The κ-Framework Meets the Data: An Out-of-Sample Empirical Validation Across Seventeen Data Providers

Shehzad Ahmed* Rafiqul Bhuyan†

June 2026 — Preliminary draft

Abstract

The κ-rate research programme (Papers 1–14) proposes the convergence intensity κ of Ackerer et al. [2025] as a riba-free pricing primitive and develops it across perpetual futures, sukuk, takaful, credit protection, and a monetary architecture. This paper subjects that corpus to an *independent, out-of-sample* test. We assemble a 17-provider dataset — Cbonds (sukuk, conventional bonds, CDS, sovereign indices, local-currency curves), dYdX v4 / Binance / Bybit perpetual funding, FRED, ECB, BIS, OECD, the World Bank, IMF, WHO, the Human Mortality Database, EM-DAT, FAOSTAT, NASA POWER, CoinGecko and DeFiLlama — spanning ∼1.1 GB across credit, rates, perpetual funding, commodities, tokenized real-world assets, and six cross-domain hazard channels at 150–260 countries. On this dataset we run eighty-one pre-specified analyses (specified in the replication archive's analysis plan before execution; we say "pre-specified" rather than "pre-registered" because no external registry timestamp exists — the plan, its A1–A90 numbering, and the nine analyses dropped as infeasible are all documented in the archive), each a statistic the original papers did not themselves compute. We adopt a strict discipline: before treating any disagreement as a discrepancy we reproduce the source paper's own number from its replication archive, confirming that headline figures (e.g. the cross-sectional sukuk–CDS correlation of 0.97) are correct for what they measure and that our figures are *complementary* statistics, not corrections. The corpus's central claims survive: κ is a real, venue-independent (ρ=0.70 across exchanges), non-negative (0.01% of observations negative), credit-quality-ordered, upward-term-structured (12/12 sovereigns), mean-reverting intensity that is *not* the interest rate r (it spans six orders of magnitude, 0.002 to ∼1,700/yr at the liquid-perpetual median); the interest parameter ι is an additive, removable layer worth exactly the 0.01%/8h CEX baseline; and cross-domain diversification is genuine ($\bar{\rho}_x = 0.05$ globally). We also report one finding new to the programme: a first cross-domain *tail* measurement (COVID-19) in which sovereign credit and mortality intensities co-moved, showing that crisis co-movement can exceed the small average cross-domain correlation along channel-specific lines. Finally, we supply the first independent empirical companion to the *monetary* half of the corpus (Papers 8–9), which ships replication code but had no out-of-sample test: the active-fiscal determinacy ceiling $\gamma_s^* = e^{\bar{\kappa}} - 1$ reproduces and extends to twenty-six sovereigns across two independent instruments; the κ-Standard money base contracts ∼2% per +100 bp on 860 bonds (matching the paper's −1.99%); and the benchmark-free κ estimator ranks 860 bonds identically to the spread-based one, differing only by an additive constant equal to the reserve currency's own yield layer mapped through the same transformation — which, by the

*Baraka Protocol; Department of Finance, Independent University Bangladesh. shehzadahmed@arcusquantfund.com.

†Department of Accounting and Finance, Alabama A&M University.

programme's own decomposition, is that currency's time-value and premium content, not a hazard — direct evidence that the conventional benchmark is a decomposition for *relative* pricing, not a dependency. We conclude that the framework's direction is robustly supported on independent data, and we delineate exactly which numbers are construction-specific and must be quoted with their method attached.

Keywords: convergence intensity, riba-free pricing, out-of-sample validation, perpetual futures funding, sukuk, sovereign CDS, cross-domain correlation, tokenized real-world assets.

1 Introduction

A theory earns confidence not when its authors restate it but when independent data, gathered for other purposes and analysed with statistics the authors did not choose, fail to break it. The κ-rate programme makes an unusually testable set of claims: that a single convergence intensity κ — the rate at which a perpetual contract converges to its underlying, equivalently the Duffie–Singleton default hazard at the riba-free point $\iota = 0$ [Ahmed & Bhuyan, 2026b] — is observable from markets, non-negative, credit-ordered, term-structured, mean-reverting, and categorically distinct from the interest rate r, while the interest parameter ι is a separable layer that can be removed without breaking no-arbitrage [Ackerer et al., 2025]. Each of these is a measurement, not a matter of interpretation.

This paper tests them out of sample. We assemble an independent dataset from seventeen providers (Section 2) that the original fourteen papers did not use, and run eighty-one analyses (Sections 4–9), each computing a quantity the source papers did not. Our governing methodological rule (Section 3) is that a disagreement is not a discrepancy until the source paper's own number has been reproduced from its replication archive and the two statistics shown to measure the same object; applying it dissolves three apparent disagreements into benign differences of construction. The headline result is that the corpus survives: the confirmations are broad and the two genuine refinements are (i) several headline magnitudes are construction-specific and must be quoted with their method, and (ii) average cross-domain correlation understates crisis co-movement, a point we document with the first cross-domain tail measurement and feed back into Paper 14.

2 Data

Table 1 summarises the dataset (full catalogue and per-series schema in the replication archive's `MASTER_DATA_MAP`). Coverage is deliberately broad: three perpetual-futures venues (one at $\iota = 0$, two at $\iota > 0$); the full sukuk and conventional sovereign credit stack in USD and four local currencies; three independent risk-free curve systems (US Treasury, euro-area AAA, and local-currency government curves); sovereign-credit fundamentals; six cross-domain hazard channels at 150–260 countries; and the tokenized real-world-asset market as a live foil for r.

3 Method and the reproduce-first discipline

We follow three rules. *(1) Out-of-sample.* Every series is independent of the data the source papers used (e.g. the perpetual-funding panel is collected fresh from exchange indexers; the cross-domain panels are global rather than the original four-country subsets). *(2) Complementary statistics.* We deliberately compute quantities the papers did not: time-series rather than cross-sectional correlations, multi-venue rather than single-venue funding, global rather than regional cross-domain matrices. *(3) Reproduce-first.* Before treating any disagreement as a discrepancy we reproduce the

Layer	Providers	Content
Credit (USD)	Cbonds	921 sukuk+conventional bond G-spread histories; 51 sovereign CDS 40 sovereign USD T-spread/YTM indices; full sovereign universe m
Credit (local)	Cbonds	215 OIC local-ccy sovereign bonds (IDR/MYR/TRY/SAR); India/I Brazil/Mexico/China local yield curves
Perp funding	dYdX v4 (ι=0), Binance, Bybit (ι>0)	43 dYdX markets (hourly, 804k intervals); 40 Binance, 30 Bybit (8h); incl. PAXG/XAUT/XAG/WTI commodity perps
Risk-free curves	FRED, ECB, Cbonds	US Treasury (DGS, full curve); euro AAA + all-issuer (8 tenors); lo
Credit fundamentals	BIS, OECD, IMF	debt-service ratios, credit/GDP, REER, 49-country policy rates, 46-country long/short rates, debt/current-account/GDP
Cross-domain	World Bank, WHO, HMD, EM-DAT, FAOSTAT, NASA	cereal yield, death rate, adult mortality, 41-country life tables, 16,853 disaster records, crop yields, climate (211 countries)
Tokenized RWA	CoinGecko, DeFiLlama	PAXG/XAUT gold, BUIDL/USYC/USDY treasuries; 149 RWA pro
Microstructure	dYdX indexer, Binance	live order-book depth (BTC/ETH); dYdX-v4 TVL history (821 day

Table 1: The seventeen-provider out-of-sample dataset (∼1.1 GB; all keyless or user-licensed; every series reproducible from a per-provider puller).

source paper's own figure from its replication archive and verify the two statistics measure the same object. This rule matters: it dissolved three apparent "corrections" (Section 10).

On independence. One disclosure governs how this paper should be read: it is written by the corpus's own authors. "Independent" throughout modifies the *data* — seventeen providers the source papers did not use, statistics they did not compute — not the investigators. A validation run by the proposers, however disciplined, is not external replication, and we do not present it as such; it is the strongest test the authors can run on themselves, published with the archive that lets anyone else run a stronger one. External replication by parties with no stake in the programme remains the open invitation, and nothing in this paper substitutes for it.

