

The Rate That Wasn't There

The Rate That Wasn't There

*How a missing line of code became a financial primitive,
and where the search for it had to stop*

Shehzad Ahmed
with Rafiqul Bhuyan

Baraka Protocol · Arcus Quant Fund
Independent University Bangladesh · Alabama A&M University

July 2026

A narrative companion to *The κ -Rate Research Programme*: seventeen collected papers

This book makes no claim of overall Shariah compliance and asserts no fatwa.
The jurisprudential question it isolates is reserved, explicitly, to qualified scholars.

*"The contract is not the only place riba can enter.
It may also enter through the rate used to price it."*

*"The most load-bearing number in finance
turned out to be optional."*

*"I did not earn independence from the scholars.
I built the clearest possible question to bring them."*

Contents

The Argument in Brief	6
Author's Note	8
1 Two in the Morning	10
2 The Theorem Under It	12
3 Is the Rate an Axiom?	14
4 A Chain With No Rate	16
5 Meeting the Data	18
6 Where Riba Enters the Price	20
7 What It Means	22
8 A Reader's Guide to the Seventeen Papers	24
9 Key Results at a Glance	26
10 Glossary of Key Terms	27
About the Authors	29

The Argument in Brief

For the reader who wants the whole thing on one page before the story, here is the spine of the argument in five steps.

1. **The observation.** A live derivatives venue (dYdX v4) prices perpetual futures with *no interest term* in its funding formula, only a premium measuring how far the contract has drifted from the index. The most load-bearing number in finance, the risk-free rate, turns out to be optional in at least one working system.
2. **The theorem.** A 2025 result (Ackerer, Hugonnier, Jermann, *Mathematical Finance*) proves a perpetual contract has a *unique* no-arbitrage price with the interest parameter set to zero, expressed not as a discounted cashflow but as an expectation at a random event time. We show this is formally the same equation as the Duffie–Singleton credit model, with a convergence intensity κ in the slot the interest rate occupies. Interest and risk are *separable*; the interest part is *removable*. Mathematically, a positive κ moves the pole of the pricing operator off zero, leaving the *riba*-free point regular where the conventional model diverges.
3. **The primitive.** What remains, κ , is not an interest rate but a *hazard*: the rate of a real event (a default, a loss, a price converging), read from markets, never stipulated. One primitive prices an entire product line: credit spreads, takaful contributions, a benchmark curve; and, we show, a determinate monetary economy with money backed by equity rather than debt.
4. **The evidence.** Across an 81-analysis validation ledger, comprising primary tests, secondary analyses, and separately labelled exploratory work on an externally sourced seventeen-provider dataset, κ behaves like a hazard and nothing like a rate: positive, ordered by credit quality, recovered as the same kind of object from credit spreads, perpetual mean-reversion, and insurance loss records (three channels that share no inputs), and spanning four orders of magnitude that no single interest rate could.
5. **The reserved question.** Decomposing κ into realised hazard and risk premium shows the *riba* question reduces to one measured quantity, which varies from $\sim 18\%$ (takaful) to the great majority (sovereign credit). Removing the interest parameter removes the *predetermined* charge that *defines* *riba*; whether the contingent, loss-bearing premium that remains is itself *riba* is a question of *fiqh*, not of data, and it is reserved, in the clearest form we could build it, to qualified scholars.

The chapters that follow tell how this was found; the seventeen papers behind Part II hold the

proofs, the measurements, and the boundary.

Author's Note

This is the story of a number that everyone in finance treats as bedrock, the proof that it was never bedrock at all, and the honest place where that proof runs out and a question begins that no proof can finish.

I did not set out to write seventeen papers. I set out, at two in the morning, reading exchange code I had no business reading, to understand why one line was missing. The chapters that follow are the narrative: how the missing line led to a theorem, the theorem to a primitive, the primitive to a chain, the chain to an avalanche of evidence, and the evidence, finally, to a single question I am not qualified to answer and neither, it turns out, is the data.

Part I tells that journey in seven chapters, each through the lens that the work demanded at that stage: the builder's, the mathematician's, the economist's, the crypto-native's, the empiricist's, the jurist's, and at the end, plainly, the human one. Part II is the work itself: the seventeen papers, bound here in reading order, from the perpetual-futures foundation to the separation theorem that holds the whole thing up.

I have tried, throughout, to honour one rule above all others: to state the limits as loudly as the claims. The space this work touches, interest-free finance, has earned its skepticism through fifty years of overclaiming. I would rather be the person who marked the boundary than the one who pretended there wasn't one. Where I am sure, I say so. Where the question belongs to others, I have tried to hand it to them in the clearest form I could build, and then to step back.

A note on how to read this. If you are a builder, Chapter 1 is your door; a mathematician, Chapter 2 and Paper 17; an economist, Chapter 3 and Papers 8–9; a student of *fiqh*, Chapter 6 and Paper 11. But the book is built to be read straight through, because the meaning is in the sequence: the way each lens hands the problem to the next, and the way they all arrive, together, at the same reserved question. A glossary at the end of this front matter defines the terms of art from all five fields; a table of key results maps each claim to the paper that proves it.

S.A., July 2026

Part I

The Journey

Seven chapters, seven lenses, one question.

Chapter 1

Two in the Morning

The interest term in pricing is a configuration, not a law of nature.

Fintech has spent fifteen years rebuilding the front of finance. Better onboarding, instant rails, clean APIs laid over a tired core. What almost nobody touches is the core itself: the pricing engine. Every product anyone ships, a loan, a savings yield, an insurance premium, a structured note, is priced off one number underneath: a risk-free interest rate. It is the deepest dependency in the stack, and it is treated as immovable.

I found a reason to touch it by accident.

Perpetual futures use a payment called *funding* to keep the contract tethered to the real price. On Binance, the funding formula has two parts: a premium, measuring how far the contract has drifted from the index, and a small fixed interest rate, around eleven percent a year, charged just for holding the position over time. On dYdX v4, the second part is gone. Funding is pure premium: the gap between contract and spot, divided by spot. The code does not even have a field for an interest rate. A live venue, real volume, pricing a derivative with no interest in the formula, and nothing breaks.

THE LINE THAT WASN'T THERE

$$\underbrace{F^{\text{Binance}} = \underbrace{\frac{P_{\text{perp}} - P_{\text{index}}}{P_{\text{index}}}}_{\text{premium}} + \underbrace{\frac{\approx 0.01\%}{8h}}_{\text{interest}}}_{\text{interest charged for holding over time}} \implies \underbrace{F^{\text{dYdX v4}} = \frac{P_{\text{perp}} - P_{\text{index}}}{P_{\text{index}}}}_{\text{premium only; no interest field}}$$

One term, deleted. Everything in this book follows from asking what could be built on its absence.

That is not trivia. It means the interest term in pricing is a configuration, not a law of nature. And if you can remove it cleanly from one instrument, the builder's question becomes obvious: what else can you rebuild on whatever replaces it?

What replaces it I call *kappa*. Where an interest rate charges for time, kappa measures the rate of a real event: a default, a loss, a price converging. It is read from the market, not stipulated in

a contract. I did not yet know, that night, that there was a theorem under it. I knew only that a number I had been taught to treat as the price of time had just been deleted from a working system, and the system had not noticed.

The reason this mattered to me was not abstract. Roughly \$3.9 trillion of Islamic finance exists today, growing ten to fifteen percent a year, serving close to 1.8 billion people. And it runs on workarounds: products structured to look interest-free while still pricing off SOFR underneath. The form is clean; the pricing is not. I had spent years close enough to that industry to feel the quiet discomfort of it. And here, in a funding formula on a crypto exchange, was the seam. The defect everyone wraps around lives in the pricing layer. The pricing layer is exactly what the missing line touched.

So I started pulling the thread.

Chapter 2

The Theorem Under It

The discount was never load-bearing for finiteness. The hazard was.

A missing line is an accident. A primitive needs a proof. The thread led, within weeks, to a small mathematical fact with large consequences, and to the realisation that the fact had only just become provable.

Take the reduced-form credit model of Duffie and Singleton. A defaultable claim is priced as an expectation, under the pricing measure, of a discounted recovery paid at a random default time τ with hazard rate λ . In the clean case it collapses to

$$P_t = \frac{\lambda}{r + \lambda} X_t.$$

Now set $r = 0$. The price does not diverge. It becomes X_t . It stays finite, because the stopping time τ , not the discount factor, is what made the integral converge. The discount was never load-bearing for finiteness. *The hazard was.*

In 2025, Akerer, Hugonnier and Jermann proved the perpetual-futures analogue from first principles, in *Mathematical Finance*. With the interest parameter $\iota = 0$ and a convergence intensity $\kappa > 0$, a perpetual has a unique no-arbitrage price, and it admits a random-time representation: an *undiscounted* expectation $f_t = \mathbb{E}[X_\theta]$, evaluated at an event time θ whose law is governed entirely by κ . No factor of the form e^{-rt} appears anywhere.

Place the two side by side and they are the same equation. The map is exact: $\theta \leftrightarrow \tau$, the event time; $\kappa \leftrightarrow \lambda$, the event intensity; $\iota \leftrightarrow r$, the charge for time; $X_\theta \leftrightarrow R_\tau$, the value at the event. My co-author and I proved this isomorphism and drew from it the statement that became the spine of everything after: *the risk content and the interest content of a no-arbitrage price are separable, and the second is removable*. Sixty years of reduced-form credit theory had fused a hazard and a discount into one short-rate-style expression purely out of convention, not necessity.

Why was this recent rather than ancient? Because the *tool* was recent. In a world with no second source of convergence, removing the discount is fatal: the consol, worth c/r , diverges as $r \rightarrow 0$, and that single fact is why “just set the rate to zero” looked impossible for fifty years. What I would only fully understand much later, when I sat down to write the separation theorem properly (Paper 17 of this volume), is *why* the perpetual escapes it. Hold the real-event intensity fixed and vary the pure time-charge (the credit form of the operator, which coincides with the perpetual exactly at the riba-free point), and the pricing operator has a pole, a singularity, as a

function of the interest parameter. In the conventional perpetual that pole sits exactly at $r = 0$, which is why you cannot remove it. A positive convergence intensity moves the pole. It sits at $\iota = -\kappa$, a strictly positive distance away, leaving the riba-free point $\iota = 0$ as an ordinary, regular, removable point of the operator. That displacement of the pole off the origin, by the event intensity κ , is the separation theorem. Within this class of event-time pricing models, the pure time-charge is not required for finiteness or uniqueness: with $\kappa > 0$ supplying the stopping mechanism, $\iota = 0$ is a regular point and the price stays well defined. Interest is not structural to the no-arbitrage pricing of event-anchored claims. It is a separable term the standard models happened to fuse.

I had a primitive with a proof under it. The builder's question, what else can you build on it, was now open in earnest.

Chapter 3

Is the Rate an Axiom?

The axiom was a choice.

Open any asset-pricing text and the first object you meet is r , the risk-free rate. It is treated less like a parameter than like an axiom: the price of time, assumed positive, the bedrock everything else sits on. The theorem had just told me it was a choice. The economist in me wanted to know how far the choice reached. Could you run not just an instrument but an *economy* without it?

Start with what r actually bundles. Since Fisher and Wicksell we have known the interest rate is not one thing: a pure time-preference component, a charge for deferral, fused with a risk component, compensation for real uncertainty. The fusion has been invisible ever since, because nobody had a reason to pull them apart. The theorem was the reason. What is left, when you pull the time-charge out, is close to Wicksell's natural rate (a rate grounded in real economic dynamics rather than monetary policy) but made rigorous, endogenous, read from markets rather than set by a committee, non-negative by construction, and zero when no real event can occur.

So we built the monetary model. Base money is kappa-priced sovereign equity rather than interest-bearing debt. Policy is conducted by quantity rather than a policy rate. Credit is intermediated through profit-and-loss sharing. It has a closed-form determinacy condition, $\gamma_s < e^{\bar{\kappa}} - 1$, in which kappa occupies the precise slot the interest rate occupies in conventional fiscal theory. Time preference stays in household preferences; what leaves is the interest *contract* from the asset space. The economy is determinate. Because money is backed by real equity, it escapes the real-bills indeterminacy that sinks naive proposals. And it is Cantillon-neutral: since money's value is pinned to its backing, issuance tethered to output transfers no purchasing power to first recipients. The results are Dynare-validated and disciplined against real sovereign data. Those are Papers 8 and 9 of this volume.

I am not claiming a free lunch, and I held that line from the start. An interest-free economy has, by construction, no instrument that pays the same in every state; the riskless intertemporal trade is gone. For a representative agent that costs nothing. Under heterogeneity it costs a measurable but small amount: on the order of single basis points to low tens of basis points of lifetime consumption in a two-agent general equilibrium, borne mostly by the more risk-averse agent. The long-horizon behaviour is the deepest open question, and I have held it open the entire way.

But the claims an economist should sit with are narrow and large at once. First, the risk-free interest parameter is removable within the event-time pricing architecture developed here.

Second, the monetary model admits a determinate equilibrium without an interest-bearing asset. Within those structures, the conventional rate is a design choice rather than a mathematical necessity. The axiom was a choice.

Chapter 4

A Chain With No Rate

$\iota = 0$ is the substrate, not a setting.

Theory is one thing. The thing that makes interest-free pricing actually *work* is a measured event rate, not a committee. And there is exactly one place you can compute that rate in the open, block by block, with no one to trust: a blockchain.

So we built the obvious thing. A Layer 1 where kappa is not an imported oracle but a quantity the chain computes from its own activity and writes into consensus, into every block. No explicit interest parameter appears in the implemented consensus, reward or contract-pricing formulas: not in how contracts are priced, not in the validator rewards, not in any instrument on top. Where interest could re-enter indirectly, and what stands in the way, is traced in the exhibit below. $\iota = 0$ is the substrate, not a setting. On the public testnet, kappa sits in consensus and prototype contracts read it directly through a `block.kappa` precompile, pricing credit and takaful off it. Escrow in the chain's native coin, testnet assets throughout. A funded and solvent perpetual market. Manipulation defences written into the chain itself. It runs, end to end, under test conditions.

This part of the journey taught me something the papers don't fully convey: how much of the work of an interest-free system is *defence*. When your key number is a measurement rather than a decree, the adversary stops attacking the contract and starts attacking the estimator. A frequent-batch auction instead of a last-trade mark. A thirty-block time-weighted average. A clamp. A stake-weighted median. A change-point monitor watching the oracle for the signature of manipulation. A depth rule that caps how much value a manipulated print can reprice, relative to bonded stake. I simulated all of it (cadCAD for the macro dynamics, reinforcement learning for the adversaries, game theory for the equilibria) and the manipulation frontier turned out to have a closed form. That is Paper 3, and Papers 12 through 14, the frontier of the thing.

WHERE INTEREST COULD RE-ENTER, AND WHAT STANDS IN THE WAY

Layer		Re-entry channel	Defence	Class
Funding formula	for-	an explicit interest term	none exists; the dYdX v4 form has no r field	proof
κ estimator & consensus		rate inputs in the median	funding-implied intensities only; integer arithmetic; no external curve enters	proof
Validator rewards	re-	inflationary block reward (structural interest)	no inflationary reward in the protocol	proof
Market data beneath κ		venue prices that embed rates	contamination test: κ loads -0.3 to -0.75 on the risk-free curve, not $+1$; dYdX funding oscillates through zero (P2, P15)	measured
Benchmark curve		anchoring to the Treasury curve	the κ -curve is built benchmark-free; the reserve currency's time-value layer is disclosed, not removed (P2a)	measured
Collateral & escrow		yield on posted collateral	native-coin escrow pays none; the external-collateral path is gated until solvency and governance exist	governed
Settlement assets	as-	stablecoins or bridged assets importing external yield	none in the substrate; any future bridge passes the $\iota = 0$ ante-gate and the board allowlist	governed

No explicit interest parameter appears in any implemented formula (the proof rows). Independence from interest-bearing markets is measured, not assumed (the measured rows). What cannot be proven or measured is governed (the governed rows). Full economic independence from interest is partly a theorem, partly an evidence file, and partly a standing governance obligation, and the chain treats it as all three.

The crypto-native reason to care, beyond the elegance: some of the best-funded “Islamic chains” raised hundreds of millions and still shipped an inflationary block reward that is structurally interest, sitting quietly under a Shariah oracle. They spent at the app layer. The defect was at the monetary-primitive layer. That is the layer this rebuilds. A missing parameter in a funding formula turned out to be the seam where you can pull interest out of money itself, and then out of the network that mints it.

Chapter 5

Meeting the Data

A quantity you can read off a register of crop failures cannot be interest.

By this point I had a theorem, an economy, and a chain. What I did not yet have was an answer to the only question a skeptic should ask: is kappa a real thing, or is it the interest rate wearing a new word?

So I tried to break it. I built an externally sourced holdout dataset—seventeen data providers, about 1.1 gigabytes of credit spreads, perpetual funding, insurance loss records, mortality tables, and disaster registries—and ran, in the end, eighty-one analyses against it under one discipline I held without exception: *reproduce first*. Before treating any disagreement with the corpus as a discrepancy, reproduce the original paper’s own number, with its own method, from its own data. That rule mattered. It dissolved three of my own early “corrections”: the corpus was more careful than my red-teaming first assumed. One more disclosure governs how all of it should be read: this is not independent replication by unaffiliated investigators. It is adversarial internal validation, run on data the original results never touched.

THE DATASET: ~1.1 GB, 17 PROVIDERS, BUILT TO BREAK THE CLAIM

Credit & rates (Cbonds): 921 USD bonds / 900k bond-days, 121 issuers, sukuk *and* conventional; 51 sovereign CDS series; 215 local-currency bonds (IDR, MYR, TRY, SAR); local yield curves (India, Malaysia, Brazil, Mexico, China).

Cross-domain (World Bank, WHO, HMD, EM-DAT, FAOSTAT, NASA, IMF): mortality (life tables to 1751), agriculture (226 countries), disasters (16,853 records), climate (211 countries).

Markets (dYdX, Binance, Bybit, Yahoo): perpetual funding on three venues, gold/commodities, FX for 14 currencies, live order-book depth.

Contamination test (FRED): risk-free curves, VIX, credit OAS, to check κ is not r relabeled.

81 analyses across the primary validation ledger, with additional exploratory extensions reported separately; reproduce-first throughout.

What came back was the empirical signature of a hazard, not a rate. Domain-specific kappa intensities are recovered through three independently constructed channels that share no inputs: from credit spreads, as $\hat{\kappa} = s/(1 - \delta)$; from the mean-reversion of perpetual prices; and from

realised insurance loss frequency. The three return the same *kind* of object, ordered by credit quality, distinct from, but not economically independent of, the risk-free curve. On real sovereign sukuk it steps clearly from investment grade to high yield exactly as a default hazard should. Its magnitudes span four orders: roughly 0.02 to 0.2 per year for credit and insurance, against 10^2 to 10^3 per year for liquid-perpetual convergence. No single interest rate is at once two percent and eighty thousand percent per annum. A quantity you can read off a register of crop failures, where there is no bond market and no interest rate at all, cannot be interest. That is Paper 15.

