The Riba Question, Reduced: A Risk-Premium Decomposition of the Convergence Intensity

Shehzad Ahmed* Rafiqul Bhuyan†

June 2026 — Preliminary draft

Abstract

The κ-rate research programme proposes the convergence intensity κ of Ackerer et al. [2025] as a riba-free replacement for the interest rate r. A companion paper shows, across eighty-one out-of-sample analyses, that κ is a real, measurable intensity and that the interest parameter ι — the predetermined, guaranteed component that defines riba al-nasiah — is empirically removable. That establishes κ is riba-free in *form*. It does not establish it is riba-free in *substance*, and the present paper isolates exactly what remains. We decompose κ into the part that materialises (the physical intensity κ_P, realised defaults and losses) and the part the market prices ($\kappa_Q = s/(1-\delta)$). The risk-premium share $1 - \kappa_P/\kappa_Q$ varies across domains by almost two orders of magnitude: it is a measured ∼18% for takaful (premiums equal expected loss plus a margin) and the great majority for sovereign credit (≥75% at the sample median, and ≥∼45% even under the most conservative pairing of the Poisson default bound with the tightest investment-grade spread; point ∼85–95%), highest for the safest credit, and largest in crises. We argue this reduces the entire riba verdict for the credit and monetary layer to a single, genuinely unsettled question of *fiqh*: whether a fully contingent, loss-bearing risk premium, stripped of every predetermined charge, is riba in substance, or the permitted form of return under *al-ghunm bil-ghurm*. We show why competent scholars divide on it, why the data cannot adjudicate between them, and why the strongest practical ground is the realised-hazard products (takaful) while the monetary ambition is the contested frontier. The contribution is not a verdict but a reduction: we turn a fifty-year slogan into a measured quantity and locate, to the basis point, where the question lives.

Keywords: riba, convergence intensity, risk premium, al-ghunm bil-ghurm, credit-spread puzzle, Islamic finance, takaful, physical vs risk-neutral measure.

1 The question that survives

The prohibition of riba is the organising constraint of Islamic finance, and the κ-rate programme's claim is that the convergence intensity κ prices time-and-risk without it. The empirical case for that claim is now substantial: on an independent, seventeen-provider dataset, κ is venue-independent, non-negative, credit-ordered, term-structured, mean-reverting, distinct from r across six orders of magnitude, and identifiable as the same object in credit, mortality, disaster and perpetual markets [Ahmed & Bhuyan, 2026, P15]. Crucially, the interest parameter ι — the layer the programme says

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

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

is removable — is removable in the data: on conventional venues the funding rate is pinned at a fixed floor that vanishes when $\iota = 0$, and the systematic transfer it induces goes to zero.

What riba *is*, in the classical definition, is the predetermined, guaranteed increment — *al-ribā huwa al-ziyāda al-mashrūṭa* — decoupled from the outcome of the underlying venture. Setting $\iota = 0$ removes precisely that. So the programme does not, and need not, claim that κ is not a return. It claims that κ is a return of the *permitted kind*. The question that survives the empirical work is therefore not "is κ interest?" but the sharper one this paper isolates and measures: *when the predetermined component is gone, what is the nature of what remains?*

2 Decomposition: realised hazard versus risk premium

For any instrument priced off a hazard intensity, the risk-neutral intensity the market charges,

$$\kappa_Q \;=\; \frac{s}{1-\delta}, \tag{1}$$

with s the credit spread (or premium rate) and δ the recovery, exceeds the physical intensity κ_P at which the insured event actually occurs. The wedge $\kappa_Q - \kappa_P$ is the risk premium: the compensation the holder earns for *bearing* the risk, over and above the loss that materialises. We summarise it by the *risk-premium share*

$$\Pi \;=\; 1 - \frac{\kappa_P}{\kappa_Q} \;=\; 1 - \frac{\mathrm{EL}}{s}, \qquad \mathrm{EL} = \kappa_P(1-\delta), \tag{2}$$

the fraction of the priced intensity that is *not* expected loss. Its structural form—under a Cox–Ingersoll–Ross intensity and the change of measure, $\Pi = |\lambda|/\alpha$, the ratio of the market price of κ-risk to the speed of mean reversion—is derived in the companion paper [Ahmed & Bhuyan, 2026, P17]; here Π is the quantity we estimate. Two properties of Π govern its interpretation and we state them before the numbers. First, Π is the *non-expected-loss* share of the spread, which also contains liquidity and tax components we do not strip out; it is therefore an *upper bound* on the pure risk premium, and our estimates should — and do — sit at the upper end of the published credit-spread-puzzle range [Elton et al., 2001, Huang & Huang, 2012]. Second, where realised events are rare κ_P is not zero but bounded: with d events in N exposure-years the Poisson 95% upper bound on κ_P is, for $d = 0$, the rule of three $3/N$, so Π is a *bound*, not a point.