On the number of analyses. A large analysis count invites the suspicion of multiple comparisons — that with enough tests, some confirm by chance. Three features of the design answer it. First, the analyses are overwhelmingly *confirmations of the corpus's own pre-existing claims*, not hypotheses fished from the data: each targets a specific quantity a source paper already asserts (a slope, a level, a ratio, a determinacy ceiling), so a "hit" is a reproduction out of sample, not a discovery, which is not a multiple-comparisons problem. Second, we report the failures: a null where the corpus's claim does not survive (the credit-to-mortality lead), an inconclusive test, and several figures we label construction-specific rather than confirmed — a protocol that reports its nulls is not selecting on significance. Third, the load-bearing conclusions rest not on any single test at the 5% margin but on directions confirmed by many independent instruments at once (venue-independence across three exchanges, term structure across twelve sovereigns, the keystone across forty-one markets). The reader who prefers a smaller paper may take Sections 4–7 and the scorecard as the result and the remainder as a robustness appendix; the substantive findings do not change.

4 κ is a real, observable, venue-independent intensity

Venue independence. If κ were an exchange artifact it would not agree across independent venues. Aggregating each market's funding to a daily $\hat{\kappa}$ and correlating across exchanges gives a mean cross-venue correlation of 0.70 (CEX∼CEX 0.85, DEX∼CEX 0.65; n=8 liquid markets). The same convergence signal appears on dYdX, Binance and Bybit; the lower DEX–CEX figure relative to CEX–CEX is itself the ι layer (Section 5).

It is the tradable market's credit signal. The aggregate sovereign $\hat{\kappa}$ (mean across 20 Cbonds T-spread indices) correlates -0.80 in levels with every tradable EM-credit ETF (EMHY -0.87, HYG -0.80, EMB -0.80, EMLC -0.80, VWOB -0.77): higher κ, wider spreads, lower bond prices, as a credit intensity must. κ is not a model construct; it moves with the market that prices the risk.

Non-negativity. The programme stresses $\kappa \geq 0$ by construction, contrasting it with conventional rates that went negative. Empirically the credit-spread carrier of κ is negative in 7 of 54,047 sovereign T-spread observations (0.01%, and only for Sri Lanka mid-default); sovereign CDS and the funding-magnitude $\hat{\kappa} = |\,\text{basis}\,|$ are never negative. By contrast the euro AAA risk-free reached -1.00% and stayed negative 1,644 days (Section 7). The differentiator holds with hard numbers; we recommend phrasing it as a property of the *hazard*, a rare negative G-spread being a technical floor-breach rather than a negative intensity.

κ predicts realized default. The sharpest test of the claim that $\hat{\kappa} = s/(1-\delta)$ is a forward *default hazard* is whether it anticipates actual defaults. Ranking the twenty sovereign T-spread indices by their *early* (2016–18) mean $\hat{\kappa}$, the three issuers that subsequently defaulted or restructured occupy the top three places: Ukraine (13.1%), Lebanon (13.0%) and Sri Lanka (10.0%), with the rest of the field below 9.7%. Lebanon, the highest pure-credit intensity, defaulted first (March 2020); Sri Lanka followed in 2022. The $\hat{\kappa}$ ladder front-loaded the realised defaulters out of sample — exactly the behaviour a hazard rate should exhibit. The sample is small (three events) and Ukraine's 2022 default is partly war-driven, but Lebanon and Sri Lanka are clean fiscal defaults that the intensity flagged years ahead.

The keystone, measured directly. The parameter the entire programme rests on is the AHJ convergence intensity itself — the rate at which a perpetual converges to spot, $\theta \sim \text{Exp}(\kappa)$. We measure it directly from the basis: fitting an AR(1) to each market's funding (premium) series and reading $\kappa = -\ln\phi$ per hour gives, across 41 dYdX markets, a median convergence intensity of $\sim$1,700/yr — a median basis half-life of 3.5 hours (from 12.7 hours on a thin commodity perp to 0.5 hours on the most liquid). The perp basis closes in hours; the keystone is not an assumption but a measured rate of order 10^3/yr, consistent with the magnitude argument of Section 7 and distinct from the funding-magnitude reading (this is the *speed* of convergence, not the size of the premium).

5 The interest parameter ι is a separable, removable layer

The CEX funding formula adds an interest component (the 0.01%/8h baseline); the $\iota = 0$ premium-only rule does not. The *signed* mean funding rate isolates it cleanly. Across nine liquid markets the dYdX ($\iota = 0$) mean drift is $+0.10\,\text{bp/interval}$; the CEX ($\iota > 0$) mean drift is $+1.01\,\text{bp/interval}$ — precisely the $1\,\text{bp/8h}$ interest baseline (Table 2). This is a sharper statistic than the funding-*magnitude* comparison, which is asset-class dependent (clear positive wedge for mid-cap alts, near zero for BTC/ETH and for metals). The interpretation is exactly the Separation Theorem: ι is an additive layer of size equal to the interest baseline, removable without disturbing the convergence signal κ.

	dYdX (ι=0)	Binance (ι>0)	Bybit (ι>0)
BTC	+0.13	+1.07	+1.20
ETH	+0.15	+1.29	+1.32
XRP	+0.08	+1.38	+0.96
mean (9 mkts)	**+0.10**	**+1.01**	

Table 2: Signed-mean funding drift (bp per interval). The CEX drift equals the 1 bp/8h interest baseline; the $\iota = 0$ venue is ∼zero. The wedge is the riba.

Distributional distinctness (Paper 1). The same separation shows in the full funding *distribution*, not just its mean. Recomputing Paper 1's distributional test on the larger multi-venue panel, the CEX-versus-dYdX funding distributions differ with Cohen's $d = 0.64$ (BTC 0.89, ETH 0.84, XRP 0.73; Kolmogorov–Smirnov $D \approx 0.6$), reproducing the paper's reported $d = 0.782$ to within sampling and venue differences — a large effect size, with the CEX distribution shifted up by the interest component (SOL the lone exception, owing to its negative-mean Binance funding).

The CEX floor is a literal interest pin. The mechanism is even starker in the raw distribution. On Binance ($\iota > 0$) the *median* funding rate across BTC, ETH, XRP and SOL is *exactly* 0.0001 — the 0.01%/8h interest baseline written into the formula — and 44–48% of all observations sit within a hair of that floor; only 14–24% are negative. On dYdX ($\iota = 0$) the median is ≈ 0, just 3–5% of observations cluster near any floor, and 24–32% are negative — a genuinely two-sided premium. The contrast is the riba made visible: the CEX rate is pinned to a predetermined interest constant, the $\iota = 0$ rate is centred on zero.

The transfer nets to zero at $\iota = 0$. What ultimately matters is not the funding rate but the wealth *transfer* it produces between the two sides — the system-level riba test of Paper 3. A Monte-Carlo of a balanced long–short market (2,000 seeds, ±75 bp clamp) makes the point: at $\iota = 0$ the cumulative net transfer from longs to shorts over a year is statistically zero (mean 0.86% of notional, a confidence interval straddling zero, the small residual being a known Itô discretisation drift), whereas with $\iota > 0$ a systematic transfer accrues that grows linearly with ι — 5.4% of notional at the $\iota = 0.01\%$/8h CEX floor (about 11%/yr), rising to 46% at ten times that rate. The interest leg is exactly the riba: present when $\iota > 0$, gone when $\iota = 0$. (We do not reproduce the paper's interior Nash leverage equilibrium; our simplified leverage game collapses to a corner, and we defer to Paper 3's full bimatrix solution for that figure.)

6 κ is credit-ordered and term-structured

Levels. With $\delta = 0.6$, $\hat{\kappa} = \text{OAS}/(1 - \delta)$ from FRED ICE BofA indices reproduces the corpus's credit levels exactly: US high-yield mean $\hat{\kappa} = 0.081$ (claimed 0.081), EM high-yield 0.102 (claimed 0.104), EM corporate 0.048, US investment grade 0.023 — a clean ordering.

Term structure. Multi-tenor CDS (6M → 3Y) gives an upward-sloping κ-curve for *all twelve* sovereigns tested (mean +0.39/yr; Turkey +1.08, Egypt +0.90), independently confirming Paper 2a's slope claim that the single-date sukuk snapshot could only support weakly. The result also survives Paper 2a's own panel specification on a larger, fresher sample: on a 179,618 bond-day USD-sukuk panel (96 issuers, versus the original 273-bond pull), the issuer-plus-month fixed-effects slope is +2.85 bp/yr for investment-grade sukuk (+2.86 for the full sample), and — adding the inference the original table omitted — it is strongly significant under issuer-clustered standard errors ($t = 4.1$ and $t = 4.7$).