THREE INDEPENDENT CHANNELS, ONE KIND OF OBJECT, AND IT IS NOT r

Channel (no shared inputs)	κ estimated from	Magnitude / year
Credit spreads	$\hat{\kappa} = s / (1 - \delta)$	0.02–0.2
Perpetual basis	mean-reversion half-life	10^2 – 10^3
Insurance losses	realised claim frequency	0.07–0.15

Ordered by credit rating; loading on the risk-free rate = -0.3 to -0.75 (not $+1$). An interest rate is not at once 2% and 80,000% per annum; a hazard can be, because it measures the event, not the time.

I want to be honest about how this chapter actually felt, because it is the turn of the whole book. I had built all of this hoping the data would hand me a verdict: proof that kappa is clean, that I could carry to the world and to the scholars and say, *see, it is done*. The data did something else. It confirmed, overwhelmingly, that kappa is real and that interest is removable. And then, when I pushed it to answer the deeper question (not “is kappa real?” but “is what remains, after you remove interest, itself a forbidden gain?”), it gave me a number instead of an answer. And the number pointed somewhere I had not expected to go.

Chapter 6

Where Riba Enters the Price

The riba is not in the contract terms. It is in the rate used to price it.

Islamic finance has done extraordinary work on the *form* of contracts. Murabaha, ijara, sukuk: the whole apparatus exists to keep the letter of a transaction free of stipulated interest. And yet a sukuk is structured impeccably, then priced as a spread over LIBOR, its cashflows discounted with an interest-bearing factor. The form is clean. The pricing is not.

Naming the problem precisely is what makes it solvable. The contract is not the only place riba can enter. It can also enter through the pricing architecture beneath an otherwise permissible contract, and that is where this defect lives: in the rate used to price the contract. *Riba al-nasiah* is a charge for deferment: a charge for the mere passage of time, independent of outcome. A discount-rate parameter encodes exactly such a predetermined charge for time, although whether benchmark use alone changes a contract's legal character remains a question of *fiqh*, not of mathematics. No amount of contractual structuring reaches a defect that sits one layer below the contract. The theorem removes exactly that charge, and only that charge. I should name what kind of claim this is: that riba's substance can live in the pricing layer, below the contract, is itself a jurisprudential position, not a theorem, and it is the central claim this programme carries to the scholars.

Set kappa against El-Gamal's three marks of riba (a charge that is predetermined, owed regardless of outcome, and levied for time alone) and it fails all three. It is measured after the fact, not stipulated. It is contingent on the outcome; indeed it *is* the outcome. And it prices an event, not the passage of time. On the marks that define the prohibition, kappa is not interest.

But the data, when I decomposed it, would not let me stop there. This is the heart of what I learned. Split kappa into the part that actually *materialises* (realised defaults and losses) and the part the market merely *prices* (the risk premium). The risk-premium share varies across domains by almost two orders of magnitude. In takaful it is about eighteen percent: you pay, very nearly, for what actually happens. In sovereign credit it is the great majority: no sovereign sukuk has ever defaulted, yet a convergence intensity of order four percent is charged, which is, functionally, a return for bearing risk over time. I had found, and measured, the place where the easy story breaks. Where kappa is realised hazard, the case that it is riba-free in substance is at its strongest. Where it is mostly risk premium, the harder question opens. That is Paper 16, the reduction.

THE RIBA QUESTION, MEASURED: WHERE THE LINE ACTUALLY FALLS

Domain	priced (κ_Q)	realised (κ_P)	risk-premium share
Takaful (life / crop)	9.6% / 0.18%	7.9% / 0.15%	~18% (measured)
Corporate credit ($\delta = 0.60$ convention)	2.2%	0.19%	~91% (upper bound)
at measured recovery $\delta \approx 0.10$	0.98%	0.19%	~81% (upper bound)
Sovereign credit	4.0%	$\leq 1.0\%$	$\geq 75\%$, point ~85–95%

Realised hazard in insurance (you pay for what happens); mostly risk premium in credit. Each share is computed from its own row ($1 - \kappa_P / \kappa_Q$); even the most conservative pairing (the Poisson default bound against the tightest investment-grade spread) keeps the sovereign share above ~45% (Paper 16). The corporate row is re-run at the recovery my own default study measured on asset-based sukuk ($\delta \approx 0.10$, Paper 2c): the share falls to ~81%, so the convention overstates it, and honesty here costs nothing and buys the conservative direction; the ~45% sovereign floor survives the same substitution. The riba question is quietest for takaful and loudest for sovereign credit: located, not answered. Every share is conditional on the stated recovery convention, physical-event estimator and observation window; the point estimates should be read alongside the conservative bounds and sensitivity analyses of Papers 15 and 16.

I make no claim of overall compliance and assert no fatwa. This addresses *riba al-nasiah*, at the pricing layer, and nothing more. It does not, on its own, settle *gharar*, *maysir*, or *qabd*: distinct pillars that belong to qualified scholars, not to a formula. For the asset-backed class we propose four product-layer conditions; whether they *satisfy* the pillars is a board's determination, not mine. The cash-settled synthetic is the hardest case of all, and we treat it last, not first. As I write, the argument is in front of scholars, and what I have asked of them is not a blessing but a critique: to find where it errs.

Chapter 7

What It Means

I did not earn independence from the scholars. I built the clearest question to bring them.

I set out to be the builder who could prove the thing halal and not need the scholars' permission. I want to tell you, plainly, where that ended, because the ending is the most valuable thing in this book.

Two results are established within the models and evidence presented here, and they are solid. Interest is *removable* within the event-time pricing architecture, not rhetorically but mathematically, because a positive convergence intensity pushes the pole of the pricing operator off the origin, where the conventional model's pole sits at zero and cannot be moved. And κ is *real*: positive, ordered by risk, the same class of event intensity across credit and mortality and disaster and perpetual basis, behaving like a hazard and nothing like a rate. The ledger can run without *riba al-nasiah*. That is a genuine result, and it is not small.

One question could not be settled, and showing why the evidence cannot settle it is the actual discovery. Whether the risk premium that remains, after interest is gone, is itself *riba* in substance. I went looking for data to settle it, and the data told me, to the basis point, that it cannot. The question is jurisprudential, not empirical. There is no eighty-second analysis that would change that.

For a while I read that as a defeat. It is not, and here is why. The reflex ("a risk premium is just interest in disguise") is the incomplete half. The older and load-bearing half is *al-ghunm bil-ghurm*: entitlement to gain is tied to liability for loss. It is the principle on which *mudarabah* and *musharakah* are lawful. A contingent, loss-bearing return is the *permitted* form of profit. *Riba* was never "return that accrues over time"; it was the *predetermined, guaranteed* increment, decoupled from outcome, and $\iota = 0$ removes exactly that. So I did not fail to remove the *riba*. I removed the part that is definitionally *riba*, and what remains is a return of the permitted kind, whose permissibility, in the contested credit case, is a real open question between serious scholars, on which the data is silent because the data must be silent.

So what do seventeen papers and a chain and a hundred experiments amount to? Not the answer I wanted. Something I have come to think is better. I did not earn independence from the scholars. I built the most precisely formed version of the question anyone has carried into that room: the mathematics that proves the removal works, the evidence that locates the residual to the basis point, the framing that names exactly what is at stake. And in doing so I proved that the tradition was right to reserve the verdict. Not "please bless our chain," but "here, to the

basis point, is where the line falls; rule on it.” And underneath all of it, a fallback that does not wait on the contested ruling: takaful, strong on form *and* substance, the least glamorous product in the stack turning out to be the strongest ground in it.

The honest weight of what remains is real. The grandest piece (the monetary vision, the kappa-Standard) stands squarely on the contested ground, and by its own logic on the most premium-laden part of it. It is a frontier gated on one ruling that could go against it. The seventeen papers do not dissolve that risk. They make it precise, which is the most a body of work can honestly do.

The first technical phase is complete: the mathematics identifies the removable parameter, the empirical work measures the residual, and the protocol demonstrates an implementable architecture. What remains cannot be supplied by another regression or simulation alone: a ruling, an institution, independent replication, the people willing to stake their judgment on it. I reached the boundary between what the present programme can establish and what must be decided by others, and I labelled it. Most work either never reaches that boundary or pretends it isn't there.

The number that wasn't there turned out to be optional. What I found in its place was real, and measured, and removable. And the last question, the one that matters most, I am handing, in the clearest form I could build, to the people whose question it always was.

wa-Llāhu a'lam.

Chapter 8

A Reader's Guide to the Seventeen Papers

What follows in Part II is the work itself, in reading order, from the perpetual-futures foundation to the separation theorem that holds it up. At the midpoint sits the κ -Framework synthesis (Paper 7), the one-document overview of the whole programme, placed just before the monetary papers it leads into. For the reader who wants a path through them:

The foundation. Paper 1 (perpetual-futures pricing) and Paper 2 (the credit equivalence: the keystone isomorphism $\kappa \leftrightarrow \lambda$). Start here; everything rests on these.

The instruments. Paper 2a (the κ -yield curve), 2b (stochastic takaful), 2c (the Islamic CDS, $\kappa(1 - \delta)$). One primitive, three markets.

The synthesis. Paper 7 (the κ -Framework), the single-document overview of the whole programme, and the natural place to start if you want the argument end to end before the full treatments.

The system. Paper 3 (the simulation framework: cadCAD, reinforcement learning, game theory, mechanism design), and Papers 8 and 9 (the κ -Standard and the κ -DSGE: money without a rate).

The jurisprudence and its roots. Paper 10 (the classical foundations: al-Ghazali, Ibn Taymiyyah) and Paper 11 (the Shariah boundary and the four product-layer conditions, now carrying the 2026 risk-premium addendum).

The frontier. Papers 12, 13, and 14 (liveness and macro robustness; oracle manipulation and the depth rule; the safe asset and the κ -ladder), the hard problems of running a measured primitive in adversarial conditions.

The capstone quartet. The newest work, and the spine of the argument as it now stands. Paper 15 *measures* the risk-premium share; Paper 16 *reduces* the entire riba verdict to that one measured quantity; Paper 17 *proves* its structural form and the separation theorem beneath everything; and Paper 11's addendum *anchors* the result in *fiqh*. If you read only four, read these.

The papers were written over roughly two years and are working drafts; they argue with each other in places and the later ones correct the earlier. I have left those seams visible. For

the record, the seams run mostly forward: Paper 16 re-prices the corporate convention at the recovery Paper 2c actually measured; Paper 17 restates Paper 2's pole discussion in the event parametrisation and reconciles the two; the 2026 addendum to Paper 11 absorbs the risk-premium result into the *fiqh* analysis. Every paper was revised most recently in July 2026, and every number in every paper carries a replication archive, published permanently at DOI 10.5281/zenodo.21290557. In this printing, each paper is preceded by a one-page editorial note stating its role in the programme, what it establishes, what later papers corrected, and what it does not establish; the notes are this volume's, not the papers'. A body of work that never revised itself would not be worth your scrutiny. And scrutiny, adversarial and unsparing, is the only thing I have ever asked of any reader of this.

Corpus index. The papers keep their original numbers (which match the SSRN drafts and the filenames); the figure in brackets is the position in this volume's reading order. The numbering skips 4–6 and uses 2a/2b/2c, so the seventeen papers run as follows.

No. (order)	Paper	No. (order)	Paper
1 (1)	Perpetual Futures	9 (9)	κ -DSGE
2 (2)	Credit Equivalence	10 (10)	Classical Foundations
2a (3)	κ -Yield Curve	11 (11)	Shariah Boundaries
2b (4)	Stochastic Takaful	12 (12)	Liveness & Macro
2c (5)	Islamic CDS	13 (13)	Warping the Glass
3 (6)	Simulation Framework	14 (14)	Safe Asset
7 (7)	κ -Framework (synthesis)	15 (15)	Event Intensity vs. Funding Carry
8 (8)	κ -Standard	16 (16)	Riba, Reduced
		17 (17)	Separation Theorem

Paper 7 is the former " κ -Framework" capstone (SSRN 6845098), here given its number and placed at the midpoint, just before the monetary papers. The numbering now skips only 4–6; from Paper 7 onward, the original number and the reading-order position coincide.

Chapter 9

Key Results at a Glance

Each load-bearing claim, in one line, with the paper that establishes it.

Result	Paper
A perpetual has a unique no-arbitrage price at $\iota = 0$ (the enabling theorem)	AHJ 2025
Perpetual price $\equiv$ Duffie–Singleton credit price; $\kappa \leftrightarrow \lambda$, $\iota \leftrightarrow r$	P2
Separation: the interest content is removable; pole displaced to $\iota = -\kappa$	P17
The Islamic CDS spread $s^* = \kappa(1 - \delta)$ contains no interest rate	P2c
κ recovered identically from 3 independent channels, spanning 4 orders of magnitude	P15
κ ordered by credit rating, distinct from, but not economically independent of, the risk-free curve	P2a, P15
Determinate interest-free monetary equilibrium; condition $\gamma_s < e^{\bar{\kappa}} - 1$; Cantillon-neutral	P8, P9
Risk-premium share of κ : $\sim 18\%$ (takaful) to the great majority (sovereign credit)	P15, P16
Structural form of the risk-premium share: $\Pi = \lambda /\alpha$	P17
Welfare cost of the missing riskless bond: $\sim 0.4\text{--}24$ bp of lifetime consumption	P14
Oracle-manipulation frontier (closed form) and the depth rule	P13
The riba verdict reduces to one measured, jurisprudentially-reserved quantity	P11, P16

Chapter 10

Glossary of Key Terms

This book crosses five fields. Below are the terms of art a reader from any one of them may meet from another. Definitions are working ones, meant to let the narrative read smoothly; the papers give the rigorous statements.

Kappa (κ), convergence intensity. The central object. The rate of a real event: a default, a loss, a perpetual price converging to its index. A hazard intensity, read from markets, never stipulated by contract. The proposed replacement for the interest rate in pricing.

The κ -family. Domain-specific event intensities sharing one mathematical role: κ^{credit} for credit events, κ^{loss} for insured losses, $\kappa^{\text{mortality}}$ for mortality, $\kappa^{\text{convergence}}$ for perpetual-price convergence. The claim throughout this book is structural equivalence, one pricing grammar, not numerical or economic identity across domains.

Iota (ι), interest parameter. The discount-rate slot in the perpetual-pricing equation; the charge for the mere passage of time. Setting $\iota = 0$ is the mathematical act of removing interest. The conventional risk-free rate r is its analogue in the credit model.

Riba al-nasiah. The prohibited charge for deferment: a predetermined increment owed for the passage of time, independent of outcome. The specific form of riba this work removes at the pricing layer.

Al-ghunm bil-ghurm. “Gain accompanies liability for loss.” The principle that entitlement to profit is tied to bearing real risk, the basis on which profit-and-loss-sharing contracts are lawful. The crux of the reserved question.

Takaful. Cooperative (mutual) insurance structured on risk-sharing rather than risk-transfer for premium. The domain where κ is most clearly realised hazard, and the work’s strongest ground.

Gharar. Excessive uncertainty or ambiguity in a contract; a distinct prohibition from riba, reserved here to scholars.

Maysir. Gambling; speculative gain detached from real economic activity. A distinct prohibition, reserved to scholars.

Qabd. Possession or constructive control of an asset before profiting from it; the hardest pillar for cash-settled synthetics.

Sukuk. Islamic financial certificates, often described as the asset-backed counterpart of bonds; the credit instruments measured in the empirical work.

Mudarabah / Musharakah. Profit-sharing and joint-venture partnership contracts: the canonical lawful forms of risk-bearing return.

Murabaha / Ijara. Cost-plus sale and leasing contracts: the structures whose *form* is interest-free even when the *pricing* is not.

Perpetual futures / funding rate. A derivative with no expiry, kept near its index by periodic “funding” payments between longs and shorts. The instrument where the missing interest term was first observed.

Hazard rate / intensity. The instantaneous rate at which a random event occurs. κ and the credit hazard λ are of this kind.

Reduced-form credit model (Duffie–Singleton). A framework that prices defaultable claims via a default *intensity* rather than a structural model of the firm. Shown here to be the same equation as the perpetual price.

Cox process (doubly-stochastic Poisson). A Poisson process whose rate is itself random; the formal object whose first jump is the convergence/default time θ .

Cox–Ingersoll–Ross (CIR). A square-root mean-reverting diffusion used to model κ over time; non-negative, with an affine term structure.

Feller condition. The inequality ($2\alpha\bar{\kappa} \geq \nu^2$) guaranteeing a CIR process stays strictly positive, the continuous-time form of κ ’s non-negativity.

Physical vs. risk-neutral measure. The “physical” measure describes events as they actually occur; the “risk-neutral” (pricing) measure is the one under which prices are expectations. The gap between κ under the two is the risk premium.

Risk-premium share (Π). The fraction of the priced intensity that is *not* realised expected loss; the quantity the *riba* question reduces to. Equals $|\lambda|/\alpha$ structurally.

Recovery rate (δ). The fraction of value recovered when a credit event occurs; appears in $\hat{\kappa} = s/(1 - \delta)$.

Separation theorem. The result that interest and risk are separable components of a no-arbitrage price, and the interest component is removable without the price collapsing.

Determinacy / Cantillon neutrality. *Determinacy*: the monetary equilibrium has a unique price level. *Cantillon neutrality*: new money issuance transfers no purchasing power to whoever receives it first.

The κ -Standard. The proposed monetary architecture in which base money is κ -priced equity rather than interest-bearing debt.

About the Authors

Shehzad Ahmed is the founder of Arcus Quant Fund and Baraka Protocol. He built the κ -rate research programme, the live κ -Chain prototype, and the empirical validation described in this book. His work sits at the intersection of quantitative finance, distributed systems, and Islamic commercial law.

Rafiqul Bhuyan is a professor of finance (Department of Accounting and Finance, Alabama A&M University), co-author of the programme, and the economist's check on its claims. His work spans asset pricing, derivatives, and Islamic finance.

The seventeen original working papers are on SSRN (author id 10250376).

The live prototype is at `kappachain.site`.

*The authors welcome the kind of scrutiny that tries to break the argument,
and, above all, the scrutiny of scholars and students of fiqh al-mu'amalat.*

Part II

The Research Corpus

Seventeen working papers, in reading order, from the foundation to the theorem.

The capstone quartet (P15 measures, P16 reduces, P17 proves, P11 anchors) is the spine.

EDITORIAL NOTE TO THIS PRINTING

Paper 1 · Perpetual Futures

Working paper, revised July 2026 · SSRN 6322778

Role in the programme. The foundation. The instrument where the missing interest term was first observed, and where the 2025 Akerer–Hugonnier–Jermann theorem is taken seriously as a pricing basis rather than a curiosity. **What it establishes.** A perpetual future admits a unique no-arbitrage price with the interest parameter at zero; the funding formula decomposes into a premium and a pure time-charge; live venues realise both designs (a fixed interest component pinned at 0.01% per 8 hours against a pure-premium formula whose funding oscillates through zero). **Corrected or extended later.** The novelty claim was scoped in July 2026 to the Islamic-finance literature; the general perpetual-pricing comparison has precedents. Paper 2 supplies the credit isomorphism that turns this observation into a primitive. **Not established here.** The Shariah status of perpetual contracts themselves; the ruling objections are quoted and reserved, not resolved.

