistent with ordered boundary transport. Transitions toward the long-boundary family are associated with falling χ and falling 2Y/30Y volatility ratios, while transitions back toward short/belly states are associated with rising χ and renewed short-end volatility dominance. This pattern appears in both medians and means, indicating that the result is not driven by a skewed transition distribution.

Transition	Mean $\Delta\chi$	Median $\Delta\chi$	Mean $\Delta(2Y/30Y)$	Median $\Delta(2Y/30Y)$
Short/belly $\rightarrow$ Intermediate	-0.164	-0.147	-0.373	-0.320
Intermediate $\rightarrow$ Short/belly	0.160	0.147	0.374	0.320
Intermediate $\rightarrow$ Long-boundary	-0.077	-0.084	-0.372	-0.353
Long-boundary $\rightarrow$ Intermediate	0.074	0.084	0.372	0.353

Table 8: Coordinate changes across family-level episode transitions




Figure 9: Episode-level transition graph for volatility-shape families. Direct short/belly–long-boundary transitions are absent in the sample; all endpoint movement passes through the intermediate/balanced family.

The balance between transitions toward the long-boundary family and transitions back toward the short/belly family should not be read as a lack of structure. The counts are nearly balanced, 31 versus 32, which is consistent with a stationary full-sample classification rather than a one-way drift in the boundary layer. The structure is not directional trend; it is the sequential constraint on the transition graph and the associated movement of χ and relative short-long volatility.

This empirical adjacency constraint is consistent with, but not derived from, the Paper 3 [3] maturity-loading geometry, in which the short-boundary, belly, and long-boundary channels form a maturity continuum rather than disconnected regimes. The intermediate/balanced family can therefore be interpreted as an empirical passage state between the short-boundary and long-boundary limits, in line with the role of the $\tau e^{-a^2\tau}$ interior loading in the maturity-space Jacobian $J(\tau)$, which carries boundary deformation between the short and long endpoints without reaching either limit.

9 What the nested states are not

The results above support a structural and descriptive interpretation. They do not yet support all possible stronger claims.

First, the nested states are not standalone forecasting regimes. They are optimized to classify the concurrent location and geometry of the Robin boundary layer, not to act as directional momentum signals or unconditional predictors of the next volatility block. In-sample state dummies explain a material share of next-60-day average curve volatility, but time-block cross-validation is negative. This means that the historical state-volatility association does not yet translate into a robust out-of-sample forecasting claim. The likely reason is that states are persistent and partially macro-overlapping: the same state label can recur in different macro contexts with different subsequent

volatility outcomes.

Second, the nested states are not covariance-homogeneous blocks in Robin-coordinate innovation space. Here again, “innovation” refers to the local movement of the curve around the current fitted boundary configuration. The earlier rank-two covariance tests remain the primary evidence for innovation geometry. The boundary states should be interpreted as location regimes that condition curve shape and maturity-risk surfaces, not as replacements for the covariance manifold.

Third, the long-boundary volatility family should not be described as a 30Y interior hump. A maximum at the 30Y endpoint is a terminal long-boundary maximum or long-side tilt. This language is more accurate and more consistent with the boundary interpretation.

These limits strengthen rather than weaken the note. They prevent overclaiming and allow the central result to be stated cleanly: nested Robin boundary states are persistent, macro-related but not macro-redundant, and economically meaningful because they condition curve shape, volatility location, and realized risk.

10 Conclusion

This companion note extends the empirical Robin boundary framework from innovation covariance geometry to boundary-location state structure. The evidence supports a coupled boundary-regime interpretation. Broad amplitude-slope regions contain finer short-long boundary substates. These substates persist, overlap with macro history without collapsing into macro labels, and condition observable curve and risk properties.

The volatility-shape evidence is the clearest theoretical payoff. State-conditioned volatility curves separate into short/belly-hump, intermediate/balanced, and long-boundary maximum or tilt families. The long-boundary cases are weak-anchoring states rather than literal 30Y interior humps. This is consistent with a boundary-transport interpretation: when short-boundary anchoring is weak, realized volatility can be transmitted toward the long side of the curve. The family-level transition evidence strengthens this interpretation: across 63 episode transitions, the system never jumps directly between the short/belly and long-boundary endpoints, but instead passes through the intermediate/balanced family.