Rating ordering. The cross-section is monotone in credit quality. Bucketing 691 rated bonds by rating and taking the median $\hat{\kappa}$ within each, the ordering rises AAA 1.23, AA 1.78, A 2.38, BBB 2.85, BB 8.32, B 9.23 (%/yr) — monotone increasing with the sharp, realistic jump at the investment-grade/high-yield boundary. κ behaves as a credit intensity must. The ordering extends to a full *surface*: tabulating median $\hat{\kappa}$ across rating $\times$ tenor from multi-tenor CDS, every cell rises in both directions (AA 0.30 $\to$ 0.59% over 6M $\to$ 3Y, up to B 7.28 $\to$ 12.51%), monotone in tenor for all five rating buckets and monotone in rating at all four tenors — the κ-surface is the two-dimensional generalisation of the curve. It is also internally arbitrage-free: the forward intensity implied between every adjacent tenor pair, $f(t_1, t_2) = [\kappa_2 t_2 - \kappa_1 t_1]/(t_2 - t_1)$, is non-negative in all 15 pairs across the surface — the dynamic analogue of the Feller condition, so the surface admits a valid no-arbitrage hazard-rate term structure rather than merely sloping up.

Why the curve slopes up: a term premium, not expectations. The upward slope could in principle be pure expectations (the market expecting κ to rise). It is not. For each of the twelve sovereigns with both 6-month and 3-year CDS, comparing today's κ(3Y) to the κ(6M) *actually realised* over the following 30 months gives a positive gap in all twelve cases (pooled mean +1.65 percentage points of κ; Egypt +5.0, Turkey +3.0, Korea +0.2). The long end sits systematically above the short rate that subsequently materialises, so the slope is a genuine term/risk premium — investors are compensated for holding convergence-intensity duration, exactly as a credit term structure should behave. The upward shape is not USD-specific: in the Chinese local market the AA corporate-over-government κ rises from 1.03% at 1 year to 2.11% at 10 (the AAA curve, like high-grade curves everywhere, is nearly flat), so the κ-term-structure replicates on conventional non-USD bonds as well.

The κ-curve in local currency. The strongest non-USD test uses local-currency sukuk directly. In Malaysia, where both legs are observable, the local credit intensity — $\hat{\kappa}$ = (MYR sukuk yield − MGS government yield)$/(1 - \delta)$ — is positive and upward-sloping by maturity: 0.08% at 1–4 years rising to 1.06% beyond 12 (slope +0.065%/yr). The κ-yield curve is therefore not a dollar artifact; it exists in ringgit, at a level consistent with Malaysia's USD sukuk κ ($\sim 0.9\%$). The Islamic wrapper itself adds nothing in local currency: the Malaysian Islamic government curve (MGII) and the conventional one (MGS) differ by a mean of 3.5 bp across tenors, the same wrapper-independence the USD matched panel shows. (The local nominal *level* is, as expected, inflation-dominated — Turkish lira sukuk yield 33–47% on an inverted curve, Malaysian and Saudi 3–5% — so a local yield is $r + \kappa$ with κ the small credit part.)

Dynamics. All twenty sovereign $\hat{\kappa}$ series mean-revert (AR(1) $\rho = 0.77$–0.98); the T-spread-*index* half-life (median 7.1 months) runs longer than Paper 2a's per-issuer constant-maturity 1.4–4.3 months because index aggregation smooths — consistent in direction. $\hat{\kappa}$ scales by $1/(1 - \delta)$, so the level depends on the recovery assumption, but the cross-sectional ordering is δ-invariant. A principal-components decomposition of 1,524 four-tenor κ-curves returns the textbook term-structure factor structure: a dominant level factor (98.7% of variance, all-tenor loadings of one sign), then slope (1.2%, opposite-signed ends) and curvature (0.1%) — κ is a genuine term-structure object with the same Litterman–Scheinkman decomposition as a yield curve, not noise. The non-negativity the corpus builds in via the CIR Feller condition $2\alpha\bar{\kappa} \geq \nu^2$ also holds empirically: fitting the condition to each sovereign $\hat{\kappa}$ series, it is satisfied for 18 of 20 sovereigns, failing only for Ukraine and Sri Lanka — the two defaulted issuers, where κ is large and volatile. For every performing sovereign $\kappa \geq 0$ is a measured property, not just an assumption. That recovery assumption is itself empirically grounded, not arbitrary: the five USD sukuk with realised post-default recoveries cluster at a median of 10% for asset-based structures (Serba Dinamik 5%, Swiber 10%, NMC 20%), validating the corpus's $\delta = 0.10$ corporate convention, while the one sovereign-linked default (Bank Asya, a Turkish-government backstop) recovered near par at 95%, consistent with the much higher sovereign δ. The two recovery conventions bracket the observed defaults.

iCDS pricing and the riba premium (Paper 2c). Reproducing Paper 2c's present-value riba premium $\Pi^{\text{PV}} = (1-\delta)\big[(1-e^{-\kappa T}) - \frac{\kappa}{\kappa+\iota}(1-e^{-(\kappa+\iota)T})\big]$ ($T = 7$, $\iota = 4.24\%$) over the empirical sovereign range $\kappa \in [1\%, 10\%]$ yields $\Pi^{\text{PV}} = 0.36$–2.41% of notional (maximum 2.6% over $[0.5\%, 12\%]$) — confirming the paper's "under 4%" claim. The fair iCDS spread $s^* = \kappa(1-\delta)$ contains no ι term, so it is invariant to the interest rate: doubling ι ($4.24\% \to 8.48\%$) leaves s^* at 160 bp for $\kappa = 4\%$, whereas a conventional spread embeds the funding leg. The riba-free spread is, by construction, the part of credit protection that does not depend on r.

Robustness to recovery and regime. Because $\hat{\kappa} = s/(1-\delta)$, the recovery assumption δ scales every level by $1/(1-\delta)$ but leaves orderings unchanged: across $\delta \in \{0.4, 0.6, 0.8\}$ the rating-bucket levels move by a factor of three while the AAA-to-B ordering is identical (rank correlation 1.0), and the determinacy ceiling of Section 9 rises with δ but preserves its credit ordering. The directional findings are also regime-stable: splitting at the COVID onset (2020-03), the rating ordering is monotone in both sub-periods and the multi-tenor CDS κ-curve slopes upward for all thirteen sovereigns *both* before and after — the slope claim does not depend on the epoch. Only levels move with δ and regime; signs and orderings do not.

7 κ is not r — and the relationship is a decomposition

Four independent measurements separate κ from r without claiming they are unrelated. *(1) Magnitude.* The perpetual convergence intensity (from the hourly funding half-life, median ~3.5 h) is ~1,700/yr at the median — the same figure as the keystone measurement of §4 — with the most liquid markets running to ~10^4/yr; credit κ is 0.01–0.10/yr; mortality $\kappa \approx 0.002$–0.01/yr — a range of *six orders of magnitude.* No interest rate is simultaneously 0.2% and 170,000% per year. *(2) Policy rate.* Across twelve sovereigns the ratio r/κ of the BIS policy rate to the USD credit $\hat{\kappa}$ ranges 1.6× to 6.6× (Turkey: policy 37% vs credit $\hat{\kappa} = 5.6\%$); the two are cross-sectionally correlated (+0.77, both rising with country risk) but the four-fold ratio spread and the inflation-driven Turkish outlier show r carries a monetary/time component κ lacks. *(3) Inflation.* $\hat{\kappa}$ correlates only +0.41 with inflation; Turkey's 61.6% inflation sits beside a 5.6% credit $\hat{\kappa}$. *(4) Negative rates.* The euro AAA risk-free was negative for 1,644 days at the 5-year tenor while credit priced throughout — the Separation Theorem's empirical corroboration, quantified. *(5) The tokenized foil.* On-chain tokenized treasuries pay r directly (BUIDL 3.5%, USYC 3.15%, USDY 3.55%, against a 3.83% 3-month T-bill), while the tokenized-gold perpetual pays $\kappa \approx 0.18$ with no interest — two on-chain assets, one paying r and one paying κ, in the same market. The defensible statement is not "κ and r are unrelated" but "$r = \kappa +$ (separable time/inflation component)": they share a credit-risk core and differ by the layer the programme says is removable. That separable component is, on the data, inflation: across twelve emerging-market central banks the gap between the policy rate and the credit $\hat{\kappa}$ correlates with national inflation at $\rho = 0.96$ ($p < 0.001$), with a slope of 0.46 — so what r carries that κ does not is, to first order, the inflation/monetary layer. (The correlation is anchored by Turkey's 62% inflation; excluding it the relation is positive but noisier, and the slope below one reflects incomplete real-rate pass-through. The direction is the point: the Separation Theorem's "removable layer" is empirically the inflation wedge.)