Replication for every number in this paper: DOI 10.5281/zenodo.21290557 · github.com/shehzadahmed-xx/kappa-replication

The Interest Parameter in Perpetual Futures: Shariah Analysis and Empirical Evidence from Centralized and Decentralized Exchanges

Shehzad Ahmed* Rafiqul Bhuyan[†]

February 2026

Abstract

Purpose — This paper investigates whether perpetual futures can satisfy Islamic prohibitions on *riba* (interest), *gharar* (uncertainty), and *maysir* (speculation) — a question unaddressed by prior literature.

Design/methodology/approach — The study develops a four-category taxonomy of perpetual futures by funding-formula structure, collects 39,406 unique funding-rate intervals across four platforms over 365 days (February 2025 – February 2026), and analyses a cross-platform comparison — with a placebo design — between interest-bearing and interest-free DEX protocols. Shariah analysis applies El-Gamal’s *riba* conditions, Kamali’s standardisation criterion, and Salamon’s *maysir* test.

*Department of Finance, Independent University Bangladesh (IUB), Dhaka, Bangladesh.
2120922@iub.edu.bd

[†]Corresponding author. Department of Accounting and Finance, Alabama A&M University, Normal, AL, USA. rafiqul.bhuyan@aamu.edu

Findings — The interest parameter in CEX funding formulas is mathematically unnecessary. dYdX v4 operates with $I = 0$ and produces a funding distribution distinct from CEX platforms (Cohen’s $d = 0.782$ at native cadence, 0.703 cadence-matched; Kolmogorov–Smirnov $D = 0.506$). Hyperliquid — a DEX using the CEX formula — clusters with CEX. Formula structure rather than exchange architecture governs the *riba* dimension of the Shariah classification.

Originality/value — The first empirically validated framework demonstrating that the *riba* component of perpetual futures funding is structurally removable, correcting an error of extrapolation in Islamic finance scholarship that treated a single CEX implementation as representative of the entire instrument class. Ownership (*qabdh*) and delivery-intent are scoped as open questions for cash-settled synthetic perpetuals.

Research limitations/implications — The empirical sample covers a single 365-day period; extension to additional time windows would strengthen external validity.

Practical implications — The taxonomy provides a screening tool for Shariah boards, fund managers, and exchange designers evaluating Shariah-compliant perpetual futures instruments.

Keywords: *riba*, *gharar*, *maysir*, perpetual futures, Islamic finance, DeFi, DEX, Shariah compliance, funding rate

JEL classification: G11, G12, G13, Z12

Article classification: Research paper

1 Introduction

Perpetual futures are the defining instrument of the cryptocurrency era, with daily trading volumes exceeding \$100 billion ([Chainalysis, 2025](#); [CoinGecko, 2025](#)). Yet for 1.8 billion Muslims, this market has been off-limits because Binance and Bybit embed a fixed, predetermined interest charge — *riba* — in their funding mechanisms, prohibited under every school

of Islamic jurisprudence. Binance labels this charge `interestRate` in its public API: 0.01% per 8 hours (10.95% per year) (Musaffa Academy, 2024; AAOIFI, 2022; Zulkarnain & Syibly, 2024). The conservative consensus is correct for CEX perpetuals.

This paper provides evidence consistent with the interest component being a design choice, not a technical requirement. Premium-only DEX protocols dYdX v4 and D8X operate with $I = 0$, and Akerer et al. (2025) prove that $I = 0$ yields spot convergence. The *riba* dimension of Shariah compliance in perpetual futures is therefore a replicable design criterion, not a platform-specific accident. (Ownership and delivery-intent are separate questions, addressed in Section 6.)

A consistent error runs through the prior literature — treating the CEX perpetual as *the* perpetual, with the DEX alternative invisible despite operating at scale since 2023 (a review of 57 sources is reported in Section 3). We correct it through four contributions:

1. **Four-category taxonomy** of leveraged crypto derivatives by the interest parameter I and settlement date: CEX perpetual, Premium-only DEX, Hybrid DEX, and DEX expiring futures.
2. **Natural experiment** comparing dYdX ($I = 0$) and Hyperliquid ($I = 0.01\%/8h$) — both decentralised venues in the same market, differing in funding formula (with residual protocol-design confounds disclosed in Section 4).
3. **365-day empirical validation** across 39,406 funding rate intervals on four platforms.
4. **Scoped Shariah verdict.** The *riba* dimension is resolved structurally by the premium-only formula; ownership (*qabdh*) and delivery-intent remain open for cash-settled synthetic perpetuals and bound the affirmative case to RWA-backed and DEX expiring-future structures.

2 How Perpetual Futures Work and Where Riba Enters

2.1 The Funding Mechanism

A perpetual futures contract has no expiry date. Without scheduled settlement, the contract price would drift arbitrarily from the underlying spot price. The funding mechanism is the engineering solution: every 8 hours (on CEX) or every hour (on most DEX), the party whose position is “on the expensive side” pays the counterparty an amount proportional to the price gap. This continuous payment keeps the perpetual price tethered to spot.

The decisive question for Shariah compliance is not whether a payment occurs — it is *how the payment is calculated*.

2.2 The CEX Formula and Its Riba

Binance (2019) and Bybit both implement the same formula, originating from BitMEX’s 2016 launch (Hayes, 2016):

$$F = P + \text{clamp}(I - P, -0.05\%, +0.05\%) \quad (1)$$

where F is the funding rate, P is the premium index (perpetual price minus spot, divided by spot), and $I = 0.01\%$ per 8 hours is the **fixed interest rate**. Binance’s public API returns:

```
{  
  "interestRate": "0.00010000",  
  "lastFundingRate": "0.00005102"  
}
```

The field is named `interestRate`. The value is always 0.00010000. This is not interpretation; it is documentation. Three Shariah characteristics of I make it unambiguous *riba* (El-Gamal, 2006; Uddin, 2015):

- **Predetermined:** I is fixed at 0.01%/8h regardless of market conditions, counterparty creditworthiness, or economic outcome.

- **Independent:** I operates identically whether the market rises, falls, or is flat. No market outcome can eliminate it.
- **Labeled:** Binance itself calls it `interestRate` in its API. No scholarly inference is required.

Figure 1 illustrates the API evidence. The structural *riba* component is $I \times 3 \times 365 = 10.95\%/year$ — an interest term present in every funding interval that no trading strategy can remove within CEX architectures. (Realised net funding can be smaller, or negative, when the market premium offsets it, and short positions receive rather than pay the term; the Shariah objection attaches to the presence of the predetermined charge, not to its net cost.)




Figure 1: Binance `/fapi/v1/premiumIndex` API response. The `interestRate` field — always `0.00010000` ($0.01\%/8h$) — is the documentary source of *riba*. This is Binance’s own designation, not a Shariah inference. The structural annual burden is $0.01\% \times 3 \times 365 = 10.95\%/year$. Source: Binance public API documentation (Binance, 2019), reproduced for scholarly analysis.

2.3 The Mathematical Proof That Riba Is Unnecessary

Akerer et al. (2025) prove that a perpetual futures contract has a unique no-arbitrage price

— equal to the spot price when account rates are aligned — when $I = 0$.

Theorem ([Ackerer et al. 2025](#)). In both discrete and continuous time, with premium intensity $\kappa_t > 0$ and zero interest parameter $\iota = 0$, the no-arbitrage perpetual futures price exists and is unique. In the constant-parameter continuous case it equals

$$f_t = \frac{\kappa}{\kappa + r_b - r_a} x_t,$$

where r_b and r_a denote the rates of return on the margin (funding) account and on the underlying asset account respectively, in Ackerer et al.’s notation; the price equals the spot price x_t whenever $r_a = r_b$.

(We write I for the interest parameter as implemented in exchange funding formulas and ι for the same parameter in [Ackerer et al.](#)’s notation; $I = 0$ corresponds to $\iota = 0$.)

[He et al. \(2024\)](#) extend this framework by deriving no-arbitrage pricing bounds for perpetual futures under frictionless and frictional conditions, showing that the implied arbitrage Sharpe ratio is determined primarily by the funding-rate spread — the same I parameter that defines the *riba* floor.

Convergence conditions for inverse and quanto contracts follow the same pattern; the full table is reported in supplementary §S5. With $\iota = 0$, each formula reduces to a ratio involving only κ and the rate differential $(r_b - r_a)$; when rates are aligned, $f_t = x_t$ exactly.

2.4 The DEX Formula: Premium Only

dYdX v4 and D8X ([D8X Exchange, 2024](#)) implement the degenerate case of Equation (1) with $I = 0$:

$$F = P = \frac{P_{\text{perp}} - P_{\text{spot}}}{P_{\text{spot}}} \tag{2}$$

The funding rate is the signed market premium: longs pay shorts when the perpetual trades above spot, shorts pay longs when below. The payment is bilateral, conditional on market conditions, and can be zero or negative — none of the three *riba* characteristics are present.

The dYdX API returns no `interestRate` field (dYdX, 2024). D8X is a secondary reference (lower volume); empirical analysis uses dYdX v4 as the primary exemplar. Note that dYdX v4’s funding parameters are governance-configurable via validator votes; the $I = 0$ specification is therefore a current policy choice rather than a protocol-immutable property, and has been retained across every market since launch through the period of analysis.

2.5 Platform Taxonomy

Not all perpetual platforms are structurally equivalent. The four-category taxonomy (Table 1) classifies instruments by two dimensions: the interest parameter I and whether a settlement date exists.

Table 1: Four-category platform taxonomy. Shariah classification is determined by the interest parameter I and the presence of a settlement date, not by CEX/DEX architecture alone.

Category	Platforms	Riba source	Expiry	Shariah Status
CEX perpetual	Binance, Bybit	0.01%/8h	No	<i>Riba-bearing</i>
Prem-only DEX	dYdX v4, D8X	None	No	Halal candidate
Hybrid DEX	Hyperliquid	0.01%/8h	No	<i>Riba-bearing</i>
DEX expiring	MCDEX, Serum (historical)	None	Yes	Strongest

Two entries deserve particular comment. **Hyperliquid** is structurally a DEX — non-custodial, on-chain settlement — yet its funding formula is analogous to Binance’s (Hyperliquid, 2025): a premium component plus a fixed interest-rate term, clamped within a symmetric band (Hyperliquid pays hourly and uses an impact-price premium). Its Shariah status is therefore identical to Binance’s. This exposes the literature’s architectural error: calling something a DEX does not eliminate *riba*. Reading the formula does.

DEX expiring futures represent the theoretically strongest Shariah case, combining (i) no interest parameter (I is structurally inapplicable — no periodic funding mechanism); (ii) a fixed settlement date, addressing the *gharar* objection about perpetual indefiniteness; and (iii) smart-contract self-custody. Implementations have included MCDEX and Serum-based order books on Solana (both since inactive — Serum was succeeded by community

forks such as OpenBook after 2022); no venue currently operates expiring on-chain futures at the scale of the perpetual venues studied here, and liquidity fragmentation across expiry dates remains the principal practical obstacle. This paper focuses on the empirically active perpetuals market (Binance, Bybit, dYdX v4, Hyperliquid); DEX expiring futures await systematic empirical validation and are flagged in the conclusion as the most important direction for future market development.

2.6 Everlasting Options: A Forward-Looking Extension

Before turning to the literature, we note one forward-looking extension of the same $\iota = 0$ logic; readers focused on perpetual futures may skip to Section 3. Ackerer et al. (2025) extend their perpetual futures framework to derive pricing for **everlasting options** — perpetual payoff instruments that carry the right (but not the obligation) to exercise a defined payoff at any point, with no fixed expiry date. Like perpetual futures, everlasting options are sustained by a periodic premium paid by one party to the other, eliminating the need for a settlement date while preserving a well-defined payoff structure.

The premium-only pricing rule ($\iota = 0$) applies directly. Under the Ackerer et al. framework, the no-arbitrage price of an everlasting option is $f_{o,t} = E[\phi(x_{\theta_t})]$, where $\phi(\cdot)$ is the payoff function, x_{θ_t} is the underlying asset price at the random stopping time θ_t — exponentially distributed with the funding intensity in the constant-parameter case — and the expectation is taken under the risk-neutral measure. When $\iota = 0$, no interest enters the pricing equation: the premium paid between counterparties is purely market-driven, determined by the option’s time value and the underlying’s dynamics rather than any fixed financing charge.

From a Shariah perspective, everlasting options on RWA-backed asset tokens possess several potentially favourable properties. First, the payoff is *defined* — the option buyer’s maximum loss is the premium paid, and the payoff function ϕ is contractually specified. This bounded, asymmetric risk profile addresses *gharar* concerns more directly than perpetual futures, where both gain and loss are unbounded. Second, when priced with $\iota = 0$, no

fixed predetermined charge exists, satisfying El-Gamal’s *riba* conditions. Third, everlasting options on commodity tokens (gold, silver, Shariah-compliant commodity indices) could enable halal hedging with more defined risk profiles than perpetual futures — the buyer knows at entry the maximum loss, while the seller is compensated by market premium, not interest.

This is explicitly a **forward-looking extension**. No *fatwa* or formal Shariah ruling currently addresses everlasting options. The standard prohibitions on conventional options — rooted in the uncertainty of the premium and the sale of a non-existent asset — may require re-examination when applied to the everlasting structure, where the premium is market-determined and no principal is transferred without delivery of the contractual payoff. Dedicated jurisprudential research is warranted before any institutional implementation. We flag this as the most theoretically significant near-term extension of the halal perpetuals research agenda, alongside DEX expiring futures (Section 8).

3 Literature Review

No paper in the prior *Islamic-finance* literature compares dYdX v4’s premium-only formula to the CEX interest-bearing formula, or draws the Shariah consequence of the difference — the central gap addressed here. (The market-microstructure literature does document funding-formula variation across venue types (He et al., 2024; Chen et al., 2024); what is absent there is the interest-parameter taxonomy and its *riba* implication, and what is absent in the Islamic-finance literature is the DEX formula itself.) This finding is uniform across the 57 sources reviewed, covering perpetual-futures mechanics (Hayes, 2016; Shiller, 1993; Akerer et al., 2025; He et al., 2024; Chen et al., 2024; Capponi & Jia, 2024; Kim & Park, 2025), *riba* (El-Gamal, 2006; Kamali, 2000; Ayub, 2007; Jobst, 2007; Solé & Jobst, 2012; Malkawi, 2014; Tashfeen, 2022; Uluyol, 2024), *gharar* (El-Gamal, 2001; Alias, 2015; Elhessi, 2018; Kasri, 2016; Alias & Othman, 2024), *maysir* (Salamon, 2015; Alias, 2021, 2023; Obaidullah, 2005), and contemporary *fatwas* and institutional standards (Musaffa Academy,

2024; SeekersGuidance, 2025; AAOIFI, 2022, 2015; DSN-MUI, 2021; Amin, 2024; Zulkarnain & Syibly, 2024).

3.1 The Progressive-Theoretical Tradition

The progressive tradition assembled, over two decades, frameworks that collectively imply premium-only DEX permissibility.

On *riba*, El-Gamal (2006) defines it as a predetermined excess independent of economic outcome; CEX funding satisfies both conditions, dYdX v4 ($I = 0$) satisfies neither. Solé & Jobst (2012) argue for the absence of a predetermined return as a primary Shariah condition for compliant derivatives (they did not consider perpetual structures, which post-date their paper). Jobst (2007) independently proposed *riba*-free derivative structures; dYdX v4 is the empirical realisation. Ayub (2007) contributes a three-part permissibility test (no *riba*, no excessive *gharar*, legitimate function): premium-only DEX perpetuals satisfy the *riba* count structurally; the remaining counts are use- and underlying-contingent (Sections 5 and 6). Kim & Park (2025) derive optimal funding-rate designs via path-dependent BSDEs, treating the interest parameter as a designable variable rather than a fixed constraint.

On *gharar*, El-Gamal (2001) reframes it as exploitative uncertainty rooted in information asymmetry — zero in DEX protocols where every participant observes the same on-chain feed and funding formula. Kamali (2000) holds that standardisation resolves *gharar*; DEX smart contracts are more standardised than any regulated exchange. Alias & Othman (2024) develop a five-dimension *gharar* framework but apply it only to CEX platforms; on DEX, every dimension yields a different verdict. On *maysir*, Salamon (2015) specifies five criteria; DEX perpetuals used for hedging satisfy none. Tashfeen (2022) explicitly calls for *riba*-free hedging instruments; dYdX v4 provides the first live demonstration.

3.2 The Conservative-Prohibitive Tradition

The conservative tradition is internally consistent and correct for CEX perpetuals. Obaidullah (2005) provides the foundational position that all interest-bearing components constitute

riba; Alias (2015) identifies information asymmetry as operative gharar; Kasri (2016) shows regulatory frameworks systematically fail to eliminate it. The instructive cases are post-dYdX papers reaching universal prohibition without awareness of the new instrument: Alias (2023) updates the prohibitive verdict in dYdX’s launch year without mentioning it; Musaffa Academy (2024) reaches prohibition one year later, correctly identifying CEX riba but generalising to all perpetuals; AAOIFI (2022) and Zulkarnain & Syibly (2024) confirm the CEX diagnosis and apply it to the instrument class — the categorical extrapolation error.

3.3 SeekersGuidance (2025) and the Synthesis

SeekersGuidance’s 2025 ruling (2025), answered by Shaykh Muhammad Carr and approved by Shaykh Faraz Rabbani, directly poses the question addressed here: “*Are crypto perpetual contracts still haram if there’s no interest?*” The ruling maintains they remain impermissible because the trader still holds a position in an asset he does not own (*qabdh*). Once interest is set aside, the conservative objection shifts to ownership: *riba* is treated structurally in Section 5, *qabdh* in Section 6. DSN-MUI Fatwa 144 (DSN-MUI, 2021) classifies crypto assets as Shariah-legitimate commodities, providing institutional grounding for the RWA-backed case.

The two traditions converge on a single literature gap. The progressive tradition specified what compliant derivatives would require; the conservative tradition diagnosed CEX perpetuals as non-compliant. Neither identified that $I = 0$ was already operationally available. Chen et al. (2024) provide the missing microstructural foundation by showing that VAMM-based and oracle-pricing DEX architectures produce distinct trader behaviour; this paper makes the empirical extension to Shariah-compliant funding-rate design. Recent design-native scholarship is consistent with that direction (Ghazali et al., 2025; Alsmadi et al., 2025; S&P Global, 2025), alongside surveys of Islamic-fintech and DeFi (Hasan et al., 2025; Siyaddis et al., 2025; Al-Samhi & Al-Malki, 2024; Sarea & Grassa, 2024; Rahman et al., 2024; Naqshbandi & Shibly, 2024; Mirza et al., 2025), crypto-asset compatibility (Khalid et al., 2024; Islamic Development Bank, 2024), derivatives in Islamic capital markets (Zainudin et

al., 2024), and Shariah-compliant firms in digital settings (Farhat, 2025; Modjo et al., 2025).