The broader empirical architecture is therefore:

Robin boundary layer $\rightarrow$ nested boundary-location states $\rightarrow$ curve shape, volatility location, covariance surface, a

Together with Paper IV, the result suggests that the Robin representation contains two complementary empirical objects: a stable covariance geometry for local curve innovations and a persistent nested state geometry for the boundary layer itself.

Appendix A Summary of state-level volatility groups

Table 9: State-level volatility-shape classification and anchoring summary

State	Family	Max maturity	2Y/30Y ratio	χ	Anchor R^2
2.2	short/belly hump	3Y	1.10	-0.011	0.040
2.3	short/belly hump	3Y	1.17	0.155	0.564
4.2	short/belly hump	3Y	1.43	0.203	0.920
4.3	short/belly hump	4Y	1.11	0.056	0.088
3.2	short/belly hump	5Y	1.01	-0.046	0.112
1.1	short/belly hump	5Y	1.02	0.143	0.872
4.1	intermediate/balanced	5Y	0.93	-0.013	0.004
1.2	intermediate/balanced	7Y	0.91	0.051	0.229
2.1	long-boundary max/tilt	14Y	0.51	-0.013	0.003
3.1	long-boundary max/tilt	30Y	0.56	-0.003	0.000
1.3	long-boundary max/tilt	30Y	0.67	0.056	0.099
3.3	long-boundary max/tilt	30Y	0.91	0.081	0.138

Small negative estimates of χ in a few short/belly-hump states should not be read as evidence of strongly reversed anchoring. Their magnitudes are close to zero and their anchoring R^2 values are weak. They indicate locally weak or unstable short-boundary anchoring, while the short/belly classification is based on the observed maturity location of realised volatility rather than on χ alone.

Appendix B State-construction procedure

Step	Specification
First-stage regions	Use the pre-existing Paper IV empirical Robin segmentation in the slope-amplitude plane (s, A) as fixed input labels. These labels are not re-estimated in this companion note.
Second-stage variables	Within each first-stage region, subdivide observations using the short-long boundary projection (s, X) only. The amplitude coordinate A is not used in the local second-stage split.
Standardisation	Standardise (s, X) before fitting local mixture components.
Local mixture fit	Fit full-covariance Gaussian mixture models with candidate component counts $K = 1, 2, 3$, three initialisations, BIC selection, fixed random seed 2026, and a 3% minimum component-share screen.
Label display	Relabel components by increasing mean X and then mean s to obtain stable display labels $i.j$.
Post-label diagnostics	Evaluate persistence, macro overlap, volatility-shape families, portfolio risk, curve-shape explanatory power, and ordered transition structure only after the labels have been fixed.

Table 10: Nested boundary-state construction procedure

Appendix C Reproducibility notes

The empirical package accompanying this draft contains the fitted Robin coordinates, nested state labels, state summaries, volatility-curve grouping outputs, ordered family-transition checks, macro-adjusted explanatory-power checks, transition and dwell-time tables, and plotting scripts. All random-label and clustering procedures should be run with fixed seeds in the public replication package. The preferred figure for the volatility-shape section is the 1Y-20Y window; the 30Y endpoint should remain in the classification tables but should not be used to suggest a literal terminal interior hump.

References

- [1] Oza, S. Unified Boundary Field Theory of Interest Rates I: An Axiomatically Consistent Single-Factor Term Structure Model with Admissible Humped Volatility, June 2026, <https://doi.org/10.5281/zenodo.20943574>.
- [2] Oza, S. Unified Boundary Field Theory of Interest Rates II: Monetary Policy as a Short-End Boundary Condition, June 2026, <https://doi.org/10.5281/zenodo.20963161>
- [3] Oza, S. 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, June 2026, [10.5281/zenodo.21169397](https://doi.org/10.5281/zenodo.21169397).
- [4] Oza, S. Unified Boundary Field Theory of Interest Rates IV: Robin Short-Rate Boundary Geometry and Empirical Low-Rank Covariance Structure, July 2026, <https://doi.org/10.5281/zenodo.21193851>.