The benchmark is a decomposition, not a dependency. A standing objection is that the spread-based $\hat{\kappa} = s/(1-\delta)$ is "contaminated" because the spread s is quoted against a Treasury curve. Paper 2a's benchmark-free estimator removes the curve entirely: $\hat{\kappa}_{\text{price}} = -\frac{1}{T}\ln\big((e^{-yT} - \delta)/(1-\delta)\big)$ uses only the bond's own yield y and maturity T. Computing it on 860 bonds and ranking against the spread-based $\hat{\kappa}$ gives a Spearman correlation of 0.84 (full panel), 0.89 (sukuk) and 0.82 (conventional) at the measured recovery $\delta = 0.10$ — κ orders bonds identically with or without a benchmark. Crucially, the *only* difference is an additive level shift, and that shift is near-exactly the reserve currency's own yield mapped through the same transformation: the median wedge is 473 bp of κ versus the predicted $y_{\text{UST}}/(1-\delta) = 472$ bp. Moreover the wedge tracks the Treasury yield *through time*: year by year (2015–26) the median wedge and the 10-year UST yield correlate at $+0.93$ ($p < 0.001$), moving together from 120 bp of κ at the 0.89% yield of 2020 to 500 bp at the 4.2% yield of 2024. One interpretive discipline is required here, and the programme's own decomposition supplies it: the US realised default intensity is ≈ 0 on the record, and this section separately finds the $r - \kappa$ gap to be the inflation/monetary layer ($\rho = 0.96$ below) — so the wedge is the reserve currency's *time-value and premium layer made visible*, not a measured default hazard of the United States. Calling it "the reserve sovereign's own κ" would relabel the removable layer rather than remove it, and we do not. What the benchmark-free estimator establishes is therefore precisely scoped: all *relative* pricing — orderings, curve shapes, spreads between issuers — is recoverable with the Treasury curve deleted, because the wedge is a common additive shift; the Treasury curve is a *time-varying* decomposition choice for isolating issuer excess over the dollar system's base yield layer, not an operational dependency. That is the defensible form of the corpus's "not contaminated" claim, confirmed on three times the bonds of the original and dynamically.

The commodity-κ level is a regime average. The programme's headline gold intensity $\kappa_{\text{Au}} \approx 0.17$ holds on the deepest venues as a *window average* (dYdX PAXG 0.183, XAUT 0.169, Binance XAU 0.174; silver dYdX XAG 0.264, Binance 0.234 against the 0.25 headline), but two qualifications are warranted: it is venue-dependent (Bybit gold runs near 0.06), and it drifted down 40–60% within the 2024–26 sample (first-half ≈ 0.22 to second-half ≈ 0.11–0.14). The figure is best read as a regime average rather than a fixed constant; the pre-2024 macro episodes are not testable because the commodity perpetuals post-date them.

The formula consequence: interest-rate risk and its removal (the 2022 rate shock, read through the model). That κ is not r has a consequence that 2022 makes vivid — provided it is stated for what it is. First, an empirical fact: the two objects are not merely distinct but *oppositely* signed over the cycle: the policy rate (effective federal funds, FRED `DFF`) is held unchanged in 47% of months and then moved in discrete committee steps (0.08% in 2021 to 5.02% in 2023 *by decree*), while κ moves continuously (flat in only 4% of months); their correlation over 2015–2026 is -0.76, because rate cuts coincide with the risk spikes that raise κ. The rate is a *choice*; κ is a *measurement*. Second, a model property: the Separation Theorem (Paper 17) says a κ-priced claim at $\iota = 0$ carries *no* discount factor, so $\partial\text{price}/\partial r = 0$ *within the pricing formula*: the direct rate channel that revalues conventional bonds is absent by construction. The 2022 rate shock lets us quantify what that channel cost. Repricing par bonds at the realised 2021-low-to-2022-high Treasury yields (full revaluation, including convexity, Table 3) gives a steep *duration gradient* for conventional, interest-priced bonds — from −4.5% at the one-year to −45.2% at the thirty-year — whereas the $\iota = 0$ formula, applied to the same shock, has no rate term to revalue. What the κ instrument retains is a far shallower *event* gradient (price $\propto e^{-\kappa(1-\delta)T}$): applying the realised 2021–22 credit-κ move (+0.003/yr) gives losses of −0.1% to −3.8% across the curve — between 12 and 35 times smaller than the conventional rate loss at every maturity, because the +2.7 percentage-point rate move dwarfs the tiny κ move. Two scope disciplines keep this honest. *This is a model computation, not an observation*: no κ-priced instrument traded through 2022, so the "0.0%" column is the formula's property, not a realised return. And *in a mixed economy the equilibrium channel remains*: a traded κ-claim competes with r-bearing assets, and the market-clearing level of κ_Q itself responds to the rate environment — the benchmark-free wedge above tracks the UST yield at $\rho = 0.93$ — so holders of a real κ-instrument in 2022 would have borne a repricing through κ_Q, not none at all. What the table isolates is the *direct* rate-revaluation channel, which the $\iota = 0$ design deletes; the indirect channel is an equilibrium fact about any asset in an r-world, and only a closed κ-world removes it.

Tenor	Δr (2021→2022)	mod. dur.	**conventional**	**κ-priced (rate)**	κ-priced (event)
1y	+4.76	1.0	−4.5%	0.0%	−0.1%
2y	+4.63	2.0	−8.6%	0.0%	−0.3%
5y	+4.09	5.0	−18.0%	0.0%	−0.6%
10y	+3.32	9.5	−26.6%	0.0%	−1.3%
20y	+3.13	17.2	−40.4%	0.0%	−2.5%
30y	+2.74	23.5	**−45.2**%	**0.0**%	−3.8%

Table 3: The 2022 rate shock across the curve. Conventional interest-priced bonds lose value in proportion to duration (a single rate decision revalues the long end by −45%); the $\iota = 0$ pricing formula has no rate term at any tenor (Separation Theorem), retaining only a far shallower exposure to the realised hazard κ. Real FRED yields; par-bond full revaluation. The two κ-priced columns are *computed from the $\iota = 0$ formula*, not observed on a traded instrument — they isolate the direct rate channel and do not include the equilibrium repricing of κ_Q that a κ-claim traded in a mixed economy would bear (see text).

What this shows — and what it does not. Interest pricing bakes a duration gradient — pure interest-*rate* risk, the price of *time* moving — into every fixed-income market; in 2022 that gradient cost a thirty-year bond nearly half its value on a policy decision, not on any real event. Pricing at $\iota = 0$ deletes that channel from the formula and keeps only the shallow exposure to the real hazard the investor actually wishes to bear. This is the Separation Theorem read across the curve against the cleanest available rate shock: *remove the price of time, keep the price of the event.* What it does *not* show is that a κ-instrument holder would have sailed through 2022 unmarked — that claim would require a traded instrument and a closed κ-world, and we have neither; it shows what the direct rate channel cost and that the $\iota = 0$ design carries no such channel. (Further scope: κ removes rate risk, not all risk — the event column is real; par-bond losses depend on a portfolio's duration mix; and the long-horizon κ dynamics remain the programme's deepest open question.)

8 Cross-domain diversification, measured

We build the global cross-domain shock matrix the programme flagged as its standing empirical task, across six domains (agriculture, mortality $\times 2$ sources, disaster, climate, macro) at 150–260 countries. The mean off-diagonal $|\rho|$ is 0.05 over 58,665 country-years — vindicating the small-ρ_x assumption on far more data than the original four-country panels, with sensible physical signs (climate$\times$agriculture $+0.10$, GDP$\times$mortality -0.08). The within-domain correlations are high (credit 0.50, mortality 0.47) — "domains, not names." Feeding the *measured* blocks ($\bar{\rho}_w = 0.43$, $\rho_x = 0.07$) into Paper 14's four-sleeve floor gives $\approx 1.6\%$, recovery near the bottom ($\sim$80%) of the paper's 80–98.5% band rather than the top — which reproduces Paper 14's own honest accounting (its 98.5% is the labelled $\rho_x \leq 0.02$ best case).

The new finding: tail dependence. Average correlation can understate *crisis* co-movement. In the 2020–21 pandemic the aggregate sovereign credit intensity rose +167% and the cross-country death rate +14% — together — while disaster impact was flat (−5%). Credit and mortality are tail-coupled through the real pandemic channel; disaster and agriculture are not. The κ-ladder therefore diversifies in normal times but less in a systemic tail, and along channel-specific lines: a tail-robust ladder should over-weight domains with disjoint real-event channels. This refines, rather than overturns, Paper 14's structural argument and is now folded into its §"What is ρ_x?" (the result is new to the corpus). A second episode confirms the coupling is channel-specific rather than a general crisis feature: in the 2022 rate-shock-and-war year the global credit intensity rose only +18% (against COVID's +167%) while the death rate *fell* (−8%, as the pandemic receded) and cereal yields were flat. A purely financial crisis did not co-move the real domains at all — so the credit$\leftrightarrow$mortality coupling needed the shared real channel of the virus, and the κ-ladder's diversification holds in financial crises, weakening only in shared real-event tails. This strengthens rather than qualifies the safe-asset argument.