4 Empirical Evidence: A Natural Experiment

4.1 Design and Data

Data were collected for 365 days (February 2025 – February 2026) from four platforms, yielding 39,406 unique intervals (Table 2). The raw dYdX v4 API export contains 88 exact duplicate records per asset, introduced by pagination overlap at page boundaries; all counts and statistics reported here use the deduplicated series. The dataset constitutes a **quasi-experimental comparison** (platforms differ in clientele, leverage, fees, and liquidity alongside the funding formula, so we claim consistency rather than identification): within the DEX category, dYdX v4 ($I = 0$) and Hyperliquid ($I = 0.01\%/8h$) operate in the same market over the same period with different values of I . Several confounds are acknowledged. The CEX-vs-DEX comparison is affected by sampling-cadence asymmetry (8-hour vs hourly) and by different fee-revenue and trader-demographic conditions. The dYdX-vs-Hyperliquid comparison is cleaner on those dimensions but carries protocol-design confounds: Hyperliquid uses an impact-price premium and an HLP house market-maker vault, while dYdX v4 uses an oracle-pricing premium with an external market-maker model. We therefore report a within-DEX distributional comparison whose primary differentiator is the interest parameter, not a controlled experiment. If the interest parameter matters, the two DEXs should produce different distributions; if architecture matters, similar ones.

Data were collected via public APIs at each platform’s native interval: 8-hourly for CEX; hourly for dYdX v4 and Hyperliquid. Figure 2 shows the complete 365-day time series.

A scope limitation should be noted explicitly. The empirical evidence in this paper is drawn entirely from cash-settled BTC and ETH perpetuals on dYdX v4 — the case the Shariah analysis treats as *open* on *qabdh* and *kali bil kali* grounds (Section 6). The affirmative Shariah case the paper builds (RWA-backed perpetuals; DEX expiring futures) does not yet have comparable 365-day operational data. Our empirical claim is therefore

limited to the *riba* dimension — where formula structure determines distributional outcome — and does not extend to the broader Shariah verdict for the affirmative categories. We treat the empirical-vs-affirmative-case asymmetry as a stated limitation; closing it through 365-day datasets on RWA-backed and DEX expiring contracts is the most important direction for future empirical work.

Table 2: Funding rate descriptive statistics, 365-day dataset (February 2025 – February 2026). CEX figures are per 8-hour interval. dYdX v4 and Hyperliquid figures are per native 1-hour interval. Annualised comparisons: CEX mean $\approx 0.0039\% \times 3 \times 365 \approx 4.3\%$ /year; dYdX v4 mean $\approx 0.0008\% \times 24 \times 365 \approx 7.0\%$ /year. Shariah classification is determined by the interest parameter I , not by CEX/DEX architecture.

Platform	Asset	Category	N	Mean	Std	Min	Max	Neg.%
Binance	BTC	CEX	1,095	0.0041%/8h	0.0043%	−0.015%	+0.010%	15.7%
Binance	ETH	CEX	1,095	0.0037%/8h	0.0050%	−0.037%	+0.010%	19.7%
Bybit	BTC	CEX	1,094	0.0040%/8h	0.0050%	−0.029%	+0.011%	21.6%
Bybit	ETH	CEX	1,094	0.0038%/8h	0.0067%	−0.118% ^a	+0.010%	25.8%
<i>CEX pooled</i>			<i>4,378</i>	<i>0.0039%/8h</i>				<i>20.7%^b</i>
dYdX v4	BTC	Prem-DEX ($I = 0$)	8,754	0.0006%/1h	0.0019%	−0.013%	+0.054%	30.7%
dYdX v4	ETH	Prem-DEX ($I = 0$)	8,754	0.0010%/1h	0.0019%	−0.014%	+0.022%	20.5%
<i>dYdX pooled</i>			<i>17,508</i>	<i>0.0008%/1h</i>				<i>25.6%^b</i>
Hyperliquid	BTC	Hybrid-DEX ($I > 0$)	8,760	0.0011%/1h	0.0013%	−0.009%	+0.023%	11.8%
Hyperliquid	ETH	Hybrid-DEX ($I > 0$)	8,760	0.0009%/1h	0.0015%	−0.039%	+0.015%	18.2%
<i>Hyp. pooled</i>			<i>17,520</i>	<i>0.0010%/1h</i>				<i>15.0%</i>

^aSingle-event extreme during an acute market backwardation episode (intra-sample; full date and supporting data in the supplementary materials and Zenodo dataset). Excluding this observation does not qualitatively alter the results. N (total) = 39,406.

^bPooled CEX mean and Neg. pct are simple averages across the four CEX rows (sample sizes are nearly equal: $n = 1,094$ – $1,095$); pooled dYdX figures average across the two dYdX rows ($n = 8,754$ each, after removal of 88 exact duplicate records per asset introduced by API pagination overlap).

4.2 The Natural Experiment: Results

4.2.1 The Primary Test: CEX vs. Premium-only DEX

Table 3 presents three tests comparing CEX (Binance + Bybit) against dYdX v4 under the null that the two distributions are drawn from the same population. Welch’s t , Mann–Whitney U , and Kolmogorov–Smirnov all reject decisively.



365-day dataset | CEX interest: 0.01%/8h = 10.95%/year | n(CEX)=4,378 n(dYdX)=17,684 n(Hyperliquid)=17,520

Figure 2: 365-day funding rate time series: CEX (Binance, Bybit) cluster above zero, anchored by I ; dYdX v4 oscillates around zero; Hyperliquid mirrors CEX despite DEX architecture, confirming formula determines outcome. Source: authors' calculations from the Zenodo dataset cited under Data Availability.

Table 3: Statistical tests: CEX ($n = 4,378$ per 8h) vs. dYdX v4 ($n = 17,508$ per hour). Hyperliquid is excluded from this comparison as it uses $I = 0.01\%/8h$. Statistics are computed on per-interval funding rates at each venue’s native cadence (8h vs. 1h) and therefore measure joint formula-plus-cadence distributional separation; cadence-matched statistics (dYdX aggregated to 8h: Welch’s $t = 35.69$, Cohen’s $d = 0.703$) are reported in supplementary §S2.

Test	Statistic	p -value
Welch’s t -test ¹	$t = 38.25$	$< 10^{-277}$
Mann–Whitney U	$U = 56,395,039$	$< 10^{-300}$
Kolmogorov–Smirnov ²	$D = 0.506$	$< 10^{-300}$
Cohen’s d (effect size) ³	$d = 0.782$	Medium–Large

¹Welch’s t -test (unequal variances) reported given unequal sample sizes (CEX $n = 4,378$ vs dYdX $n = 17,508$); Mann–Whitney and KS confirm.

²KS p -value at computational lower bound; scipy reports $p < 6.4 \times 10^{-306}$.

³Computed with the unweighted root-mean-square of the two sample variances, the conservative choice under the pronounced variance heterogeneity here (CEX variance $\approx 7.7\times$ dYdX); the n -weighted pooled-SD convention yields $d = 1.07$. The cadence-matched $d = 0.703$ in supplementary §S2 uses the pooled-SD convention.

The KS statistic ($D = 0.506$) indicates that at the point of maximal separation the two empirical CDFs differ by 50.6 percentage points; Cohen’s $d = 0.782$ sits at the boundary of medium and large under Cohen’s 1988 convention, persisting across the full sample window.

The interest floor, measured directly. The cleanest signature of the interest parameter is not the mean but how often the rate rests *exactly* at the interest component. Whenever the premium lies within the clamp band of the interest rate ($P \in [I - 0.05\%, I + 0.05\%]$, i.e. $P \in [-0.04\%, +0.06\%]$ at $I = 0.01\%$), the CEX formula (1) collapses to $F = I$ — the funding pins to the interest rate itself. On the study dataset, CEX BTC/ETH funding sits at the I -floor in a substantial fraction of intervals (Binance 13.7%, Bybit 23.4%), whereas dYdX v4 — with no interest term to pin to — does so in 0%. This is the mechanism in raw form: the $I > 0$ venues cannot fund below I unless the premium falls more than the clamp width below it ($P < -0.04\%$) — producing a mass point at $F = I$ and a suppressed negative tail — while the $I = 0$ venue oscillates symmetrically through zero (negative in 25.6% of intervals, against 17.7% for Binance). A floor-binding frequency of exactly zero on the $I = 0$ venue is direct confirmation that the *riba* is in the *formula*, not the venue.

Two features warrant robustness checks: cadence asymmetry (CEX 8h vs DEX 1h) and funding-rate autocorrelation. Frequency-matching the dYdX series to 8h means yields Welch’s $t = 35.69$ and Cohen’s $d = 0.703$; a moving block-bootstrap (Politis–White $b = 50$, 2,000 replicates) yields a 95% CI on Welch’s t of $[30.3, 47.6]$. Direction, order of magnitude, and effect size survive both checks (full statistics in supplementary §S2). The extreme p -values in Table 3 should be read as lower bounds on statistical distinguishability.

4.2.2 The Control: Hyperliquid vs. dYdX

If DEX structure were the operative variable, Hyperliquid and dYdX v4 — both DEXs, same market, same period — should produce similar distributions. They do not. Pooled negative-rate frequency (BTC+ETH) is 25.6% for dYdX v4 versus 15.0% for Hyperliquid: the platform with $I = 0$ shows *higher* negative frequency, consistent with a symmetric mechanism unanchored from zero, while Hyperliquid’s interest floor creates an upward bias identical in character to CEX (BTC-only values are reported in Figure 4). Hyperliquid’s distribution clusters with CEX (both positively biased), not with dYdX v4 (symmetric). Architecture does not determine shape; the interest parameter does. [Chen et al. \(2024\)](#) document the same pattern across VAMM-based and oracle-pricing DEX architectures; [Ruan & Streltsov \(2025\)](#) provide causal evidence from the Huobi perpetual removal that perpetual futures contracts reshape market microstructure through the funding mechanism.

4.2.3 Distributional Evidence

Figure 3 compares distributions across the three observed categories (DEX expiring futures excluded for lack of 365-day data); a per-platform boxplot is provided in supplementary §S4. Figure 4 shows negative-rate frequencies: platforms with $I > 0$ are suppressed in the negative tail because the interest component offsets downward premiums.

4.3 The Structural vs. Observed Cost Distinction

A crucial Shariah distinction operates at the cost level: structural riba versus observed market premium. CEX structural riba (unavoidable) is $0.01\% \times 3 \times 365 = \mathbf{10.95\%/year}$; CEX





Figure 3: Funding rate distributions by category. CEX (red, $n = 4,378$) is skewed positive by the interest floor. dYdX v4 (green, $n = 17,508$) is centred near zero, consistent with $I = 0$. Hyperliquid (amber, $n = 17,520$) clusters with CEX despite its DEX architecture. Cohen's $d = 0.782$, $KS = 0.506$ (CEX vs. dYdX). Source: authors' calculations from the Zenodo dataset cited under Data Availability.



Higher negative-funding % (dYdX) consistent with $I=0$ premium-only mechanism; lower frequency (Hyperliquid) consistent with $I>0$ floor

Figure 4: Negative funding rate frequency by platform (BTC, 365 days). dYdX v4 ($\approx 30.7\%$, $I = 0$) shows the highest negative-rate frequency — the signature of an unconstrained symmetric mechanism. Hyperliquid (11.8% , $I > 0$) shows the lowest — the interest floor suppresses negative rates below the BTC-only CEX average (18.6% ; Binance 15.7% , Bybit 21.6%), confirming the formula dominates over architecture. Source: authors' calculations from the Zenodo dataset cited under Data Availability.

observed total (interest + premium) is $\approx 4.3\%$ /year over the 365-day sample (pooled mean, Table 2), below 10.95% because negative premiums during downward-premium episodes within the sample partially offset the interest term — but the *riba* component is present in every interval. dYdX v4 structural *riba* is 0% /year; observed total $\approx 7.0\%$ /year during the 2025 bull period [1]. This exceeds CEX’s observed total, but every dollar is bilateral market transfer between traders, not interest. The Shariah position of a Muslim paying 7.0% in halal premium is categorically different from paying 4.3% containing inextricable *riba*.

The Shariah verdict therefore does not depend on which number is larger: it depends on whether a fixed predetermined charge is present at all (graphical comparison in supplementary §S4). One residual conservative objection deserves direct address: cross-venue arbitrage might cause the premium on a premium-only DEX to partially internalise the risk-free rate, so $I = 0$ may not eliminate implicit cost-of-carry economics, only relabel them. The Shariah test under El-Gamal (2006) and Solé & Jobst (2012) is contractual structure (presence of a predetermined excess), not equilibrium economics; the dYdX v4 equilibrium premium is in fact negative in 30.7% of BTC intervals, which pure cost-of-carry internalisation does not predict; and any residual implicit-rate component is bilateral between traders, which the classical *hamish jiddiyyah* structure permits.

5 Shariah Analysis

This study applies the three canonical Islamic finance prohibitions across all four platform categories. Table 4 summarises the verdicts. The *riba* verdict follows from reading the funding formula directly; the empirical evidence in Section 4 illustrates the distributional consequences of that formula choice rather than constituting a free-standing proof. Ownership (*qabdh*) and delivery intent apply at the level of the specific underlying rather than the funding formula, and are treated in Section 6 (Objections 4 and 5). Primary analysis focuses on Premium-only DEX perpetuals (dYdX v4, $I = 0$) as the empirically validated category; DEX expiring futures are discussed as a theoretically superior alternative pending liquidity

development.

Table 4: Shariah verdicts by category and prohibition.

Category	Riba	Gharar	Maysir	Overall
CEX perpetual	Haram	Concern	Struct.	Haram
Hybrid DEX	Haram	Reduced	Struct.	Haram
Prem-only DEX	Clear	Reduced	Low	Candidate
DEX expiring	N/A	Minimal	Low+	Strongest

Struct. = structural maysir pressure from riba floor (leverage + carry cost $\Rightarrow$ speculative incentive).

Low/Low+ = structural maysir pressure absent; verdict depends on trader intent and documentation.

Overall verdict for the Prem-only DEX row refers to the three canonical prohibitions; the ownership (*qabdh*) and delivery-intent considerations bound the “Candidate” verdict and are addressed in Section 6 (Objections 4 and 5).

5.1 Riba: The Decisive Prohibition

Riba is the decisive prohibition because it operates structurally, not conditionally. [El-Gamal \(2006\)](#)’s operational reformulation of the classical *riba al-nasi’ah* doctrine identifies two conditions for an excess to constitute *riba*: the excess is predetermined, and it operates independently of any economic outcome (this is El-Gamal’s economic operationalisation; the classical juristic definition additionally requires deferred exchange of *ribawi* counter-values, see also [Chapra 1985](#)). CEX interest ($I = 0.01\%/8h$) satisfies both operational conditions. The dYdX v4 premium satisfies neither: it is set at the time of payment by the perpetual-vs-spot price gap, and can be zero, positive, or negative.

[Solé & Jobst \(2012\)](#) stress the absence of a predetermined return as the primary Shariah condition for compliant derivatives, a principle the dYdX v4 formula satisfies. Notably, the [SeekersGuidance \(2025\)](#) ruling holds perpetuums impermissible on ownership rather than *riba* grounds: it sets *riba* aside as moot once *qabdh* is found wanting, rather than affirmatively conceding the *riba* point. We address the ownership question in Section 6.

5.2 Gharar: Substantially Reduced

The *gharar* prohibition derives from the hadith in *Sahih Muslim, Kitab al-Buyu’* (no. 1513 in the modern Dar al-Salam numbering) forbidding the sale of *gharar*, operationalised by [El-Gamal \(2001\)](#) as exploitative uncertainty — one party exploiting superior information

against another (see also [Alias, 2015](#); [Alias & Othman, 2024](#); [Kamali, 2000](#)). The question is whether any uncertainty is of the prohibited exploitative type.

A specific subcategory relevant to perpetuals is *bay' ma'dum* — the sale of a non-existent or undeliverable asset (*Sunan al-Tirmidhi, Kitab al-Buyu'*, no. 1232). Cash-settled CEX perpetuals raise this concern: the trader never possesses the underlying, and settlement is in a numeraire disconnected from physical delivery. DEX perpetuals on RWA-backed tokens address it by tying the position to a genuinely existent, deliverable underlying.

DEX smart contracts address *gharar* more completely than unregulated or offshore derivatives platforms: contractual terms are published in immutable on-chain code, publicly auditable; price discovery occurs through public order books or AMM formulas; loss is bounded automatically by liquidation at the margin threshold, eliminating the unbounded-loss *gharar* of uncollateralised OTC contracts ([Obaidullah, 2005](#)). DEX protocols achieve zero information asymmetry by design, and [Kamali \(2000\)](#)'s standardisation criterion is satisfied more strictly by smart-contract rules than by any regulatory standard — mathematically enforced rather than institutionally promised.

We conclude that *gharar* is substantially reduced in DEX perpetuals relative to CEX and OTC alternatives when assessed against the El-Gamal and Kamali frameworks; whether the residual uncertainty falls below the prohibition threshold is a determination for qualified Shariah scholars, to which this structural analysis is input. The residual technical concern — oracle risk in price-feed computation — warrants separate Shariah analysis as the oracle ecosystem matures.

5.3 *Maysir: Conditional on Use*

The *maysir* prohibition is grounded in Qur'an 5:90: classical jurists interpreted it as any transaction in which wealth transfers through pure chance without productive economic function. In a derivatives context this requires zero productive function, pure chance determination, and one-sided wealth transfer ([Salamon, 2015](#); [Alias, 2021](#)). The prohibition is use-contingent rather than structural: the same instrument can constitute *maysir* in specu-

lation and legitimate commerce in hedging.

Applied across categories: **CEX and Hybrid DEX perpetuals** embed a structural interest component of 10.95%/year of notional — at conservative $3\times$ leverage, an interest term of 32.85%/year of equity embedded in the funding formula — structural pressure toward longer holds and directional speculation. **Premium-only DEX perpetuals** have no such floor: a hedger offsetting a BTC spot position pays only the market premium, making short-horizon hedging genuinely feasible and satisfying the productive-economic-function criterion; the bilateral transfer structure mirrors established markets not classified as *maysir*. **DEX expiring futures** provide the strongest *maysir* distinction: fixed settlement dates require speculative positions to be re-established at each expiry, providing transaction-level evidence of intentionality (*niyyah*) absent in perpetually-held positions.

The zero-sum critique raised by [AAOIFI \(2022\)](#) applies equally to insurance contracts and secondary-market equity trading, neither of which classical scholars prohibit as *maysir*. Salamon’s *maysir* test — productive economic function, skill vs. chance, bilateral transfer, documented intent, absence of excessive leverage — yields the same pattern: CEX perpetuals score poorly on structural grounds; premium-only DEX perpetuals on economic-function and bilateral-transfer grounds; DEX expiring futures on documented-intent grounds.

6 Conservative Objections and Rebuttals

We address six objections systematically. Where possible, we cite the objection’s own scholarship to demonstrate that the concern is addressed rather than dismissed.

6.1 *Objection 1: Riba Persists Through Leverage and Implicit Credit*

“Even without the explicit interest parameter, leverage implies a loan that bears implicit interest.” — SeekersGuidance ([SeekersGuidance, 2025](#)).

Response. In DEX perpetuals, no loan is extended. The trader deposits collateral (margin); the protocol does not advance credit. This is structurally equivalent to the clas-

sical *hamish jiddiyyah* (security deposit), not a loan. Premium payments are bilateral market transfers between counterparties — not financing charges from a lender (Solé & Jobst, 2012; Malkawi, 2014). Once the explicit interest parameter is removed, the implicit-interest-through-leverage argument loses its foothold.

6.2 Objection 2: Gharar From Perpetual Nature and Price Uncertainty

“The absence of a settlement date creates excessive gharar.” — Musaffa (2024); Alias (2015).