3 The measured ground

Estimating κ_Q from observed spreads and κ_P from realised loss experience across four domains gives a risk-premium share that ranges over almost two orders of magnitude (Table 1).

In insurance, κ is overwhelmingly realised hazard — one pays approximately for events that occur. In credit, and above all in sovereign credit, κ is mostly a risk premium: no sovereign USD sukuk has ever defaulted, yet sovereigns are charged a convergence intensity of order 4%. How much of that is premium is bounded, not pinned, by the default record — $\geq$75% at the sovereign-median intensity, and at least $\sim$45% even under the most conservative reading (the Poisson bound against the tightest investment-grade spread; Table 1), plausibly far more — but the order-of-magnitude gap between the domains is robust to the recovery assumption and to the bounding, and the argument rests only on that gap.

Three further measurements sharpen the credit case rather than soften it. *(i)* The premium is most dominant in crises: the 2020 credit-κ index rose seven percentage points while realised defaults

Domain	κ_Q (priced)	κ_P (realised)	risk-premium share Π
Life / agricultural takaful	0.18% / 9.6%	0.15% / 7.9%	$\sim$**18%** (measured)
Corporate credit (sukuk)	2.2%	0.19%	$\approx$**91%** (upper bound)
Sovereign credit (sukuk)	4.0%	$\leq$ 1.0% (Poisson UB)	$\geq$**75%**, point $\sim$85–95%

Table 1: The risk-premium share of κ by domain, with each row computed from its own columns ($\Pi = 1 - \kappa_P/\kappa_Q$). Takaful is a direct measurement (the USDA Risk Management Agency loss ratio is 0.82: premiums equal expected loss plus an $\sim$18% margin, on \$1.25 trn of real liability). The credit figures are upper bounds on the pure premium (the spread also carries liquidity and tax); the sovereign figure is a lower bound, because zero sovereign-sukuk defaults in a short record place only a Poisson upper bound near 1% on the default rate: $1 - 1.0/4.0 = 75\%$ at the sovereign-median $\kappa_Q = 4.0\%$. Even the most conservative pairing — the same Poisson bound set against the *tightest* investment-grade sovereign spread ($\sim$70 bp, $\kappa_Q = 1.75\%$ at $\delta = 0.60$) — still gives $\Pi \geq 1 - 1.0/1.75 \approx 43\%$; that floor, not the median-based 75%, is the number quoted when we say "at least $\sim$45% under the most conservative reading."

did not surge, so the spike was a repricing of risk, not of expected loss — the excess-bond-premium of Gilchrist & Zakrajšek [2012] in κ form. *(ii)* The risk-premium share is *highest for the safest credit* — near the upper bound at AAA/AA, falling toward high-yield where default is real — so a κ-priced monetary base built on the safest assets, which is what a safe-asset base is, carries the *highest* premium share by construction. *(iii)* The premium is, in kind, a market price of bearing risk — the same compensation equity, and hence *mudarabah*, earns — but per unit of risk it earns a higher Sharpe ratio than equity (the credit-spread puzzle), so even the equity analogy is only partial.

4 The jurisprudential reduction

The economics can carry the question no further, because what remains turns on a point of law that economics does not contain. We state the two readings as their serious proponents would, because the division between them is the result.

The permissive reading (al-ghunm bil-ghurm). Islamic commercial law is built on the principle that entitlement to gain is tied to liability for loss — *al-ghunm bil-ghurm*. It is the basis on which *mudarabah* and *musharakah* are lawful: the financier earns precisely *for bearing* the venture's risk. What the prohibition forbids is not that a return accrues, or is positive in expectation, or compensates the passage of time; it is the predetermined and guaranteed increment, decoupled from outcome. The κ-holder is genuinely exposed to default — that exposure materialises in the high-yield names and is real, not nominal, even where it has not yet been realised in the safe ones — and rarity of loss does not convert a risk-bearing return into a stipulated one. On this reading κ-credit is *mudarabah*- like and permitted, and the programme is the most rigorous riba-free pricing foundation built to date.