Direction of the coupling: a null. A lead–lag check on the eight-point *global* annual series appeared to show credit-κ leading mortality by a year, but this does not survive a properly-powered test. Pooling across twenty countries (120–200 country-year change-points), the "credit leads mortality" lag is insignificant ($\rho = +0.05$, $p = 0.58$ at $k{=}{+1}$); the only marginally significant cell is the opposite direction at a two-year lag ($\rho = 0.16$, $p = 0.02$). We therefore record the credit$\leftrightarrow$mortality coupling as weak and *not* clearly credit-leading — the global hint was a small-sample artifact. The contemporaneous COVID co-spike (above) stands; its temporal direction does not.

The κ-index as a systemic-stress gauge. The same cross-sectional median that drives the ladder argument is what the liveness mechanism (Paper 12) watches for macro winters. Rebuilding the κ-index from our 921-bond panel (121 issuers, 138 months) reproduces Paper 12's stress calibration on independent data: the 99th percentile is 10.1% (paper 9.5%), the index is breached in 2 of 138 months — a 1.4% gate frequency against the paper's 1.6% — and those two months are exactly the COVID peak (2020-04/05), the same episode the paper identifies. The persistence ($\rho = 0.94$, half-life 10.5 months) matches as well. The sleeve's conjunction trigger is therefore not an artifact of the paper's particular panel.

Why diversification needs domains, not names. A principal-components decomposition of monthly $\Delta\hat{\kappa}$ across the twenty sovereigns (124 months) puts 63% of the variance on the first component — a single common credit factor — leaving only 37% idiosyncratic. Sovereign credit κ is therefore *majority systematic*, so diversifying *within* the credit sleeve buys little (consistent with the measured within-credit correlation of 0.50). That systematic factor is, however, only about half a global-risk phenomenon: regressing the cross-sectional mean $\Delta\hat{\kappa}$ on contemporaneous changes in the VIX and the US high-yield OAS gives a joint R^2 of 0.44 (each ~ 0.4 alone), so roughly 44% of the systematic credit-κ factor is global risk-off and the remaining $\sim 56\%$ is emerging-market-credit-specific. Either way the lesson for Paper 14's ladder is the same: it diversifies across *domains* rather than names, because the gains come from the small cross-domain ρ_x, not from holding more sovereigns whose common factor — global or credit-specific — does not wash out. That common factor is, moreover, *state-dependent*: the rolling 24-month correlation of the κ-index with the VIX rises from +0.39 in calm years (2017–19) to +0.79 in stress (2020–21), peaking at +0.89 at the April-2020 nadir. κ's systematic loading is largest exactly when a buffer is most needed — the within-credit diversification weakens in the tail, which is precisely why the cross-*domain* ladder, not more names, is the right construction.

κ as the force of mortality. The programme treats mortality as another κ-domain, with $\kappa_{\text{mortality}}(\text{age }35) \approx 0.002$. The Human Mortality Database confirms the structure directly: across 41 national life tables the central death rate rises from 0.0007 at age 35 to 0.080 at age 85, and $\ln \mu_{\text{age}}$ is linear in age with slope 0.094/yr — the death intensity doubles every 7.4 years. This is the Gompertz law, a clean exponential hazard term-structure, which is exactly what a κ should be: the mortality domain furnishes an independent, non-financial confirmation that κ behaves as a loss intensity with a well-defined term structure. (The level sits below the corpus's 0.002 because the HMD panel is developed, low-mortality countries; the order of magnitude and the structure are the point.)

κ as the disaster arrival rate. The disaster sleeve is the same: across 194 countries with at least five EM-DAT events (2000–24) the disaster arrival rate is a country- and peril-specific Poisson intensity, median 1.68/yr, with global per-peril rates of 0.74 (flood), 0.48 (storm) and 0.16 (epidemic). Disasters, like defaults and deaths, are a hazard with a measurable intensity — the three real-economy sleeves of Paper 14's ladder (credit, mortality, disaster) are each a κ-domain with its own arrival rate, which is exactly the structure the safe-asset argument assumes.

Takaful premiums equal expected loss. The actuarial-fairness claim — that at $\iota = 0$ the fair premium is $\pi^* = \mathbb{E}[\text{loss}]$ — is testable on a real, functioning market. Across the 139 commodity lines of the USDA Risk Management Agency crop-insurance book ($1.25 trn of liability), the aggregate premium rate is 9.6% against an expected loss-cost (indemnity over liability) of 7.9% — a loss ratio of 0.82 — and premium rates track loss-costs across commodities at $\rho = 0.69$. Premiums equal expected loss plus a modest loading, exactly the $\iota = 0$ fairness the takaful paper derives, here on a live insurance market rather than a model.

More sleeves keep ρ_x small — but only real ones. Extending the cross-domain matrix from four sleeves to seven (adding climate, the real effective exchange rate, and the debt-service ratio), the mean off-diagonal $|\rho|$ is 0.072. The genuinely real domains stay near zero against credit (agriculture +0.01, disaster +0.00, climate −0.04, mortality −0.14), so the ladder scales with more real sleeves. The exception is instructive: the real exchange rate couples to credit at −0.49 — a credit deterioration depreciates the currency — because it is a *financial*, not a real-event, variable. The lesson sharpens Paper 14's: diversify across domains with disjoint *real* channels; financial aggregates are not independent sleeves.

The measurement-integrity monitor reproduces. Paper 13's defense of the published κ against tampering has two parts; the oracle depth rule needs live order-book depth and is out of our reach, but its cross-sectional CUSUM monitor is testable on our panel. Reproducing it on 66 issuers (versus the paper's 37) — fitting each issuer's κ to the cross-sectional median, then running a Page CUSUM on the level residuals — an injected warp is detected with a median lag of 11 months at 10 bp, 8 at 25 bp and 5.5 at 50 bp (the paper reports 11.5, 9.5 and 5.0), catching 92% of issuers. Targeted κ-warps are caught in months, bigger ones faster — the integrity mechanism holds on independent data.

The depth rule and the manipulation frontier, on live order books. Paper 13's other defense — the depth rule, capping repriced notional so a trade cannot move the mark past the validity band — needs order-book depth, which we pulled directly (dYdX and Binance, BTC and ETH). On the deep venues a trade sized at 7% of one-sided book depth moves the mark only 3–13 bp, inside the ±75 bp clamp, and the dollar cost to push the mark to the 25 bp detection threshold is $1.3 M on dYdX rising to $26 M on Binance. A full Monte-Carlo of the frequent-batch-auction mark against a continuous last-trade mark, on books calibrated to that live snapshot, confirms the FBA mark is orders of magnitude costlier to warp (a multiple of ∼1,700× at 10 bp rising past 10,000× at 200 bp, 95% CIs excluding 1), reproducing the programme's 72–2,739× claim in direction and order. Layering the 30-block TWAP, the clamp and the CUSUM monitor on top, a sustained "grind" warp is detected within 0–12 blocks and its net extraction is zero or negative once the cost to hold the position is counted: manipulation does not pay. The measurement-integrity stack, previously argued, is now validated on both live order-book data and simulation.

The liveness trigger does not false-fire. The validator-liveness sleeve (Paper 12) arms only on a *conjunction* — low transaction volume *and* macro stress — precisely so it cannot be triggered by an ordinary downturn. Pulling dYdX's TVL history (a $390 M→ $96 M drawdown, a genuine crypto volume winter) and crossing it with our credit κ-stress index, the two legs never coincide: the volume winter is 2025–26, the κ stress is the 2020 pandemic spike, and the conjunction fires in zero months. The trigger correctly stays shut during a crypto-only winter — the false-positive avoidance the design intends, confirmed on data.

The system-level simulations reproduce on their own engines. Finally, the two purely simulated claims — protocol solvency and censorship-resistance — reproduce when their original code is re-run. Executing Paper 3's cadCAD market (version 0.5.3) over thirty Monte-Carlo runs of 720 steps each gives zero insolvency events in 30/30 runs, an insurance fund that grows from its $100 k seed to $100.9 k (self-sustaining), and a funding rate that never leaves the ±75 bp clamp — the solvency claim exactly. Re-training Paper 12's censor-cartel with 150,000 PPO steps, the learned censorship policy earns −0.39 less than honest behaviour and beats honesty in 0 of 30 matched episodes, with suppression converging to 1.4%: censorship is strictly dominated and the liveness sleeve is censorship-proof, as claimed. These are reproductions rather than independent tests — they re-run the authors' own simulators — but together with the data-grounded manipulation and conjunction results above they close the last open claims with live code.