Response. Perpetual nature does not constitute prohibited *gharar*. Contract terms are fully specified: funding formula hardcoded, margin requirements known, liquidation threshold public. The absence of a settlement date is a feature, not an ambiguity — and the absence of expiry removes the *bay’ al-kali bil kali* gharar that classical scholars identified. Automated liquidation bounds loss to deposited margin (Kamali, 2000; Elhessi, 2018).

6.3 Objection 3: Maysir Through Zero-Sum Structure

“One party wins only because another loses — this is gambling.” — AAOIFI (2022); Alias (2021).

Response. Every secondary-market transaction has buyer and seller; every insurance contract has premium-payer and beneficiary. If zero-sum payoff structure constitutes *maysir*, all secondary-market activity is prohibited — a conclusion no classical scholar has reached. The operative test is *economic utility*, not aggregate payoff symmetry (Ayub, 2007; El-Gamal, 2006). Derivatives markets transfer unwanted risk from hedgers to willing risk-takers; Kamali (2000) and Ayub (2007) permit instruments serving this function.

6.4 Objection 4: Non-Ownership and Qabdh (Possession) Requirement

“Do not sell what you do not possess” (*Sunan al-Tirmidhi, Kitab al-Buyu’*, no. 1232) — SeekersGuidance (SeekersGuidance, 2025).

Response. Once the *riba* parameter is neutralised by a premium-only formula, *qabdh* becomes the binding constraint. We treat it in two layers: contract classes where the Shariah

case is strong, and the empirical exemplar where it is partial.

Strong cases (RWA-backed perpetuals; DEX expiring futures). DEX perpetuals written on tokenised gold or commodity claims represent a verified legal claim on a physical asset held in custody, satisfying the ownership requirement directly; [DSN-MUI \(2021\)](#) provides the institutional foundation. (Tokenised-gold cases raise a separate *sarf taqabud* constraint, outside the scope of this paper.) Classical precedent permits *salam* and *istisna*, both involving payment before delivery; the *qabdh* requirement is not absolute, and DEX expiring futures are structurally analogous to *salam*. AAOIFI’s communications on digital assets ([AAOIFI, 2022](#)) and the growing *fiqh* literature on constructive *qabdh* (exclusive, transferable on-chain entitlements) increasingly support treating DEX token positions as satisfying ownership requirements.

Hard case (cash-settled BTC and ETH perpetuals). The trader’s position is not possession of the underlying; it is an entitlement to a cash-settled price claim against the protocol. The *salam* analogy is imperfect, since *salam* contemplates delivery of a specific real commodity. Two partial defences exist: the constructive-*qabdh* literature noted above, which treats on-chain entitlements as an electronic possession analogue ([AAOIFI, 2022](#)); and the narrower scoping claim that the Shariah-compliant application of the premium-only formula is most defensible on RWA-backed and DEX expiring-future cases. The cash-settled synthetic case is treated as open, not as resolved by this paper.

6.5 Objection 5: Usmani’s Twin Objections and OIC Fiqh Academy Resolution 63

Usmani ([2015](#)) raises two distinct arguments against futures, both of which apply with full force to cash-settled synthetic perpetuals and warrant separate treatment.

Objection 5a (delivery intent). Futures are not Shariah-compatible because parties typically do not intend delivery; the exercise collapses into an exchange of price differentials rather than a transfer of the underlying. OIC Fiqh Academy Resolution 63 ([OIC Fiqh Academy, 1992](#)) reaches a compatible conclusion on *gharar* and gambling grounds, holding that standard futures and options transactions are impermissible because the subject of the

contract is neither a sum of money, a utility, nor a financial right that can be waived.

Response (5a). The delivery-intent objection is correctly directed at contracts used purely for speculative price-differential settlement and applies with equal force to the CEX case analysed in this paper. For hedging-motivated use on RWA-backed perpetuals or DEX expiring-future instruments, delivery intent is present by construction at the expiration of the contract. The premium-only cash-settled perpetual on BTC or ETH does not in general involve delivery intent, which reinforces the scoping set out in Objection 4: the paper’s affirmative case runs through RWA-backed and DEX expiring structures.

Objection 5b (bay’ al-kali bil kali). A more fundamental argument: cash-settled perpetuals reduce to the exchange of one deferred payment claim for another, which the classical jurists classified as *bay’ al-kali bil kali* (sale of debt for debt). This is prohibited under a hadith reported in *Sunan al-Daraqutni*; the hadith’s *isnad* is contested (some classical muhad-dithun grade it weak), but the prohibition is widely accepted across the major schools on the basis of consensus (*ijma’*) and analogical reasoning (Usmani, 2015, 2002; AAOIFI, 2022), although Ibn Taymiyyah and Ibn Qayyim narrow the *ijma’* claim to specific debt-for-debt configurations (Ibn Taymiyyah, 1995). Each side of a perpetual position is, in this reading, a claim against the future settlement flow rather than a sale of an existent commodity.

Response (5b). The *kali bil kali* argument is the sharpest objection in the conservative literature for the cash-settled case. Three observations limit its force. First, the classical *kali bil kali* prohibition contemplates a debt-for-debt exchange where neither side possesses consideration; a DEX perpetual position is collateralised against deposited margin (the trader’s own funds), so neither side is purely deferred in the classical sense. Second, the funding mechanism is a continuous bilateral transfer rather than deferred settlement of a fixed notional; on a premium-only formula, no party owes a predetermined sum at any future date. Third, the objection retains its force for any position taken without a hedging rationale, irrespective of formula. We nevertheless read *kali bil kali* as a genuine residual constraint on cash-settled synthetic perpetuals that the affirmative RWA-backed and DEX expiring cases

do not encounter, consistent with the bounded scope: *riba* resolved structurally; ownership, delivery-intent, and *kali bil kali* open for cash-settled synthetic structures.

6.6 Objection 6: Societal Harm Through Excessive Speculation

“High-leverage derivatives cause financial instability and individual ruin.” — DSN-MUI (2021); Kasri (2016).

Response. This objection applies to *irresponsible use*, not to the instrument itself. The same critique applies to margin lending in equities, FX trading, and leveraged property investment — none categorically prohibited. DEX perpetuals used with disciplined position sizing ($\leq 3\times$ leverage), algorithmic stop-losses, and productive underlying assets are no more harmful than established leveraged instruments. The *maqasid al-Shariah* objective of protecting wealth (*hifz al-mal*) is served by appropriate risk management, not categorical exclusion.

7 Implementation Framework for Muslim Traders

This framework operationalises the paper’s findings as a screening tool. The formula test below settles the *riba* dimension; asset selection and intent documentation address the ownership (*qabdh*) and *maysir* dimensions that Section 6 leaves contingent on the underlying and the trader.

7.1 Platform Selection: The Formula Test

The single most important selection criterion is the interest parameter I in the funding formula, documented in each platform’s official API and technical documentation. Do not rely on marketing descriptions such as “decentralised” or “non-custodial.”

Riba-free formula (passes the formula test): dYdX v4 ($I=0$, documented (dYdX, 2024)) and D8X ($I=0$, documented (D8X Exchange, 2024)).

Riba-bearing despite DEX architecture: Hyperliquid ($I=0.01\%/8h$, documented (Hyperliquid, 2025)). Using Hyperliquid subjects a Muslim trader to the identical *riba* burden

as using Binance.

As new DEX platforms emerge, verify the funding formula documentation before participation. Smart contract governance may alter parameters over time; continuous monitoring is required.

7.2 Asset Selection: Productive Underlying

Prioritise perpetuals on assets with demonstrated economic utility:

- **Commodity-backed tokens** (gold, silver, oil tokens) — directly relevant to *qabdh* and *maqasid* arguments.
- **Infrastructure tokens** with significant productive usage (network fee revenue, validator staking).
- **Tokenised real-world assets** (RWA tokens representing verified physical assets).

Avoid perpetuals on purely speculative or utility-less tokens. Asset selection strengthens the *maysir* analysis by establishing genuine economic purpose (Kalimullina, 2020). The leading near-term exemplars are perpetuals on vault-secured tokenised gold (inflation and currency-risk hedging for businesses in volatile-currency jurisdictions) and on tokenised sukuk yields (balance-sheet liquidity management without interest-rate swaps) — risk-transfer applications in which the position offsets a real exposure, consistent with the asset-backing, possession, insurable-interest, and delivery conditions developed in companion work; neither instrument is claimed to exist or to be certified today.

Note on everlasting options: For certain hedging applications — particularly those requiring a defined maximum loss rather than an uncapped directional position — everlasting options on RWA-backed commodity tokens (Section 2.6) may offer superior Shariah defensibility compared to perpetual futures. The bounded payoff structure addresses *gharar* concerns more directly, and pricing without interest ($\iota = 0$) preserves the structural *riba*-free character. Pending dedicated scholarly review and formal *fatwa*, institutions with more conservative Shariah requirements should note this instrument class as a potentially stronger candidate

for board approval than outright perpetual futures positions.

7.3 Risk Management: Parameters

Implement disciplined position management consistent with *hifz al-mal*:

- Maximum leverage: 2–3 $\times$ (not the 50–100 $\times$ offered by most platforms).
- Mandatory stop-loss: maximum 10–15% drawdown from entry before automatic exit.
- Position sizing: maximum 5% of portfolio in any single perpetual position.
- Regular rebalancing: review position rationale monthly.

7.4 Intent Documentation (*Niyyah*)

Document trading intent before opening positions: the economic rationale (hedging, market-making, investment), the target asset, the risk parameters, and the exit criteria. This serves both Shariah discipline and practical risk management. A position opened without articulated rationale is more difficult to distinguish from speculative gambling (*maysir*) in retrospect.

7.5 Scholarly Consultation

This paper provides the structural analysis; it does not constitute a *fatwa*. Muslim traders and institutions should engage qualified Shariah scholars in their jurisdiction with this technical analysis as input. We particularly encourage engagement with AAOIFI and OIC Fiqh Academy on formal standard-setting for DEX derivative instruments.

8 Conclusion

The Islamic finance consensus that perpetual futures are prohibited is correct — for CEX perpetuals. The error across 57 reviewed sources is one of extrapolation: a correct *riba* diagnosis in one implementation was generalised to all implementations. That generalisation was defensible as long as all implementations shared the same formula; it is no longer defensible.

The interest parameter I in CEX funding formulas is demonstrable *riba* — predeter-

mined, independent of economic outcome, and labelled as interest in the platform’s own API. That verdict stands for all platforms with $I > 0$ (Binance, Bybit, Hyperliquid). The interest parameter is, however, mathematically unnecessary: [Ackerer et al. \(2025\)](#) prove this, and dYdX v4 and D8X have operated successfully with $I = 0$ at scale. 365 days and 39,406 funding intervals confirm that the interest parameter — not exchange architecture — determines distributional structure: Hyperliquid (DEX, $I > 0$) behaves like CEX; dYdX v4 (DEX, $I = 0$) behaves differently. The Shariah verdict follows from reading the funding formula, with the distributional evidence serving as confirmation rather than free-standing proof. Applying El-Gamal’s *riba* conditions, Kamali’s standardisation criterion, Sole & Jobst’s no-predetermined-return framework, and Salamon’s *maysir* test yields a coherent verdict on premium-only DEX perpetuals used on productive underlying assets with disciplined risk management: *riba* absent by design; *gharar* substantially reduced by on-chain transparency; *maysir* use-contingent. Ownership (*qabdh*) and delivery-intent remain open for cash-settled synthetic perpetuals.

The practical implication: structural barriers to Muslim participation in this market are not inherent to the instrument class but arise from specific funding-formula design choices. Where $I = 0$, as in dYdX v4 and D8X, the *riba* basis for prohibition is removed — a replicable design criterion, not a platform-specific accident.

The reform implication is broader: perpetual futures can be structured to remove the *riba* component by design, and the empirical evidence is consistent with premium-only ($I = 0$) funding functioning at institutional scale across the 365-day sample analysed here. The remaining Shariah dimensions — ownership (*qabdh*), delivery-intent, and product-level *gharar* — are determinations for qualified scholars, to whom this paper supplies the technical foundation. We invite Islamic scholars, Shariah boards, and standard-setting bodies — particularly AAOIFI (whose Shari’ah Standards already address derivative and commodity contracts) and the OIC Fiqh Academy — to engage with the technical specifics of DEX funding formulas and develop formula-specific guidance for institutional participants.

Limitations and future research. Four limitations bound these results. The empirical sample covers a single 365-day window; multi-cycle data would strengthen external validity. The $I = 0$ specification on premium-only DEXs is governance-configurable rather than protocol-immutable, so compliance is a monitoring requirement, not a permanent guarantee. Oracle manipulation in price-feed computation is a residual *gharar* concern distinct from the funding formula and warrants separate analysis as the oracle ecosystem matures. And no recognised authority has yet issued a formal *fatwa* on DEX perpetuals; this paper supplies the technical foundation, not the jurisprudential ruling. The most important extension is 365-day operational data on the RWA-backed and DEX expiring-future structures on which the affirmative Shariah case rests.

Funding

This research received no specific grant from any funding agency in the public, commercial, or not-for-profit sectors.

Acknowledgements

No acknowledgements are required for this work.

Declaration of generative AI use

During the preparation of this work, the authors used Claude Code (Anthropic) for copy-editing — specifically, correcting, formatting, and improving the clarity of the language and grammar of the authors’ own original text. No AI tool was used to generate or draft new content. The authors reviewed and take full responsibility for all content in the published article.

Data availability

The funding rate dataset (39,406 unique intervals across four platforms, February 2025 – February 2026) used in this study is available at: <https://doi.org/10.5281/zenodo.18817390> [dataset]. The raw dYdX v4 exports in the deposit additionally contain 88 exact duplicate records per asset from API pagination overlap; reported statistics use the deduplicated series.

Notes

[1] The 365-day window covers a predominantly bullish phase; a bear-period sample would likely show lower or negative annualised premiums on dYdX, while the CEX *riba* floor of 10.95%/year would persist unchanged.

References

- Chapra, M. U. (1985). *Towards a Just Monetary System*. The Islamic Foundation, Leicester.
- Ibn Taymiyyah, T. (1995). *Majmu' al-Fatawa*, Vol. 29. King Fahd Complex for the Printing of the Holy Qur'an, Medina.
- Usmani, M. T. (2002). *An Introduction to Islamic Finance*. Kluwer Law International, The Hague.
- AAOIFI. (2015). *Shari'ah Standards for Islamic Financial Institutions*. Accounting and Auditing Organisation for Islamic Financial Institutions. <https://www.iefpedia.com/english/wp-content/uploads/2017/12/Shariaa-Standards-ENG.pdf>
- AAOIFI. (2022). *Shari'ah Standards* (compiled edition, with AAOIFI communications on digital and crypto assets). Accounting and Auditing Organisation for Islamic Financial Institutions, Manama.

- Ackerer, D., Hugonnier, J., & Jermann, U. (2025). Perpetual futures pricing. *Mathematical Finance*. <https://doi.org/10.1111/mafi.70018>
- Al-Samhi, M. A., & Al-Malki, A. M. (2024). Crypto assets and Islamic finance: A systematic literature review. *Journal of Islamic Accounting and Business Research*, 15(4).
- Alias, M. (2015). Gharar in forward and futures contracts. *ResearchGate*. <https://www.researchgate.net/publication/273904700>
- Alias, M. (2021). Derivative instruments: Criteria for speculation in Islamic finance. *Academy of Accounting and Financial Studies Journal*, 25(S3). <https://www.abacademies.org/articles/derivative-instruments-criteria-for-speculation-in-islamic-finance-11080.html>
- Alias, M. (2023). Derivative instruments: Criteria for speculation in Islamic finance (updated). *ResearchGate*. <https://www.researchgate.net/publication/369559261>
- Alias, M., & Othman, A. (2024). A review on *gharar* dimension in modern Islamic finance transactions. *Journal of Islamic Accounting and Business Research*, 16(5), 976–994. <https://www.emerald.com/insight/content/doi/10.1108/JIABR-02-2023-0060>
- Alsmadi, A. A., et al. (2025). Beyond compliance: Islamic Fintech and CSR. *Future Business Journal*, 11(1).
- Amin, M. (2024). Are perpetual futures contracts Shariah compliant? https://www.mohammedamin.com/Islamic_finance/Perpetual-futures.html
- Ayub, M. (2007). *Understanding Islamic Finance*. John Wiley & Sons.
- Binance. (2019). Introduction to Binance Futures funding rates. *Binance Support*. <https://www.binance.com/en/support/faq/detail/360033525031>
- Capponi, A., & Jia, R. (2024). The adoption of blockchain-based decentralised exchanges. *arXiv:2103.08842* (rev. 2024). <https://doi.org/10.48550/arXiv.2103.08842>

- Chainalysis. (2025). *An Introduction to Perpetual Futures*. <https://www.chainalysis.com/blog/perpetual-futures-intro-2025/>
- Chen, E., Ma, M., & Nie, Z. (2024). Perpetual future contracts in centralised and decentralised exchanges: Mechanism and traders' behaviour. *Electronic Markets*, 34, 35. <https://doi.org/10.1007/s12525-024-00715-1>
- CoinGecko. (2025). *State of Crypto Perpetuals Market 2025*. <https://assets.coingecko.com/reports/2025/CoinGecko-State-of-Crypto-Perpetuals-Market.pdf>
- D8X Exchange. (2024). The silent difference between DeFi perpetual futures contracts. *Medium*. <https://medium.com/@d8x.exchange/the-silent-difference-between-defi-perpetual-futures-contracts-283e47a84b94>
- DSN-MUI. (2021). Fatwa No. 144/DSN-MUI/XI/2021 on crypto assets. *Indonesian Ulema Council*.
- dYdX. (2024). Perpetual funding rate. *dYdX Help Center*. <https://help.dydx.exchange/en/articles/4797443-perpetual-funding-rate>
- El-Gamal, M. A. (2001). An economic explication of the prohibition of *gharar*. *Islamic Economic Studies*, 8(2), 29–58.
- El-Gamal, M. A. (2006). *Islamic Finance: Law, Economics, and Practice*. Cambridge University Press.
- Elhessi, A. (2018). *Gharar* in futures contracts from an Islamic perspective. *Risk Governance and Control: Financial Markets and Institutions*. <https://virtusinterpress.org/Gharar-in-futures-contracts-from-Islamic-perspective-A-case-study-of-USA-Malaysia.html>
- Farhat, A. (2025). The impact of Shariah compliance on firm's performance and risk. *Review of Financial Economics*.