The strict reading. A systematic, priced, expected-positive return *is*, in modern asset pricing, the discount rate itself; and the premium in κ is largely systematic — it loads on global risk and amplifies in the tail. A discount rate, however one contracts around it, is interest's substance. Moreover, the credit application prices *default risk on debt*, which raises *bay' al-dayn* (trading debt

below par) and *gharar* independently of the riba question; on this reading [in the spirit of El-Gamal, 2006] the asset-backed, realised-hazard structures — takaful, salam, ijārah — are the only safe ground, and the credit and monetary layer re-expresses interest under a contingent label.

Why the data cannot decide. Both arguments are serious and held by competent scholars. The decomposition does not favour either: it confirms the premium is real (the permissive reading's exposure is genuine) *and* that it is systematic and, for safe credit, largely unrealised (the strict reading's concern is genuine). No additional measurement bears on the choice between them, because the choice is between two principles of law applied to the same measured object. Demonstrating that — rigorously, with the quantity that would have to be argued laid bare — is the contribution.

5 Implication: the realised-hazard ground and the contested frontier

The reduction has a direct corollary for where the programme stands on firm ground and where it does not. The takaful family is sound on both form and substance: $\iota = 0$ removes the predetermined charge, and the residual premium is a measured $\sim$18% margin over realised loss, the least contestable form of risk-bearing return. It requires no resolution of the open question. The κ-Standard and the monetary architecture, by contrast, stand on exactly the contested ground, and — by the safest-credit result — on its most premium-laden part. They are best understood not as settled applications but as a research frontier whose permissibility is gated on the one ruling. One does not build a monetary system on a question one has shown to be open; one brings the question, fully formed, to those whose office it is to answer it.

A product-layer narrowing. The decomposition itself makes available a resolution short of the ruling. Because κ separates into the realised κ_P and the priced κ_Q, a compliant credit instrument need not be priced at the market κ_Q. Structured as a cooperative guarantee — a *tabarru* pool in which members mutually indemnify realised default loss, gated on insurable interest and possession of the reference obligation — it prices its contribution at the realised hazard κ_P, the actuarially fair expected loss, exactly as a takaful contribution is set. The contested wedge $\kappa_Q - \kappa_P$ is then not *sold* to a protection writer for profit but retained by the pool and returned as surplus, with a disclosed *wakala* fee for administration. So priced, credit-loss cover is structurally the takaful case established above: it inherits that case's soundness on form and substance and needs no resolution of the open question. This does not dissolve the frontier — a κ_Q-priced, tradeable form, and any monetary base that carries a market risk premium, remain squarely on the contested ground — but it relocates and narrows it. The open question ceases to be "may one charge a priced risk premium?" for the whole credit layer and becomes the narrower pair that survives the κ_P-priced mutual structure: whether tokenised title is valid possession, and whether the *wakala* loading is a fee for service or the premium under another name. Both are narrower; both remain reserved.

6 Conclusion

The empirical programme proves κ is real and that ι — the predetermined, guaranteed charge that defines riba al-nasiah — is removable. What remains, once it is removed, is a return of the permitted kind: contingent and loss-bearing. The risk-premium decomposition shows how much of κ that return constitutes in each domain — a measured $\sim$18% in takaful, the great majority in

credit — and thereby reduces the entire riba verdict for the credit and monetary layer to a single question: whether a fully contingent, loss-bearing risk premium, stripped of every predetermined charge, is riba in substance or the lawful form of profit under *al-ghunm bil-ghurm*. The data cannot answer it; it can only state it exactly, and show that the corpus was right to reserve it. That is the most a measurement can honestly provide, and it is, we believe, the most precisely formed version of the question yet brought to the scholars. *wa-Llāhu a'lam.*

References

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

Ahmed, S., & Bhuyan, R. (2026). The Separation Theorem: Interest as a Removable Parameter in No-Arbitrage Pricing.

Ahmed, S., & Bhuyan, R. (2026). The κ-Framework Meets the Data: An Out-of-Sample Empirical Validation Across Seventeen Data Providers.

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.

El-Gamal, M. A. (2006). *Islamic Finance: Law, Economics, and Practice*. Cambridge University Press.