9 The monetary corpus meets the data (Papers 8–9)

The monetary half of the programme — the κ-Standard (Paper 8) and the interest-free DSGE (Paper 9) — ships replication code but, unlike the pricing papers, had received no independent out-of-sample test. Its quantitative claims are nonetheless measurable, and we test four.

The determinacy ceiling, on twice the sovereigns. Paper 9's price-level determinacy condition under active-fiscal closure is $\gamma_s < \gamma_s^* = e^{\bar{\kappa}} - 1$, where $\bar{\kappa}$ is the steady-state credit intensity. Reproducing its `empirical_determinacy_map.py` returns the published twelve-sovereign ceilings exactly (Qatar 0.0072 → Pakistan 0.0955, ordered by credit quality). We extend it to *twenty-six* distinct sovereigns across two independent instruments: the twenty Cbonds sovereign USD T-spread indices and the twelve sovereign CDS curves (3-year). On both legs the ceiling is small but strictly positive and monotone in credit quality (CDS leg 0.006 Korea → 0.139 Egypt, bracketing the paper's range; the index leg runs 1.3–2.6× wider for the same sovereign, the systematic wedge of an index spread over a 3-year CDS). The determinacy window is narrow but never empty for any performing sovereign — the first independent confirmation of Paper 9's central result, on an instrument the paper did not use.

Sovereigns and corporates. Extending the map to corporates (Paper 9's `sovereign_vs_corporate`) on our 107-issuer panel (versus the original 34), the corporate determinacy ceiling sits above the sovereign in five of seven countries with both, and corporate credit is *more persistent* (AR(1) $\rho = 0.865$ versus the sovereign 0.77, half-life 4.8 months), so the buffer requirement $S^* \propto \sigma^{2.3}(1-\rho)^{-2.2}$ is larger (∼25% versus ∼12%) — the paper's sovereign<corporate ordering, on three times the issuers.

Money-base auto-contraction. The κ-Standard prices base money as $M = \sum_i N_i e^{-\kappa_i(1-\delta)T_i}$, which contracts when intensities rise. Reproducing Paper 8's `full_universe_base.py` gives −1.99% per +100 bp on its 195-bond panel; on our 860-bond panel the sukuk sub-panel contracts −2.03% (an essentially exact match) and the full panel −3.08%, the elasticity scaling as $\approx -(1-\delta)\bar{T}$ with the mean maturity $\bar{T}$ — the countercyclical base mechanism, confirmed independently. The base is moreover sensitive to the curve's *shape*, not only its level: a steepening (slope) shock moves the base less than a parallel one (−2.3% vs −3.1% per 100 bp) but concentrates the impact in long maturities (the $\geq$ 10-year bucket falls 11.8% while short bonds rise), so the κ-Standard carries a term-structure exposure that a single $\bar{\kappa}$ would miss.

Half-lives, systematic share, and the buffer, as ladders. Two reconciliations emerge, each the monetary analogue of the κ–CDS ladder of Section 10. First, κ's mean-reversion half-life is construction-specific: per-issuer constant-maturity 2.7 months (Paper 9), per-sovereign index 6.1 months, cross-sectional median index 3.0 months, aggregate index 7.1–8.4 months (Paper 2a) — to be quoted with its aggregation level, not as a single number. Second, the systematic share of κ reconciles cleanly: a principal-components decomposition puts the first component at 63% on all twenty sovereigns (matching Section 8), rising to 72% on the seventeen performing sovereigns and 75% on the ten investment-grade core — so Paper 9's "80–100% systematic" (an index/core statistic) and our 63% (a broad cross-section) are the same quantity at different breadth. The buffer estimate inherits this sensitivity: because S^* depends on $(1-\rho)^{-2.2}$, it ranges from ∼12% (per-issuer persistence) to ∼45% (index persistence), so the programme's "11–13%" figure is itself method-specific and should be reported with the persistence measure it assumes.

The real-block feedback exists. Paper 9's determinacy also turns on the procyclical credit channel $\hat{\kappa} = \psi(-\hat{y})$, whose threshold it places near $\psi^* \approx 0.62$. Pooling 159 country-years of World Bank GDP growth against the next change in credit-κ, the feedback is strongly procyclical and significant ($\rho = -0.50$, $p < 0.001$; κ rises as output falls). Its raw magnitude is unit-sensitive (the paper's $\hat{y}$ is an output gap, not a growth rate), so we read the result as confirming the *sign and strength* of the channel rather than a precise ψ — and the strength is itself the point: the feedback is large enough that the paper's stabilisation devices (net-worth predetermination, Tobin's Q, a sovereign stabilisation fund), which it shows restore determinacy, are empirically relevant rather than merely formal.

10 Reconciliation: three apparent discrepancies that are not

The reproduce-first rule (Section 3) converted three would-be corrections into benign differences of construction; each source number reproduces exactly.

- **κ–CDS co-movement is a ladder of four reproducible statistics.** The single most instructive case. Paper 2c's headline 0.97 is the *cross-sectional* correlation of median sukuk G-spread against sovereign CDS across five sukuk sovereigns; we reproduced 0.972 from its `real_validation.py`. Paper 13's per-country panel of 0.989–0.998 uses CIR-smoothed constant-maturity κ. Our two independent figures fall between and below: a sovereign T-spread *index* against CDS gives time-series 0.79, and the *raw* (unsmoothed) sukuk-implied $\hat{\kappa}$ against CDS gives 0.64 (Indonesia 0.62, Malaysia 0.66; only two sovereigns clear the 24-month overlap, so this figure is thin). The four numbers — raw 0.64, index 0.79, cross-section 0.97, CIR panel 0.989 — do not contradict; each measures a different object (raw vs. smoothed, within-country vs. across-country), and all reproduce. This is the cleanest illustration of the rule: quote the κ–CDS correlation with its construction attached. The relationship also holds *dynamically*: for the three sovereigns with both legs over 24 months, the sukuk-implied and CDS-implied κ co-move at a mean level correlation of 0.71 and their basis mean-reverts (AR(1) half-life $\approx$ 10 months), so the iCDS par-spread identity $s^* = \kappa(1 - \delta)$ is arbitrage-consistent in time and not only in the cross-sectional snapshot.
- **Takaful $\hat{\kappa} = 12.06\%$ identical for India and Pakistan.** Paper 2b uses $\hat{\kappa} = -\ln(1 - n_{\text{loss}}/T)$; both countries recorded 5 loss-years in 44, so the value is identical *by the data*, not by estimator pathology, and the paper already labels it a coincidence (and shows Pakistan's premium model misses by 1.6–3.4×). Our continuous-threshold estimator answers a different question and is not a correction.
- **Spanning recovery** 98.5% **vs. our $\sim$80%.** Paper 14 reports a *grid* over $\rho_x \in [0, 0.20]$ and discloses the measured-block floor; 98.5% is its labelled best case. Our measured-ρ_x figure reproduces its own honest accounting.

The lesson generalises: several κ magnitudes are construction-specific (cross-section vs. time-series; count-MLE vs. threshold; assumed vs. measured ρ_x) and must be quoted with their method attached. Doing so leaves the corpus's claims intact.

11 Scorecard

Claim tested	Verdict	Out-of-sample evidence
κ credit levels = OAS/$(1 - \delta)$	Confirm	US-HY 0.081, EM-HY 0.102 (exact)
Sukuk κ wrapper-independent	Confirm	matched panel $t = -1.23$ (n.s.)