- Ghazali, N. M., et al. (2025). Exploring financial derivatives through Shariah principles. *Al-Qanatir: International Journal of Islamic Studies*.
- Hasan, R., et al. (2025). Smart contracts and Shariah compliance in decentralised exchanges. *FinTech*, 4(2).
- Hayes, A. (2016). Announcing the launch of the perpetual XBTUSD leveraged swap. *BitMEX Blog*. <https://www.bitmex.com/blog/announcing-the-launch-of-the-perpetual-xbtusd-leveraged-swap>
- He, S., Manela, A., Ross, O., & von Wachter, V. (2024). Fundamentals of perpetual futures. *arXiv:2212.06888* (rev. 2024). <https://doi.org/10.48550/arXiv.2212.06888>
- Hyperliquid. (2025). Funding rate mechanics. *Hyperliquid Documentation*. <https://hyperliquid.gitbook.io/hyperliquid-docs/trading/funding>
- Islamic Development Bank. (2024). *Islamic Finance and Digital Assets: Opportunities and Challenges*. Islamic Development Bank Publications.
- Jobst, A. A. (2007). Derivatives in Islamic finance. *SSRN Working Paper*. https://papers.ssrn.com/sol3/papers.cfm?abstract_id=1015615
- Kalimullina, M. (2020). Islamic finance and food security: A review of the literature. *Heliyon*, 6(11).
- Kamali, M. H. (2000). *Islamic Commercial Law: An Analysis of Futures and Options*. Islamic Texts Society.
- Kasri, N. S. (2016). *Critical Analysis of Malaysian Securities Commission SAC Resolution on Crude Palm Oil Futures*. PhD Thesis, University of Gloucestershire. <https://eprints.glos.ac.uk/3444/>
- Khalid, W., et al. (2024). Blockchain-based sukuk: Bridging traditional Islamic finance with DeFi. *International Journal of Islamic Finance*, 16(2).

- Kim, J., & Park, H. (2025). Designing funding rates for perpetual futures in cryptocurrency markets. *arXiv:2506.08573*. <https://doi.org/10.48550/arXiv.2506.08573>
- Malkawi, B. H. (2014). Financial derivatives between Western legal tradition and Islamic finance. *Journal of Banking Regulation*, 15(1).
- Mirza, N., et al. (2025). Cryptocurrency regulation and Islamic finance: A comparative study. *Emerging Markets Review*, 62.
- Modjo, M. I., et al. (2025). Weathering the storm: Shariah compliance, digital innovation, and stock performance. *Journal of Islamic Monetary Economics and Finance*, 11(1).
- Musaffa Academy. (2024). Things about perpetual futures you need to know. <https://academy.musaffa.com/things-about-perpetual-futures-you-need-to-know/>
- Naqshbandi, M., & Shibly, M. (2024). Crypto assets and Islamic commercial law: Compatibility and challenges. *International Journal of Law and Management*, 66(4).
- Obaidullah, M. (2005). *Islamic Financial Services*. King Abdulaziz University. https://www.kau.edu.sa/Files/320/Researches/50893_21007.pdf
- OIC International Islamic Fiqh Academy. (1992). Resolution No. 63 (1/7) on financial markets: Stocks, options, commodities, and credit cards. Seventh Session, Jeddah. <https://iifa-aifi.org/en/32438.html>
- Rahman, M. P., et al. (2024). Decentralised finance and Shariah compliance: Challenges and opportunities. *Journal of Financial Regulation and Compliance*, 32(3).
- Ruan, Q., & Streltsov, A. (2025). Perpetual futures contracts and cryptocurrency market quality: Insights from emerging markets. *Cornell SC Johnson College of Business Emerging Markets Institute*. Working paper.

- S&P Global Ratings. (2025). *Islamic Finance 2025–2026: Resilient Growth Amid Upcoming Headwinds*. S&P Global Ratings. <https://www.spglobal.com/ratings/en/research/articles/250417-islamic-finance-2025-2026>
- Salamon, H. (2015). Speculation: The Islamic perspective; A study on al-*maisir*. *Mediterranean Journal of Social Sciences*, 6(4). <https://www.academia.edu/25656159/>
- Sarea, A. M., & Grassa, R. (2024). The evolution of Islamic fintech: Trends and future directions. *Quarterly Review of Economics and Finance*, 93.
- SeekersGuidance. (2025, October 2). Are crypto perpetual contracts still haram if there's no interest, and the asset price tracks? Answered by Shaykh Muhammad Carr; approved by Shaykh Faraz Rabbani. <https://seekersguidance.org/answers/transactions/are-crypto-perpetual-contracts-still-haram-if-theres-no-interest-and-the-asset-price-tracks>
- Shiller, R. J. (1993). *Macro Markets: Creating Institutions for Managing Society's Largest Economic Risks*. Oxford University Press.
- Siyaddis, A., et al. (2025). Tokenisation of real-world assets: Islamic finance implications. *Journal of Islamic Monetary Economics and Finance*, 11(2).
- Solé, J., & Jobst, A. A. (2012). Operative principles of Islamic derivatives: Towards a coherent theory. *IMF Working Paper* WP/12/63. <https://www.elibrary.imf.org/view/journals/001/2012/063/article-A001-en.xml>
- Tashfeen, M. (2022). Financial derivatives hedging: A review of Islamic finance at crossroads. *Global Economics Review*, 7(2). <https://www.gerjournal.com/article/financial-derivatives-hedging-a-review-of-islamic-finance-at-cross-roads>
- Uddin, M. A. (2015). Principles of Islamic finance: Prohibition of *riba*, *gharar* and *maysir*.

MPRA Paper No. 67711. https://mpra.ub.uni-muenchen.de/67711/1/MPRA_paper_67711.pdf

Uluyol, B. (2024). Financial derivative instruments and their applications in Islamic banking and finance. *Borsa Istanbul Review*, 24(2). <https://www.sciencedirect.com/science/article/pii/S2214845024000334>

Usmani, M. T. (2015). *An Introduction to Islamic Finance*. Maktaba Ma'ariful Qur'an, Karachi.

Zainudin, M., et al. (2024). Derivatives in Islamic capital markets: Regulatory perspectives. *International Journal of Islamic and Middle Eastern Finance*, 17(1).

Zulkarnain, R., & Syibly, M. R. (2024). Analysis of futures trading on the Binance platform from the perspective of Fiqh Muamalah. *Journal of Islamic Economics Lariba*. <https://journal.uui.ac.id/JIELariba/article/download/33018/17999/132830>

EDITORIAL NOTE TO THIS PRINTING

Paper 2 · Credit Equivalence

Working paper, revised July 2026 · SSRN 6322858

Role in the programme. The keystone. The isomorphism that maps the perpetual price onto the Duffie–Singleton credit price: $\kappa \leftrightarrow \lambda, \iota \leftrightarrow r$, event time onto default time, terminal value onto recovery. **What it establishes.** The risk content and the interest content of an event-anchored no-arbitrage price are separable, and the interest content is removable. One equation, two literatures. **Corrected or extended later.** Paper 17 restates the pole discussion in the event parametrisation and reconciles which operator owns the pole. Paper 16 re-prices this paper’s corporate convention at the recovery Paper 2c actually measured. Demand-side and agri-takaful caveats were added in July 2026. **Not established here.** Whether the contingent premium that remains after removal is itself permissible; that question is defined here and reserved throughout.

Replication for every number in this paper: DOI 10.5281/zenodo.21290557 · github.com/shehzadahmed-xx/kappa-replication

From Perpetual Contracts to Islamic Credit: The Random Stopping Time Equivalence and a Riba-Free Foundation for Sukuk, Takaful, and Credit Protection

Shehzad Ahmed¹ and Rafiqul Bhuyan^{*2}

¹Department of Finance, Independent University, Bangladesh. 2120922@iub.edu.bd

²Department of Accounting and Finance, Alabama A&M University, USA. rafiqul.bhuyan@aamu.edu

February 2026 (revised June 2026)

Abstract

We establish a formal mathematical equivalence between two previously unconnected pricing frameworks — the random time representation of perpetual futures contracts [3] and the reduced-form credit risk model of Duffie & Singleton [19] — and demonstrate that this equivalence has far-reaching implications for Islamic finance.

The core observation is structural: the stopping time θ_t in the Akerer-Hugonnier-Jermann perpetual pricing formula is isomorphic to the default time τ in the Duffie-Singleton defaultable bond formula, under the identification of the convergence intensity κ with the hazard rate λ and the interest parameter ι with the risk-free discount rate r . When the interest parameter is set to zero ($\iota = 0$) — which Akerer et al. prove achieves unique no-arbitrage perpetual futures pricing — the analogous credit instrument at $r = 0$ is also no-arbitrage, with the stopping time distribution alone carrying all timing information previously encoded in the interest discount.

This equivalence supplies what Islamic finance has lacked for half a century: a mathematically rigorous, no-arbitrage pricing foundation for credit instruments — sukuk, takaful, and credit default swaps — that requires no positive interest rate. We prove the main theorem, derive closed-form prices for an Islamic perpetual sukuk and an everlasting takaful contract using Proposition 6 of Akerer et al., and show that the convergence intensity κ functions as a *riba-free* analog to the risk-free rate across all Islamic credit instruments. We term this the κ -rate. We further establish a **Separation Theorem**: at $\iota = 0$ the no-arbitrage price carries no interest term, so interest is an *eliminable* component of pricing — not an intrinsic necessity. Empirical evidence from the European negative interest rate environment (2014–2022) corroborates the theoretical finding that credit markets function without a positive discount rate. We further validate the central identification $\hat{\kappa} = s/(1 - \delta)$ on real markets: across 195 actively-quoted USD sukuk (40 sovereign, 154 corporate, one municipal; Cbonds), the implied $\hat{\kappa}$ has the predicted credit-regime magnitude (median ≈ 0.02 per annum) and broadly increases as credit quality deteriorates (apart from the Saudi-dominated AA bucket, which prices just above single-A) — so the *riba-free* intensity is a measurable property of the traded sukuk market, not only a theoretical construct.

Keywords: *riba*, Islamic finance, credit risk, sukuk, takaful, perpetual contracts, stopping time, no-arbitrage, Duffie-Singleton, κ -rate, *riba-free* pricing, blockchain sukuk

JEL Codes: G12, G13, G15, G22, Z12

1 Introduction

Pricing credit risk is arguably the central unresolved problem in Islamic finance. The \$3 trillion Islamic finance industry [32], which has grown from a niche experiment into a global system over the past half century [30], operates without any mathematical tool for pricing credit default risk that does not require a positive interest rate. Indeed, every credit-derivatives framework in use today — the

structural models of Merton [46], Black & Cox [12], and Leland [44]; the reduced-form models of Jarrow & Turnbull [33], Duffie & Singleton [19, 20], and Lando [43]; the textbook synthesis of Bielecki & Rutkowski [9] and Jeanblanc, Yor & Chesney [35] — is built on a positive risk-free interest rate $r > 0$. Islamic scholars correctly identify this r as *riba*: it is the compensation for the pure time value of money, which Islamic jurisprudence prohibits as a predetermined excess in exchange transactions [22, 55].

The consequence has been severe. Islamic institutions cannot hedge credit risk [42]. Sukuk (Islamic bonds) are priced at issuance using LIBOR/SOFR-anchored models, making the interest rate implicit even when explicit interest payments are restructured away [54, 37, 49]. Takaful (Islamic insurance) operators invest pools in sukuk priced by *riba*-based models [11]. The Islamic credit default swap does not exist. Scholars have sought workarounds for decades, but these have been cosmetic, not mathematical [40, 22].

This paper resolves the impasse. We demonstrate that a recent result from mathematical finance — the random time representation theorem for perpetual futures contracts proved by Akerer, Hugonnier & Jermann (2025) [3] — contains, as a special case, a complete and rigorous pricing framework for credit instruments that does not require $r > 0$.

The Core Observation

The Akerer-Hugonnier-Jermann theorem establishes that the unique no-arbitrage price of a perpetual futures contract on a spot asset X_t with interest parameter ι and convergence intensity κ is:

$$f_t = \mathbb{E}_t^{Q^a} \left[e^{-\int_t^{\theta_t} \iota_s ds} X_{\theta_t} \right] \quad (1)$$

where θ_t is a random stopping time with exponential density under the pricing measure Q^a . The critical result, proved in their Theorem 3, is that $\iota = 0$ satisfies all no-arbitrage conditions. When $\iota = 0$, equation (1) collapses to:

$$f_t = \mathbb{E}_t^{Q^a} [X_{\theta_t}], \quad \theta_t \sim t + \text{Exp}(\kappa) \quad (2)$$

The instrument price is the *expected spot value at a random future moment* — with the probability distribution of that moment determined entirely by κ , and no interest discount anywhere.

The Duffie-Singleton reduced-form credit model prices a defaultable instrument as:

$$P_t = \mathbb{E}_t^Q \left[e^{-\int_t^\tau r_s ds} R_\tau \right] \quad (3)$$

where τ is the default time (exponentially distributed with hazard rate λ) and R_τ is the recovery value at default.

Equations (1) and (3) are the same equation. The stopping time θ_t is τ . The convergence intensity κ is λ . The interest parameter ι is r . The spot X_{θ_t} is the recovery R_τ .

At $\iota = r = 0$: both yield $\mathbb{E}[X_\tau]$. The Akerer theorem proves this is no-arbitrage. Therefore credit pricing at $r = 0$ is no-arbitrage, provided $\kappa = \lambda > 0$ carries all timing information. *Riba* is therefore not a pricing necessity but a historical artifact of fusing timing and discounting into a single parameter.

Related Literature

This paper sits at the intersection of four bodies of work.

Reduced-form credit risk. Jarrow & Turnbull [33] introduced intensity-based default modeling; Duffie & Singleton [19] generalized the framework to affine stochastic intensities and established the risk-adjusted short-rate representation; their book [20] provides the comprehensive treatment. Lando [43] gave the definitive Cox-process (doubly stochastic Poisson) treatment. Brémaud [14] established the martingale theory for point processes underpinning survival-probability calculations. Elliott, Jeanblanc & Yor [23] resolved the filtration enlargement problem; Jeanblanc, Yor & Chesney [35] provide the definitive mathematical synthesis. Bielecki & Rutkowski [9] and Schönbucher [50] synthesized this literature into standard practitioner frameworks. Duffie & Lando [18] extended the framework to incomplete accounting information. Our contribution shows that the entire reduced-form machinery transfers to Islamic finance under $\iota = r = 0$, with κ replacing λ .

Perpetual contract theory. Akerer, Hugonnier & Jermann [3] prove that perpetual futures have unique no-arbitrage prices at $\iota = 0$ and admit the random time representation (Theorem 3) central to our argument. Merton [46] and Leland [44] priced perpetual-maturity corporate debt using structural first-passage models; Black & Cox [12] introduced the stopping-time formalism in that context. Shiller [52] proposed macro perpetuals conceptually. None connected the stopping-time structure to Islamic finance or to reduced-form credit modeling at $r = 0$.

Sukuk and takaful pricing literature. Jobst [36] and Jobst et al. [37] document the economics and issuance of Islamic securitization. Uddin et al. [54] document the LIBOR-embedding problem in sukuk pricing. Empirical studies consistently find that sukuk spreads behave similarly to conventional bond spreads (Reboredo et al. [49]; Hassan et al. [26]; Cakir & Raei [15]), meaning that the same mathematical credit risk machinery drives both — but that machinery currently requires $r > 0$. Our framework provides a *riba*-free alternative that preserves the same spread dynamics through κ . For takaful: Billah [10, 11] identifies the tension between actuarial fairness and *riba*-free discounting. Htay & Salman [27] attempt a *riba*-free contribution formula but encounter a circularity that the κ -rate framework resolves exactly, while the accounting and governance treatment of Htay et al. [28] sets out the operational standards a transparent, on-chain takaful pool must meet. The retakaful literature [2] develops contribution optimization models that are compatible with the present framework.

Islamic finance, jurisprudence, and fintech. El-Gamal [22] provides the canonical economic analysis of *riba*. Iqbal & Mirakhor [31] and Chapra [16] establish the normative framework. Khan & Ahmed [42] document that Islamic institutions are structurally underhedged. Kamali [39] provides the jurisprudential analysis of derivatives. Recent literature has begun exploring blockchain-based Islamic finance: Meera & Larbani [45] critique the ownership effects of fractional-reserve money

creation from an Islamic perspective; more recently, tokenized sukuk on Ethereum [41] demonstrate that smart contracts can enforce the asset-backing requirement of AAOIFI standards at the protocol layer — directly enabling the on-chain κ -parameterized sukuk we propose in Section 6.1. SeekersGuidance [51] provides contemporary conservative jurisprudential support for the narrower claim that the $\iota = 0$ condition removes the *riba* objection specifically; the same ruling nonetheless maintains that leveraged perpetual contracts are impermissible on separate *gharar* and ownership (*qabdh*) grounds, which concern contractual structure rather than the pricing parameter.

Contributions

This paper makes seven contributions.

1. **Formal equivalence theorem.** We prove that the Akerer-Hugonnier-Jermann perpetual pricing formula and the Duffie-Singleton defaultable bond formula are formally isomorphic (Proposition 1), and that no-arbitrage at $\iota = 0$ implies no-arbitrage for credit pricing at $r = 0$ (Corollary 1).
2. **Islamic perpetual sukuk.** We derive a closed-form no-arbitrage price for an Islamic perpetual sukuk using the κ -parameterized stopping time (Section 6.1).
3. **Actuarially fair takaful formula.** We apply Proposition 6 of Akerer et al. to derive a closed-form everlasting takaful contribution pricing formula — *riba*-free and on-chain verifiable, with product-layer questions reserved to Shariah scholars (Section 6.2).
4. **Islamic credit default swap.** We derive a *riba*-free iCDS whose par spread equals $\kappa(1 - \delta)$ and is independent of any interest rate, giving Islamic finance a credit-hedging instrument it has lacked (Section 6.3).
5. **The κ -rate.** We propose the convergence intensity κ as an Islamic alternative to the risk-free rate r , providing the first mathematically rigorous pricing foundation for a complete *riba*-free credit market (Section 7).
6. **Separation Theorem.** We formally establish that interest is an eliminable component of no-arbitrage pricing — the price at $\iota = 0$ is well-defined and arbitrage-free (Theorem 2) — providing the mathematical statement that Islamic finance has required but lacked.
7. **Shariah analysis.** We analyze the *riba* question formally and set out the *gharar*, *maysir*, and AAOIFI considerations [1] that remain for scholarly determination (Section 8).

2 The Akerer-Hugonnier-Jermann Framework

2.1 Market Setup

Let $(\Omega, \mathcal{F}, \mathbb{P})$ be a probability space with filtration $(\mathcal{F}_t)_{t \geq 0}$ satisfying the usual conditions (Protter [48]). There are two funding accounts: a long account growing at rate $r_a(t)$

and a short account growing at rate $r_b(t)$, both strictly positive. All pricing below is under the long's pricing measure Q^a , the risk-neutral (equivalent martingale) measure whose existence is equivalent to absence of arbitrage in the sense of Harrison & Kreps [25]. Let X_t denote the price of a spot asset satisfying, under Q^a :

$$\frac{dX_t}{X_t} = (r_a(t) - r_b(t)) dt + \sigma_X dW_t^{Q^a} \quad (4)$$

where W^{Q^a} is a standard Brownian motion under Q^a and $\sigma_X > 0$ is the spot volatility. A perpetual futures contract on X_t with convergence parameter $\kappa > 0$ and interest parameter $\iota \geq 0$ is a continuously-settled contract with no fixed expiry.