ι a removable interest layer	Confirm	signed drift +1 bp/8h CEX vs ∼0 dYdX
κ mean-reverts	Confirm	all 20 sovereigns $\rho < 1$
κ-curve slopes up	Confirm	12/12 sovereigns, multi-tenor CDS
κ non-negative	Confirm	0.01% negative; risk-free −1% for 1,644 d
κ venue-independent	Confirm	cross-venue $\rho = 0.70$
κ tracks tradable credit	Confirm	vs EM/HY ETFs −0.80
κ spans many orders of magnitude	Confirm	0.002 → 1,700/yr median (6 OOM)
$\kappa \neq r$	Confirm (decompose)	r/κ 1.6–6.6×; $r = \kappa +$ time-component
Cross-domain ρ_x small	Confirm	global 0.05, 58,665 country-years
κ–CDS co-movement	Construction-specific	cross-section 0.97 *and* time-series 0.79
Takaful κ India=Pakistan	Construction-specific	identical by loss-count; method-dependent
Spanning recovery	Construction-specific	98.5% best case; ∼80% at measured ρ_x
κ rating-ordered AAA→B	Confirm	monotone 1.23 → 9.23, sharp IG/HY jump
κ-curve slope (panel, significance)	Confirm	+2.85 bp/yr, t=4.7 clustered
P1 funding distribution distinct	Confirm	Cohen's $d = 0.64$ (claimed 0.782)
iCDS PV-riba < 4%; par-spread ι-free	Confirm	max 2.6%; $s^* = \kappa(1-\delta)$ ι-invariant
κ systematic vs. idiosyncratic	**Refine (new)**	first-PC 63% ⇒ domains, not names
Cross-domain *tail* independence	**Refine (new)**	COVID: credit +167%, mortality +14% together
P9 determinacy ceiling $e^{\bar{\kappa}} - 1$	Confirm+extend	12 sov reproduced; extended to 26, two instruments
P9 sovereign < corporate ceiling	Confirm+extend	107 issuers; corp more persistent (ρ 0.865)
P8 money-base auto-contraction	Confirm+extend	sukuk −2.03% vs paper −1.99%/100 bp
Benchmark-free κ: orderings (no Treasury)	Confirm	Spearman 0.84–0.89; wedge 473 ≈ 472 bp = UST yield layer
κ systematic share (monetary)	**Reconcile**	63% broad → 75% IG-core = P9's 80–100%
CDS–bond basis (dynamic iCDS)	Confirm	level corr 0.71; basis mean-reverts (HL ∼10 mo)
δ and regime robustness	Confirm	ordering δ-invariant; slope 13/13 both regimes
Commodity-κ level (gold 0.17)	Refine	holds on deep venues; regime-average, drifts 40–60%
Cross-domain lead–lag	**Null**	global hint dies when pooled (ρ=0.05, p=0.58)

Systematic factor = global risk?	Refine	R^2=0.44 on VIX+HY-OAS; ∼56% credit-specific
Benchmark-free wedge vs. UST (time)	Confirm	wedge tracks UST yield, ρ=0.93 over 2015–26
Liveness κ-index gate (P12)	Confirm	gate 1.4% vs 1.6%; breach = COVID 2020
κ predicts realized default	Confirm	3 defaulters rank top-3 of 20 on early $\hat{\kappa}$
P9 real-block feedback ψ exists	Confirm (sign)	growth$\times\Delta\kappa$ $\rho=-0.50$, $p<0.001$
κ-curve slope is a term premium	Confirm	12/12 positive, +1.65 pp vs. realized κ(6M)
$r-\kappa$ gap = inflation (Separation)	Confirm	cross-country ρ=0.96 (Turkey-anchored)
κ as force of mortality (HMD)	Confirm	Gompertz: μ doubles every 7.4 y, 41 tables
Recovery δ vs. realized defaults	Confirm	asset-based median 10% $\Rightarrow$ δ=0.10
κ as disaster arrival rate	Confirm	Poisson, median 1.68/yr, 194 countries
κ-curve in China (non-USD)	Confirm	AA corp 1.03 $\rightarrow$ 2.11% over 1–10 y
CEX funding pins to interest floor	Confirm	Binance median = 0.0001; 45% at floor
Feller $2\alpha\bar{\kappa} \geq \nu^2$ (non-neg.)	Confirm	18/20 sovereigns (fails only defaulters)
Takaful $\pi^* = \mathbb{E}$[loss] (RMA)	Confirm	premium 9.6% vs loss-cost 7.9%, ratio 0.82
Cross-domain ρ_x, 7 sleeves	Confirm	0.072; real domains ∼0, REER couples −0.49
Glass-monitor CUSUM (P13)	Confirm	warp caught in 5.5–11 mo, 66 issuers
Corporate−sovereign κ premium	Confirm	+29 bp G-spread (paper +38)
κ-curve in local currency (MYR)	Confirm	sukuk−govt 0.08 $\rightarrow$ 1.06%, upward
Wrapper-independence in local ccy	Confirm	MGII vs MGS mean 3.5 bp
Depth rule (live order book)	Confirm	7%-depth trade 3–13 bp; warp \$1.3–26 M
Manipulation cost (FBA vs cont.)	Confirm	1.7k–12.6k×; grind net extraction ≤ 0
Liveness conjunction trigger	Confirm	volume$\wedge\kappa$-stress co-occur 0 months
Riba test: net transfer at ι=0	Confirm	~ 0 (CI straddles 0); riba linear in ι
Protocol solvency (cadCAD MC)	Confirm	0/30 insolvent; fund self-sustaining
Censorship-proof sleeve (PPO)	Confirm	advantage −0.39; 0/30 beat honesty
AHJ convergence κ (keystone)	Confirm	basis half-life 3.5 h; $\kappa\sim$1,700/yr
κ-surface (rating × tenor)	Confirm	monotone both dims (5/5, 4/4)
κ–VIX loading state-dependent	Refine	+0.39 calm $\rightarrow$ +0.79 stress
κ-curve factor structure	Confirm	level 98.7%/slope/curvature (yield-curve-like)
κ-surface arbitrage-free	Confirm	forward $\kappa \geq 0$ in 15/15 tenor pairs
Base curve-shape sensitivity (P8)	Refine	slope shock: long bucket −11.8%

12 Extended program: analyses beyond the pre-specified set

Following the eighty-one pre-specified analyses above, we ran a further set (catalogued A91–A120 in the replication archive) extending the validation in three directions — the multi-modal oracle, the interest-versus-κ system comparison, and the dynamics and economics of κ. We summarise the findings here at the level of conclusions; the scripts and full output are in the archive. We report the null and refuted results alongside the confirmations, because a programme that only confirms is not testing.

κ as a hazard — confirmed. (i) *Convergence is universal*: the basis mean-reverts in 43/43 perpetual markets, median half-life 3.7 h, reproducing the BTC keystone measurement market by market (A114). (ii) *Out-of-sample default prediction*: the three sovereigns that actually defaulted (Lebanon, Sri Lanka, Ukraine) sat in the top decile of the cross-sectional κ distribution twelve months before default — 3/3 at every horizon (A113). (iii) *Mean-reverting intensity, not a random walk*: forward κ regresses on its own deviation with slope -0.26, AR(1) half-life $\sim$3.4 months (A109). (iv) *Cross-instrument and cross-asset coherence*: CDS-κ and bond-κ for the same sovereign co-move at 0.69 in changes (A111); sovereign-κ co-moves with the currency at +0.34 (A112). (v) *Within-rating information*: inside a rating bucket, κ deviation predicts the forward spread change at corr -0.47 — κ differentiates risk the static letter cannot (A119). (vi) *The risk premium is earned*: realised one-year excess returns rise monotonically with κ (low +0.6%, high +7.4%), evidence that bearing κ-risk is compensated — gain accompanying liability for loss (A116). (vii) $\sim$67% *of sovereign-κ variation is one systematic factor* (A120).

Interest is removable, with measured consequence — confirmed. The interest parameter is a committee choice (the policy rate is unchanged 47% of months, then jumps by decree) while κ is a continuous measurement; their correlation is -0.76 (A99). The practical payoff is the 2022 duration gradient (Table 3): conventional bonds lost up to -45% at thirty years to a rate decision, while the $\iota = 0$ formula carries no direct rate term at any tenor (A100; a model property, with the mixed-economy equilibrium caveat stated at the table), and across the panel rate risk is 40–52% of realised yield volatility, ballooning in the rate-shock year (A108). At $\iota = 0$ the funding drift is $\sim$zero versus the $\sim$8%/yr interest-floor drift on conventional venues, turning a one-directional charge into a two-sided contingent settlement (A110).

Multi-modal oracle security — confirmed. The three recovery channels are nearly mutually independent ($\rho \approx 0.05$, A91); a crisis moves only the channels in its causal path (2020 lit up credit and mortality but not crops; 2018 the reverse, A92–A95); and a median of three channels with two non-tradeable legs converts a $\sim$\$23M market attack into an effectively-infeasible two-institution data breach (A96).

Null and refuted results — reported honestly. The macro claim that interest-debt drags growth is *not* supported once development confounds and reverse causality are removed (within-country, lagged: the relationship collapses to zero, A107). κ does *not* lead the currency — FX moves first (A112). κ's monthly changes show *no* clean up-jump asymmetry (the extreme skew is an outlier artifact, A117). Sukuk and conventional κ are highly correlated (0.91) but *not* identical — a $\sim$1 pp level gap from maturity and demand, most of which is maturity (A118). The cross-country κ–inflation correlation (+0.61) is a fragility confound, not κ containing inflation (A103). These corrections sharpen, rather than weaken, the surviving claims.