2.2 The Perpetual Futures Price

Under constant parameters $r_a, r_b, \kappa, \iota, \sigma_X$, Proposition 3 of Akerer et al. [3] establishes that the unique no-arbitrage perpetual futures price is:

$$f_t = \frac{\kappa - \iota}{\kappa + r_b - r_a} \cdot X_t \quad (5)$$

provided $\kappa + r_b - r_a > 0$. The ratio f_t/X_t is the perpetual *basis*. When $r_a = r_b$ (symmetric funding rates):

$$f_t = \frac{\kappa - \iota}{\kappa} \cdot X_t \quad (6)$$

2.3 The Random Time Representation

The central result of Akerer et al. is a *random time representation* that reveals the economic structure of perpetual pricing. Define the **convergence stopping time** θ_t as a random variable with conditional density under Q^a :

$$q(\theta_t \in ds \mid \mathcal{F}_t) = (\kappa_s - \iota_s) \times \exp\left(-\int_t^s (\kappa_u - \iota_u) du\right) ds, \quad s \geq t \quad (7)$$

This is a Cox process (doubly stochastic Poisson process) with intensity $\kappa_t - \iota_t$, as formalized in the theory of point processes by Brémaud [14] and the credit risk context by Lando [43]. The distribution of θ_t is exponential with rate $\kappa - \iota$ in the constant-parameter case:

$$Q^a(\theta_t > s \mid \mathcal{F}_t) = e^{-(\kappa - \iota)(s - t)}, \quad s \geq t \quad (8)$$

The perpetual futures price then satisfies (Theorem 3, Akerer et al.):

$$f_t = \mathbb{E}_t^{Q^a} \left[e^{-\int_t^{\theta_t} \iota_s ds} \cdot X_{\theta_t} \right] \quad (9)$$

The price is the discounted expected spot value at the random stopping time θ_t . The discount is ι (the interest parameter), not r_a or r_b .

2.4 The Critical Case: $\iota = 0$

When $\iota = 0$, Theorem 3 of Akerer et al. establishes that no-arbitrage is satisfied. Substituting into equations (7)–(9):

$$q(\theta_t \in ds \mid \mathcal{F}_t) = \kappa e^{-\kappa(s-t)} ds \quad (10)$$

$$f_t = \mathbb{E}_t^{Q^a} [X_{\theta_t}], \quad \theta_t \sim t + \text{Exp}(\kappa) \quad (11)$$

Three consequences follow immediately:

1. **No discounting.** The expectation in (11) is taken without any interest discount. The price is the pure expected value.
2. **κ carries all timing information.** The probability distribution of *when* payoff occurs is governed entirely by κ . Larger κ means faster convergence and shorter expected stopping time; smaller κ means slower convergence.
3. **The price formula at $r_a = r_b$.** From (6) with $\iota = 0$: $f_t = X_t$. The perpetual futures price equals the spot price identically — no premium, no discount — consistent with the empirical finding that dYdX v4 funding at $\iota = 0$ oscillates through zero (negative in $\sim 26\%$ of intervals) with no structural floor, rather than pinning to an interest constant as $\iota > 0$ venues do [5]. (The *realised* sample mean over the bullish 2025 window was positive, $\approx 7\%$ /yr annualised — a market premium, not a structural charge; the companion paper reports both numbers.)

Remark 1. The CEX interest parameter $\iota = 0.0001$ per 8 hours (= 10.95%/year on Binance, Bybit, OKX) creates a permanent, market-independent wedge in equation (6): $f_t = \frac{\kappa - \iota}{\kappa} X_t \neq X_t$ whenever $\iota \neq 0$. This fixed, predetermined deviation of the perpetual from spot — present regardless of its sign — is structurally *riba al-fadl* (predetermined excess in exchange) under all four Sunni schools. Setting $\iota = 0$ eliminates it exactly ($f_t = X_t$).

3 Reduced-Form Credit Risk Models

3.1 The Duffie-Singleton Framework

The reduced-form approach to credit risk, initiated by Jarrow & Turnbull [33] and systematized by Duffie & Singleton [19, 20], models default as the first jump of a Cox process with stochastic intensity. The comprehensive mathematical treatment is provided by Bielecki & Rutkowski [9] and Jeanblanc, Yor & Chesney [35].

Let $(\Omega, \mathcal{F}, Q)$ be a risk-neutral probability space enlarged to carry the default indicator $H_t = \mathbf{1}_{\tau \leq t}$ (Jeanblanc & Rutkowski [34]; Elliott, Jeanblanc & Yor [23]). Define the **default time** τ as a stopping time with respect to $(\mathcal{F}_t)$, characterized by the Q -conditional survival probability:

$$Q(\tau > s \mid \mathcal{F}_t) = \exp\left(-\int_t^s \lambda_u du\right), \quad s \geq t \quad (12)$$

where $\lambda_t > 0$ is the **hazard rate** (default intensity). The hazard rate has the conditional density:

$$q(\tau \in ds \mid \mathcal{F}_t) = \lambda_s \exp\left(-\int_t^s \lambda_u du\right) ds, \quad s \geq t \quad (13)$$

3.2 Defaultable Instrument Pricing

For a defaultable instrument with payoff function $\phi(X_\tau)$ at the default time τ and risk-free discount rate r_t :

$$P_t = \mathbb{E}_t^Q \left[e^{-\int_t^\tau r_s ds} \cdot \phi(X_\tau) \right] \quad (14)$$

In the constant-parameter case (r, λ constant, $\tau \sim t + \text{Exp}(\lambda)$), this gives:

$$P_t = \int_t^\infty e^{-r(s-t)} \mathbb{E}_t[\phi(X_s)] \lambda e^{-\lambda(s-t)} ds \quad (15)$$

which admits the closed form $P_t = \frac{\lambda}{r+\lambda} \phi(X_t)$ when ϕ is linear and the spot is a Q -martingale. For a bond with full recovery ($\phi(x) = x$) and Q -martingale spot:

$$P_t = \frac{\lambda}{r+\lambda} \cdot X_t \quad (16)$$

In the conventional framework, $r > 0$ is what makes $P_t < X_t$: with no interest discount the bond price equals the spot, seemingly erasing the credit spread. Section 4 resolves this.

4 The Mathematical Equivalence

4.1 The Isomorphism

The resolution begins by observing that the credit and perpetual frameworks share a single mathematical skeleton. Comparing equations (7) and (13) reveals an immediate structural identity:

AHJ Perpetuals	Duffie-Singleton Credit	Role
θ_t	τ	Random stopping time
$\kappa_t - \iota_t$	λ_t	Event intensity*
ι_t	r_t	Discount parameter
X_{θ_t}	$\phi(X_\tau)$	Payoff at event
f_t	P_t	Instrument price
Q^a	Q	Pricing measure

*The AHJ stopping intensity is $\kappa - \iota$ (equation (7)); at the *riba*-free point $\iota = 0$ it coincides with κ , giving the headline identification $\kappa \leftrightarrow \lambda$ used throughout.

The two frameworks share an identical mathematical skeleton. The perpetual futures framework uses κ and ι ; the credit framework uses λ and r . The functions they play are identical.

4.2 The Main Result

Proposition 1 (Random Stopping Time Equivalence). *Let $(\Omega, \mathcal{F}, \mathbb{P})$ be a filtered probability space with the usual conditions. Let X_t be an asset price process satisfying (4) under a pricing measure $Q = Q^a$. Define:*

1. A perpetual futures contract with parameters κ, ι , priced by equation (9).
2. A defaultable financial claim with hazard rate λ , discount rate r , and recovery $\phi(X_\tau)$, priced by equation (14).

Under the exact parameter identification $\kappa - \iota \leftrightarrow \lambda$, $\iota \leftrightarrow r$, $X_{\theta_t} \leftrightarrow \phi(X_\tau)$:

$$f_t = P_t \quad (17)$$

The two pricing formulas are formally isomorphic for every $\iota \geq 0$. At the riba-free point $\iota = r = 0$ the identification reduces to the headline form $\kappa \leftrightarrow \lambda$ used throughout this paper.

Proof. The stopping time densities (7) and (13) are identical under $\kappa - \iota \leftrightarrow \lambda$: by (7) the AHJ stopping intensity is $\kappa - \iota$, which maps to the credit intensity λ . The integral representations (9) and (14) are then identical under the additional identification $\iota \leftrightarrow r$ and $X_{\theta_t} \leftrightarrow \phi(X_\tau)$. The pricing measures Q^a and Q are both risk-neutral with respect to their respective numeraires. The formal identity (17) follows directly, and at $\iota = 0$ the intensity identification is $\kappa \leftrightarrow \lambda$. $\square$

Corollary 1 (Riba-Free Credit Pricing is No-Arbitrage). *Setting $\iota = 0$ in the AHJ perpetual futures framework yields a no-arbitrage price by Theorem 3 of [3]. Under the isomorphism of Proposition 1, the corresponding credit instrument at $r = 0$ is also no-arbitrage, with pricing formula:*

$$P_t = \mathbb{E}_t^Q[\phi(X_\tau)], \quad \tau \sim t + \text{Exp}(\kappa), \quad r = 0 \quad (18)$$

The hazard rate κ alone carries all timing information. No positive interest rate is required.

Proof. By Theorem 3 of Akerer et al. [3], setting $\iota = 0$ satisfies the no-arbitrage conditions for the perpetual futures. By Proposition 1, $f_t = P_t$ under the given identification. Therefore P_t at $r = 0$ satisfies the same no-arbitrage conditions. Specifically: $P_t = \mathbb{E}_t^Q[\phi(X_\tau)]$ is arbitrage-free, being the expectation of the undiscounted payoff under an equivalent martingale measure with Q -intensity κ . $\square$

Remark 2. In this paper's applications (sukuk redemption, takaful triggers) the stopping time τ is contractually specified, mirroring AHJ's contractual stopping time, so the perpetual-style uniqueness argument applies directly. For exogenous default (the iCDS of Section 6.3) the Q -intensity is an input calibrated from market spreads, as in conventional reduced-form practice; the market is incomplete and the intensity is not identified by replication in the underlying alone.

4.3 Resolving the Apparent Paradox

Corollary 1 may appear paradoxical: how can a credit instrument with $r = 0$ have $P_t < X_t$ (i.e., trade at a

discount) without interest? The resolution requires care, because two distinct instruments are in play and they behave differently in κ . The discount-from-par comes from the recovery haircut and is κ -independent; the per-annum *spread* comes from a fixed-maturity survival claim, in which κ does enter. We treat them in turn.

From equation (6) with $\iota = 0$:

$$\frac{f_t}{X_t} = 1 \quad (r_a = r_b, \iota = 0) \quad (19)$$

First, consider a claim that pays the recovery value $\phi(X_\tau) = \delta \cdot X_\tau$ at the default time τ , where $\delta \in (0, 1)$ is the recovery fraction. Since X is a Q -martingale and τ is a stopping time, optional stopping gives $\mathbb{E}_t^Q[X_\tau] = X_t$, so

$$P_t = \delta \cdot \mathbb{E}_t^Q[X_\tau] = \delta \cdot X_t < X_t. \quad (20)$$

The discount from par is $1 - \delta$: the credit loss on default. Note that κ *does not appear* here — because the payoff is collected *at* the random default time, the timing intensity cancels, and the price is the pure recovery haircut. This is not interest; it is the recovery-rate pricing of credit risk.

The per-annum credit *spread* arises in a different instrument: a claim of fixed maturity T that pays par at T if no default occurs by T , and whose default recovery is a fraction δ of its pre-default market value — the recovery-of-market-value convention of Duffie & Singleton [19]. At $r = 0$ its price is exactly $P_t = \mathcal{P} e^{-\kappa(1-\delta)(T-t)}$, where $\mathcal{P}$ is the par value, giving a continuously-compounded spread of exactly $\kappa(1 - \delta)$. (Under recovery of face the $r = 0$ price is instead $\mathcal{P} [\delta + (1 - \delta)e^{-\kappa(T-t)}]$, whose spread equals $\kappa(1 - \delta)$ to first order in $\kappa(T - t)$.) Here κ enters directly. Both instruments are priced with $\iota = 0$ and neither uses an interest rate: the first isolates the recovery haircut, the second the hazard-driven spread. This is the separation of credit risk from time preference that Islamic finance requires.

Remark 3. In the conventional framework, the credit spread over the risk-free rate bundles two economically distinct risks: (1) the probability of default (hazard rate λ) and (2) the time value of money (risk-free rate r). The Akerer framework separates them definitively. At $\iota = r = 0$, only the hazard rate κ remains, representing pure credit risk with no interest component.

4.4 Empirical Evidence: Negative Interest Rate Environments

Our theoretical result — that credit markets can price correctly at $r = 0$ — finds empirical corroboration in the European negative interest rate experiment (2014–2022). The European Central Bank reduced its Deposit Facility Rate to -0.10% in June 2014, eventually reaching -0.50% by September 2019 (ECB [24]).

During this period, CDS markets and corporate bond markets continued to function efficiently in the eurozone: credit spreads reflected default risk accurately; sukuk issuance continued (IIFM [29]); and institutional credit

pricing did not break down. The absence of a positive r in some jurisdictions did not render credit derivatives unpriceable.

This is precisely the empirical analog of our theoretical result: the credit pricing machinery (hazard rates, survival probabilities, recovery values) does not require $r > 0$ for mathematical coherence. The interest rate is indeed eliminable from credit pricing, as our Separation Theorem (Theorem 2) shows. At $r \approx 0$ or $r < 0$, markets revealed empirically what the $\iota = r = 0$ correspondence establishes theoretically.

4.5 Direct Validation: $\hat{\kappa} = s/(1 - \delta)$ on Real Sukuk

The negative-rate episode shows credit markets functioning without a positive r ; we can go further and test the central identification of this paper — the recovery-discount relation $s = \kappa(1 - \delta)$ of Section 4, equivalently $\hat{\kappa} = s/(1 - \delta)$ — *directly* on observed sukuk. Two Cbonds pulls are used: a dedicated sovereign-screener cross-section (61 sukuk from 12 issuers, including Bahrain, Egypt, Oman, Pakistan and Türkiye, with ratings down to CCC — the sample of Table 1), and the broader active-quote snapshot of **195 instruments** (40 sovereign, 154 corporate, one municipal), spanning investment-grade and, via the corporates, high-yield ratings. For each instrument we take the observed G-spread s over the maturity-matched US Treasury curve — which, for these bullet sukuk, closely approximates the option-adjusted spread — and apply the identification with a senior recovery $\delta = 0.60$ [47].

Two predictions of the theory hold on the real cross-section (Table 1). First, **magnitude**: the implied $\hat{\kappa}$ has median 0.02 per annum (sovereign subset 0.019) — exactly the 10^{-2} order of magnitude the credit interpretation of κ requires. It is read off real prices, not chosen as a free parameter. Second, **credit monotonicity**: $\hat{\kappa}$ separates cleanly across the investment-grade/high-yield boundary, from ≈ 0.01 per annum at single-A to ≈ 0.05 – 0.07 at B/CCC, confirming that κ — identified with *zero* interest input — is genuinely measuring credit intensity, precisely as the isomorphism $\kappa \leftrightarrow \lambda$ predicts. The Separation Theorem is thus not only provable but *measurable*: the credit spread is recovered as $\kappa(1 - \delta)$ on traded instruments, with no interest term.

The recovery assumption only rescales the level: $\delta \in [0.4, 0.7]$ multiplies every $\hat{\kappa}$ by a common factor and leaves the ordering untouched. That the extracted $\hat{\kappa}$ is genuine credit risk, and not an artefact of Islamic contract structure, is confirmed by an instrument-level sukuk-versus-conventional comparison on the same issuers. Matched at maturity, the USD sukuk and conventional bonds of Indonesia, Saudi Arabia and Türkiye price within ≈ 1 bp (per-issuer median sukuk-minus-conventional G-spreads of -1.3 , -0.9 and $+0.0$ bp; the pooled median across all 48 matched pairs, which also include Bahrain’s -48 bp

Rating	Median s (bp)	Median $\hat{\kappa}$ (yr $^{-1}$)
AA (IG)	72	0.018
A (IG)	37	0.009
BBB (IG)	78	0.019
BB (HY)	195	0.049
B (HY)	215	0.054
CCC (HY)	295	0.074

Table 1: The identification $\hat{\kappa} = s/(1 - \delta)$ on the real sovereign USD-sukuk cross-section (dedicated sovereign-screener Cbonds pull, 61 bonds, 12 issuers; $\delta = 0.60$; s = observed G-spread over Treasury). $\hat{\kappa}$ sits in the predicted 10^{-2} per-annum credit regime and rises across the IG/HY boundary (the AA bucket, Saudi-dominated, prices just above single-A — the one investment-grade inversion); the separate 195-instrument active-quote snapshot (sovereign + corporate; its filter excludes the high-yield sovereigns shown here) gives the same median magnitude (0.02) and the same broadly-increasing credit ordering. Full instrument-level analysis in the companion κ -yield-curve paper [6].

captive-Islamic-bid exception, is -2.4 bp) [6]. The sukuk spread is therefore the sovereign credit spread, and $\hat{\kappa} = s/(1 - \delta)$ recovers the same hazard intensity λ the isomorphism identifies it with. The relation is dynamic as well as static: daily Cbonds histories of 54 sukuk from five issuers (2018–2026; a dedicated dynamics pull that also covers high-yield issuers outside the 195-bond active-quote snapshot), reconstructed at fixed maturity to remove roll-down, show κ mean-reverting to issuer-specific long-run levels (investment grade ≈ 1 – 3% , high yield ≈ 5 – 7% per annum; cf. the sovereign-spread dynamics literature [21]), with the Feller non-negativity condition $2\alpha\bar{\kappa} \geq \nu^2$ satisfied in every series — the empirical counterpart of the stochastic- κ extension of Section 7.3.

5 Everlasting Options and the Takaful Pricing Formula

Proposition 6 of Akerer et al. [3] extends the random time framework to arbitrary payoff functions $\phi(X_\tau)$, yielding closed-form **everlasting option** prices. Unlike the classical fixed-expiry no-arbitrage option price of Black & Scholes [13], which discounts the terminal payoff at the risk-free rate r over a deterministic horizon, the everlasting option prices the payoff at a *random* stopping time and, at $\iota = 0$, carries no interest discount at all. This extension is the key to takaful pricing.

5.1 The Everlasting Put Option

For a payoff $\phi(x) = \max(K - x, 0)$ (floor protection at strike K) with $\iota = 0$ and $r_a = r_b$, the everlasting put is the no-arbitrage value of the floor payoff delivered at the random convergence time $\tau \sim t + \text{Exp}(\kappa)$ (Proposition 6

of Akerer et al. [3]):

$$\Pi_{\text{put}}(x, K) = \mathbb{E}^{Q^a}[(K - X_\tau)^+]. \quad (21)$$

Let $\beta_p < 0 < 1 < \beta_c$ be the two roots of the characteristic equation

$$\frac{\sigma_X^2}{2} \beta(\beta - 1) - \kappa = 0 \quad (\iota = 0, r_a = r_b), \quad (22)$$

that is,

$$\beta_p = \frac{1}{2} - \sqrt{\frac{1}{4} + \frac{2\kappa}{\sigma_X^2}}, \quad \beta_c = \frac{1}{2} + \sqrt{\frac{1}{4} + \frac{2\kappa}{\sigma_X^2}}, \quad \beta_p + \beta_c = 1. \quad (23)$$

Because the genuine expectation (21) is C^1 and carries intrinsic value when in the money, the everlasting put is the *two-region* solution of Proposition 6 (value-matching and smooth-pasting at $x = K$, with the intrinsic value $K - x$ as particular solution where $x < K$):

$$\Pi_{\text{put}}(x, K) = \begin{cases} C_p x^{\beta_p}, & x \geq K, \\ (K - x) + \frac{K^{1-\beta_c}}{\beta_c - \beta_p} x^{\beta_c}, & x < K, \end{cases} \quad C_p = \frac{K^{1-\beta_p}}{\beta_c - \beta_p}. \quad (24)$$

The coefficient denominator is $\beta_c - \beta_p = 2\sqrt{\frac{1}{4} + 2\kappa/\sigma_X^2}$ (the gap between the two roots), not $1 - \beta_p$: a single power law $C_p x^{\beta_p}$ extended through $x < K$ would omit the intrinsic value and misprice the floor.