13 Limitations

Our perpetual sample begins 2019–2023 (venue inception); the cross-domain panels are annual, so the credit$\times${agriculture, mortality} cells rest on short overlaps and we report the credit row with that caveat. CDS coverage is GCC/EM sovereign; OKX funding was geo-unavailable; the CoinGecko key is a demo tier (365-day RWA history). None of these affects the qualitative conclusions, all of which rest on multiple independent series. The COVID tail result is one natural experiment; a second systemic episode is needed to generalise the channel-specific reading. One test was inconclusive by construction: whether the CEX-versus-dYdX funding wedge tracks the policy rate over time could not be resolved on BTC/ETH, whose interest wedge is near zero (the signal lives in mid-cap alts), within the short 2023–26 window.

14 Discussion: realised hazard versus risk premium

The analyses establish, beyond much doubt, that κ *exists* and behaves as an intensity. They do not by themselves answer the question the programme ultimately cares about — whether replacing r with κ yields a genuinely interest-free system or re-expresses interest's economic function under a riba-free name. That question becomes tractable once κ is split into the part that *materialises* (the physical intensity κ_P, from realised defaults and losses) and the part the market *prices* (the risk-neutral $\kappa_Q = s/(1-\delta)$). The risk-premium share $1-\kappa_P/\kappa_Q$ is not a constant of the framework; it varies across domains by almost two orders of magnitude. For life and agricultural takaful it is about 18% (the USDA loss ratio is 0.82: premiums equal expected loss plus a modest margin), so κ is overwhelmingly *realised hazard* — one pays for events that occur. For credit the share is the great majority but is bounded rather than pinned: no sovereign USD sukuk has ever defaulted, so the share is large, yet the short default record only places a Poisson 95% upper bound near 1% on the sovereign default rate, which against a ∼70 bp spread implies a risk-premium share of *at least* ∼45% (and ∼85% at a plausible 0.2% default rate). Two caveats keep this honest and are worth stating: the share $1 - \mathrm{EL}/s$ is the *non-expected-loss* portion of the spread and so also contains liquidity and tax — an *upper bound* on the pure risk premium, and indeed our investment-grade figures sit at the top of the published credit-spread-puzzle band [Elton et al., 2001, Huang & Huang, 2012]; and the share is sensitive to the recovery assumption for high-yield, though the takaful-versus-credit gap (∼18% versus the great majority) is invariant to it. What is not in doubt is the order-of-magnitude difference between the domains, and the argument rests only on that.

This locates the programme's central tension precisely. Setting $\iota = 0$ genuinely removes the *predetermined* time-charge — the CEX funding floor pinned at exactly 0.0001 and gone at $\iota = 0$ (Section 5), the net transfer that vanishes at $\iota = 0$ and grows linearly in ι otherwise, the $r - \kappa$ = inflation decomposition at $\rho = 0.96$. That is real and it is not small. But it does not remove the risk premium *inside* κ. Whether that premium is interest in substance is, however, not a question economics can settle, because it turns on a point of law that economics does not contain. One reading holds that a systematic positive expected return for deferring consumption and bearing risk is the discount rate itself, and hence interest's substance however one contracts around it (the κ-curve's term premium and the regime-amplification of κ's loading on global risk are that premium in the term structure and the tail; and pricing default risk on debt raises further questions of *bay' al-dayn* and *gharar*). The opposing reading is older and, in Islamic commercial law, foundational: under *al-ghunm bil-ghurm* — entitlement to gain is tied to liability for loss — a contingent, loss-bearing, non-predetermined return is the *permitted* form of profit, the same form on which *mudarabah* and *musharakah* rest; what the prohibition forbids is the *predetermined and* *guaranteed* increment, not the accrual of return as such. The programme's $\iota = 0$ removes precisely the predetermined component (the funding floor pinned at 0.0001 and gone at $\iota = 0$; the net transfer that vanishes at $\iota = 0$). The honest reading is therefore that κ is riba-free in *form* everywhere (market-implied, outcome-contingent, varying) and a return of the *permitted kind* (contingent, loss-bearing), and that the residual question — whether such a return, where its loss almost never materialises, is riba in substance — is genuinely open between serious scholars. It is riba-free in substance beyond dispute only where it is realised hazard. Takaful is thus the programme's strongest ground; the κ-credit and κ-monetary ambitions are its most contested, and a κ-priced monetary base inherits whichever risk-premium share its asset base carries. Our data does not, and cannot, settle whether a market-contingent risk premium stripped of predetermined interest is riba in substance — that is the jurisprudential question the programme explicitly reserves to qualified scholars. What it can do, and does, is turn that question from a slogan into a measured quantity, and confirm that the corpus drew its own honesty boundary in the right place. The decomposition also carries a design corollary the companion product papers act on: because κ separates into a realised part κ_P and a priced part κ_Q, a substance-compliant credit instrument can price its cooperative-guarantee contribution at the *realised* hazard κ_P — the actuarially fair expected loss, exactly as a takaful contribution is set — rather than at the market-priced κ_Q, so that the contested risk-premium wedge $1 - \kappa_P/\kappa_Q$ is not *sold* for profit but retained by the mutual pool and returned as surplus (with a disclosed *wakala* fee for administration), which is precisely what collapses the credit case into the takaful case that needs no resolution of the open question.

Three further cuts make the tension sharper still. First, κ is most premium-driven in crises: the COVID credit-κ index rose +7 percentage points while realised defaults did not surge, so the spike was a repricing of risk, not of expected loss — the excess-bond-premium [Gilchrist & Zakrajšek, 2012] in κ form, and the same state-dependence the VIX loading shows. Second, the risk-premium share (on the same upper-bound basis) is *highest for the safest credit*: against long-run agency default rates it is near 100% at AAA/AA and falls to 41% at single-B (the credit-spread puzzle). The implication for the monetary programme is structural rather than incidental — a κ-priced base built on the safest assets, which is what a safe-asset base is, carries the *highest* risk-premium share, so the κ-Standard is the corpus's most premium-laden application by construction. Third, that premium is a market price of bearing risk — the same kind of compensation equity (and hence *mudarabah*) earns, the strongest economic argument that it is profit-from-risk rather than riba — but it earns a higher Sharpe ratio than equity and, for safe credit, accrues without realised two-sided loss, so even the equity analogy is only partial. None of this is a verdict; it is a map of exactly where the jurisprudential question lives, which is the most a dataset can honestly provide. The full reduction — the decomposition as the statement of the riba question, the *al-ghunm bil-ghurm* framing, and why competent scholars divide on it — is developed in a companion paper [Ahmed & Bhuyan, 2026, P16].

15 Conclusion

On an independent 17-provider dataset and eighty-one analyses, the κ-rate programme's central empirical claims survive: κ is a real, observable, venue-independent, non-negative, credit-ordered, upward-term-structured, mean-reverting intensity, categorically distinct from r (six orders of magnitude; $r = \kappa+$ a separable time component), and the interest parameter ι is an additive, removable layer of exactly the CEX interest baseline. The validation now reaches the monetary half of the corpus (Papers 8–9): the determinacy ceiling, the money-base auto-contraction, the sovereign<corporate ordering, and the benchmark-free estimator all reproduce and extend on inde-

pendent, larger data. The two genuine refinements are methodological hygiene — several headline κ figures (the κ–CDS correlation, the mean-reversion half-life, the systematic share and its implied buffer) are construction-specific and must be quoted with their method — and one new substantive caveat: average cross-domain correlation understates crisis co-movement along channel-specific lines, which sharpens Paper 14's safe-asset argument. We read the exercise as the strongest available evidence that the framework's direction is sound, precisely because it was assembled to break it and reports, transparently, what held and what must be qualified.

References

Ackerer, D., Hugonnier, J., & Jermann, U. (2025). Perpetual Futures Pricing. *Mathematical Finance*, DOI 10.1111/mafi.70018.

Ahmed, S., & Bhuyan, R. (2026b). From Perpetual Contracts to Islamic Credit: The Random Stopping Time Equivalence. SSRN 6322858.

Ahmed, S., & Bhuyan, R. (2026). The κ-Rate Research Programme, Papers 1–14. SSRN author page 10250376.

Ahmed, S., & Bhuyan, R. (2026). The Riba Question, Reduced: A Risk-Premium Decomposition of the Convergence Intensity.

Elton, E. J., Gruber, M. J., Agrawal, D., & Mann, C. (2001). Explaining the Rate Spread on Corporate Bonds. *Journal of Finance* 56(1), 247–277.

Huang, J.-Z., & Huang, M. (2012). How Much of the Corporate–Treasury Yield Spread Is Due to Credit Risk? *Review of Asset Pricing Studies* 2(2), 153–202.

Gilchrist, S., & Zakrajšek, E. (2012). Credit Spreads and Business Cycle Fluctuations. *American Economic Review* 102(4), 1692–1720.