5.2 The Everlasting Call Option

For $\phi(x) = \max(x - K, 0)$, the everlasting call is, symmetrically,

$$\Pi_{\text{call}}(x, K) = \begin{cases} \frac{K^{1-\beta_c}}{\beta_c - \beta_p} x^{\beta_c}, & x \leq K, \\ (x - K) + C_p x^{\beta_p}, & x > K, \end{cases} \quad (25)$$

with the same root gap $\beta_c - \beta_p$ in every coefficient. With these (AHJ Proposition 6) coefficients the everlasting put-call parity holds exactly, $\Pi_{\text{call}} - \Pi_{\text{put}} = X_t - K$ (forward parity without discounting, since $\mathbb{E}^{Q^a}[X_\tau] = X_t$); a single power law in each leg would violate it.

6 Applications in Islamic Finance

6.1 Islamic Perpetual Sukuk

Sukuk spreads have been extensively documented in empirical literature [49, 26, 15]: they behave like credit spreads, with ratings and asset quality as primary drivers. The existing pricing models, however, embed r (LIBOR/SOFR) at every step [54, 37]. The framework below preserves the spread dynamics through κ while removing riba.

Definition 1 (Islamic Perpetual Sukuk). An Islamic Perpetual Sukuk is a financial instrument in which:

1. The sukuk holder contributes capital backed by a real asset (RWA) with market value X_t .
2. The issuer (asset operator) pays a continuous premium $\pi_t = \kappa(f_t - X_t) dt$ to the holder — the perpetual futures funding rate under the κ -parameterized formula.
3. At a random redemption time τ (oracle-triggered by asset completion, sale, or Shariah board review), the holder receives the fair market value X_τ of the underlying asset.
4. The interest parameter $\iota = 0$ at all times.

Remark 4. The random redemption time τ is analogous to the “purchase undertaking” (*wa'd*) mechanism already approved under AAOIFI Standard 17 [1, 4]. Smart-contract implementation on a public blockchain makes the redemption trigger transparent and independently verifiable, eliminating the opacity (*gharar*) that traditional *wa'd* structures carry [41].

Proposition 2 (Sukuk No-Arbitrage Price). *The fair price of the Islamic Perpetual Sukuk at time t is:*

$$P_t^{\text{sukuk}} = \mathbb{E}_t^Q[X_\tau] = \frac{\kappa}{\kappa + r_b - r_a} \cdot X_t \quad (26)$$

with $r_a = r_b$ (symmetric bilateral rates) giving $P_t = X_t$. This price is no-arbitrage by Corollary 1.

Proof. The redemption time $\tau \sim t + \text{Exp}(\kappa)$ under the pricing measure Q . With recovery $\phi(X_\tau) = X_\tau$, Corollary 1 gives $P_t = \mathbb{E}_t^Q[X_\tau]$. Under the geometric Brownian motion (4): $\mathbb{E}_t^Q[X_\tau] = X_t \cdot \kappa / (\kappa + r_b - r_a)$, which is equation (26). No-arbitrage follows from Corollary 1. $\square$

Economic interpretation. The sukuk is priced at the expected asset value at redemption, with no interest discount. The premium payments π_t transfer the basis risk from the issuer to the holder: when the asset trades at a premium to spot ($f_t > X_t$), the operator pays the holder; when at a discount, the holder pays the operator. This is risk-sharing (*mudarabah*-compatible) rather than a fixed return on capital.

The demand-side question, stated plainly. An honest reading of the constant-parameter equilibrium should be recorded here rather than left for a referee to find. At $r_a = r_b$ and $\iota = 0$ the model gives $f_t = X_t$ identically, so the equilibrium premium flow $\pi_t = \kappa(f_t - X_t) dt$ is *zero*: the holder pays X_t , receives X_τ at redemption, and earns nothing in expectation beyond exposure to the asset itself. That is exactly what an interest-free instrument *should* deliver in this benchmark — the programme’s monetary papers show the economy prices time through state-contingent equity returns, not through this instrument — but it poses a real product-design question: why hold the perpetual sukuk rather than the asset? The answers live outside the constant-parameter benchmark (liquidity and divisibility of the claim versus the asset; the funding flow when $r_a \neq r_b$ or when the basis deviates;

monetary and regulatory roles of a standardised claim), and the welfare value of what the instrument does *not* provide — a riskless intertemporal trade — is quantified in the companion safe-asset paper. We flag the question because a pricing theorem is not a demand theory, and the corpus should not read as if it were.

The κ parameter for sukuk. Two distinct intensities must not be conflated. For the perpetual sukuk above, κ is the contractual *redemption intensity* ($\kappa_{\text{red}} = 1/\text{expected horizon}$): a sukuk backed by an infrastructure asset with a 5-year expected horizon gives $\kappa_{\text{red}} \approx 0.20$; a 10-year asset, $\kappa_{\text{red}} \approx 0.10$. Credit quality — rating, asset quality, macro conditions [49, 26] — enters through the priced default-intensity component: the credit $\hat{\kappa} = s/(1 - \delta) \sim 10^{-2}$ per annum of Section 4.5, not through $1/T$. Both are directly observable and estimable without any interest rate input.

6.2 Takaful: Actuarially Fair Perpetual Insurance

Definition 2 (Everlasting Takaful Contract). An Everlasting Takaful Contract is a mutual insurance arrangement in which:

1. Participants pay a continuous premium πdt per unit time into a mutual pool until the claim event τ .
2. The pool pays $\phi(X_\tau) = \max(K - X_\tau, 0)$ — the shortfall of asset value below the insured floor K — at the claim event τ .
3. The claim event follows a Cox process with intensity κ (the historical hazard rate of the insured event: drought, flood, death, property damage).
4. No investment return is generated on the pool; it earns at rate $\iota = 0$.

Theorem 1 (Actuarially Fair Takaful Premium). *Under actuarial equilibrium (PV of contributions = PV of expected claim), the fair continuous premium rate π and the lump-sum contribution Π at inception for floor protection at level K , given current asset value $x = X_t$, intensity κ , and volatility σ_X , satisfy:*

$$\Pi_{\text{takaful}}(x, K) = \frac{K^{1-\beta_p}}{\beta_c - \beta_p} \cdot x^{\beta_p} \quad (x \geq K) \quad (27)$$

$$\pi_{\text{continuous}} = \kappa \cdot \Pi_{\text{takaful}}(x, K) \quad (28)$$

where $\beta_p = \frac{1}{2} - \sqrt{\frac{1}{4} + \frac{2\kappa}{\sigma_X^2}} < 0$. Both prices are no-arbitrage (Corollary 1) and require no positive interest rate.

Proof. Lump-sum contribution. By Corollary 1, the no-arbitrage price of $(K - X_\tau)^+$ at $r = 0$ is $\mathbb{E}_t^Q[(K - X_\tau)^+]$. By Proposition 6 of Akerer et al. [3] applied with $\iota = 0$ and $r_a = r_b$, this equals $C_p x^{\beta_p}$ with C_p, β_p as in (24) and (23), giving (27).

Continuous premium. If the participant pays continuous premium π until $\tau \sim t + \text{Exp}(\kappa)$, the present value of premium flows at $r = 0$ is $\mathbb{E}[\int_t^\tau \pi ds] = \pi/\kappa$. Setting $\pi/\kappa = \Pi_{\text{takaful}}$ gives $\pi = \kappa \cdot \Pi_{\text{takaful}}$, which is (28). $\square$

Physical vs. pricing intensity. The definition specifies κ as the historical hazard rate (a physical-measure object), while Theorem 1 prices under Q ; the actuarially fair premium equals the no-arbitrage price only if the claim intensity carries no risk premium, i.e. $\kappa^P = \kappa^Q$. This identification is appropriate for idiosyncratic, diversifiable claims (individual mortality, property damage); for systematic perils that strike many participants simultaneously (regional drought or flood), the pricing intensity κ^Q embeds an event-risk premium and must be calibrated separately from the historical hazard. This distinction governs the whole program and deserves a taxonomy: *spread- and price-implied* intensities (the sovereign and corporate $\hat{\kappa}$ of the companion yield-curve study) are Q -measure objects — the credit-spread literature finds such spreads run a multiple of actuarial expected loss — while *frequency-based* intensities (crop-loss, mortality) are P -measure objects. The two are commensurable in *kind* (both are event intensities, neither a price of time — the Shariah-relevant property) but not in *level*; downstream constructions that mix them (portfolio floors, cross-domain comparisons) should be read as conservative whenever Q -volatilities exceed P -volatilities, and takaful priced at κ^P is the actuarially fair point for a mutual pool precisely because no event-risk premium is extracted from participants.

Agricultural takaful example. Let X_t be the market price of wheat (USD/bushel). A farmer participates in a mutual takaful pool offering floor protection at $K = \$6.00/\text{bushel}$. Historical drought frequency: $\kappa = 0.08$ (one event every ~ 12.5 years). Wheat volatility: $\sigma_X = 0.25$ (25% p.a.). Current wheat price: $x = \$7.50$.

$$\beta_p = 0.5 - \sqrt{0.25 + 2.56} = -1.176, \quad \beta_c = 0.5 + 1.676 = 2.176$$

$$C_p = \frac{(6.00)^{2.176}}{\beta_c - \beta_p} = \frac{49.35}{3.352} = 14.72$$

$$\begin{aligned} \Pi &= 14.72 \times (7.50)^{-1.176} \\ &= 14.72 \times 0.0935 = \$1.38/\text{bushel} \end{aligned} \quad (29)$$

$$\begin{aligned} \pi_{\text{continuous}} &= 0.08 \times 1.38 \\ &= \$0.110/\text{bushel/year} \end{aligned}$$

The farmer contributes \$1.38 per bushel at inception (or \$0.110 per year continuously). This is actuarially fair. No interest rate was used in any step.

A design caveat this example must carry. The Cox construction models the claim event τ and the price X_t as *independent*, and for this peril they are not: a drought is a

supply shock that tends to *raise* the crop price, so a price-floor payoff $(K - X_\tau)^+$ evaluated at the drought time pays *least* exactly when the insured event strikes — while the farmer’s actual loss is in *quantity*, not price. As specified, the contract is therefore floor protection against *low-price* states (a bumper-harvest or demand-collapse hedge), not drought indemnity; a drought product should pay on a *yield* index (as the companion takaful paper’s loss-cost calibration does) or model the event–price correlation explicitly, at the cost of the closed form. The mathematics of the example is exact; its peril label is the thing to design carefully. Life takaful, next, is immune to this issue — the benefit is a fixed B , not a function of a correlated price.

Life takaful — continuous premium formula. For a fixed death benefit B payable at death $\tau \sim \text{Exp}(\kappa_{\text{mort}})$ (where κ_{mort} is the age-adjusted mortality hazard rate from actuarial tables), the fair continuous premium π_{life} satisfies:

$$\frac{\pi_{\text{life}}}{\kappa_{\text{mort}}} = B \implies \pi_{\text{life}} = \kappa_{\text{mort}} \cdot B \quad (30)$$

The participant pays $\kappa_{\text{mort}} B$ per unit time until death; the pool pays B at death. Both sides discount at $r = 0$. A 35-year-old with $\kappa_{\text{mort}} \approx 0.002$ and $B = \$100,000$ pays \$200/year — consistent with real-world life takaful contributions for that age bracket. As the participant ages (κ_{mort} rises), the formula automatically reprices contributions upward, with full on-chain transparency.

6.3 Islamic Credit Default Swaps (iCDS)

The global CDS market has a notional outstanding of approximately \$10 trillion [8], built entirely on the Duffie-Singleton framework with $r > 0$. Islamic institutions have no equivalent instrument for credit risk hedging [42].

Definition 3 (Islamic CDS (iCDS)). An iCDS is a bilateral contract in which:

1. The protection buyer pays a continuous premium π_{iCDS} per unit time to the protection seller.
2. At the credit event time τ (default, restructuring, or rating trigger), the seller pays the loss given default $\text{LGD}_\tau = (1 - \delta)X_\tau$ to the buyer, where $\delta \in [0, 1]$ is the recovery rate.
3. The interest parameter $\iota = 0$; no fixed rate premium.

Proposition 3 (iCDS Fair Spread). *The fair iCDS continuous premium (the iCDS spread) s^* when the reference sukuk has recovery rate δ , current value X_t , default intensity κ , and $\iota = 0$ is:*

$$s^* = \kappa(1 - \delta) \quad (31)$$

This is the continuous premium rate that equates the PV of protection received to the PV of premiums paid.

Proof. PV of protection received $= \mathbb{E}[(1 - \delta)X_\tau] = (1 - \delta)X_t$ (by Corollary 1 with $\phi(x) = (1 - \delta)x$). PV of premiums paid $= \mathbb{E}[\int_t^\tau s^* ds] = s^*/\kappa$. Setting equal: $s^*/\kappa = (1 - \delta)X_t/X_t = 1 - \delta$ (normalized to par $X_t = 1$), giving $s^* = \kappa(1 - \delta)$. $\square$

The iCDS spread $s^* = \kappa(1 - \delta)$ has the correct economic properties: it increases with κ (higher default risk) and decreases with δ (higher recovery). This is structurally identical to the conventional CDS spread $s = \lambda(1 - \delta)$, with κ replacing λ and no r appearing anywhere.

Pricing intensity and the substance-compliant form. Removing ι removes the price of time, but Proposition 3 as stated prices at the market intensity κ , which — by the taxonomy of §6.2 — is the Q -measure, spread-implied object κ^Q : it embeds the event-risk premium the credit-spread literature finds to run a multiple of actuarial expected loss. Whether charging that premium is itself permissible, once every predetermined charge is gone, is the substantive question the corpus reserves. There is, however, a form that does not raise it. Structured as a *cooperative guarantee* — a *kafala/tabarru* pool in which members mutually indemnify realised default loss, gated on insurable interest and possession of the reference obligation — the iCDS contribution is set at the realised hazard κ^P ,

$$s_{\text{coop}}^* = \kappa^P(1 - \delta), \quad (32)$$

the actuarially fair expected loss, exactly as the mutual-pool contribution of §6.2 is set at κ^P “precisely because no event-risk premium is extracted from participants.” The contested wedge $(\kappa^Q - \kappa^P)(1 - \delta)$ is then not *sold* to a protection writer for profit but retained by the pool and returned as surplus, net of a disclosed *wakala* administration fee. So priced, credit-loss cover is structurally the takaful case this paper has already shown to be actuarially fair and free of both *riba* and the event-risk premium — collapsing the credit case into the insurance case that needs no further resolution, and leaving only the narrower, shared questions of whether tokenised title constitutes valid possession and whether the *wakala* loading is a genuine fee.

6.4 Sovereign Perpetual Sukuk

Sovereign governments in Muslim-majority countries (Malaysia, Saudi Arabia, Indonesia, Türkiye, Pakistan) issue sukuk to fund infrastructure. These are priced against SOFR or government bond yields, embedding *riba* implicitly [29, 32]. The Islamic perpetual sukuk framework resolves this.

Example 1 (Infrastructure Perpetual Sukuk). A sovereign issues a perpetual sukuk backed by a highway project with expected completion in 7 years: the redemption intensity is $\kappa_{\text{red}} = 1/7 \approx 0.143$. The project asset value is X_t (updated quarterly by a government oracle). Under $r_a = r_b$ and $\iota = 0$: $P_t = X_t$ (at-par perpetual).

The sovereign pays premium $\pi_t = \kappa(f_t - X_t) dt$ continuously to sukuk holders. On completion (random event τ), holders receive X_τ . Secondary market price: $P_t = X_t$, updated via observed κ from reported completion timeline. *No yield-to-maturity calculation. No SOFR benchmark. No riba.*

The κ parameter functions as a *sovereign credibility signal*: stable infrastructure management exhibits high, predictable κ ; governance failures widen the basis and imply declining κ — a market-based credit signal without any interest rate.

7 The κ -Rate: An Islamic Monetary Alternative

7.1 The Problem with r

The risk-free interest rate r is the fundamental pricing parameter of conventional finance: exogenous, set by central banks, and definitionally *riba* (it compensates capital for time elapsed, regardless of real economic outcome). Every discount factor in every asset pricing model contains r . Islamic finance has no alternative [16, 31].

7.2 The κ -Rate as a Substitute

Definition 4 (κ -Rate). The κ -rate is the convergence intensity of a perpetual financial instrument linking a price claim to an underlying real asset. It satisfies:

1. **Endogenous determination.** κ is backed out from observed market prices:

$$\hat{\kappa}_t = \frac{f_t/X_t \cdot (r_b - r_a)}{1 - f_t/X_t} \quad (33)$$

valid when $r_a \neq r_b$. In the symmetric regime $r_a = r_b$ the basis is identically one and (33) is uninformative; κ is then identified from episodes of funding asymmetry, from the decay rate of transient basis deviations (whose convergence speed is κ by construction), or — as in Section 4.5 — directly from traded sukuk spreads via $\hat{\kappa} = s/(1-\delta)$. It reflects real economic convergence dynamics, not central bank policy.

2. **Observable and transparent.** κ is computable from any two market prices (f_t and X_t), publishable on-chain, and independently verifiable.
3. **No riba content.** κ is an event intensity, not a reward for the time value of money. It compensates for credit risk (sukuk), actuarial risk (takaful), or convergence risk (perpetuals).
4. **Reviewable by Shariah boards.** κ represents neither interest on capital nor predetermined excess, and is therefore submitted to Shariah-board review as an event intensity; product-level approval remains a jurisprudential determination.

7.3 Extensions: Stochastic κ and Affine Models

Definition 4 takes κ as constant. In practice, convergence intensities vary over time (as project timelines lengthen or credit conditions deteriorate). The affine credit risk framework of Duffie & Singleton [19, 20] and the general theory of Bielecki & Rutkowski [9] provide the natural extension.

Let κ_t follow a CIR process [17]:

$$d\kappa_t = \alpha(\bar{\kappa} - \kappa_t) dt + \nu\sqrt{\kappa_t} dZ_t \quad (34)$$

where $\bar{\kappa} > 0$ is the long-run mean intensity, $\alpha > 0$ is mean-reversion speed, and $\nu > 0$ is the vol-of- κ . Under this specification, the stopping time density (7) remains well-defined but the closed-form prices of Sections 6.1–6.4 require the Laplace transform of $\int_t^T \kappa_s ds$ over a fixed horizon T (available in closed form for the CIR specification via Jeanblanc, Yor & Chesney [35], Prop. 6.3.4):

$$Q(\tau > T | \mathcal{F}_t) = \mathbb{E}_t^Q \left[e^{-\int_t^T \kappa_s ds} \right] = A(T-t) e^{-B(T-t)\kappa_t} \quad (35)$$

where $A(\cdot)$ and $B(\cdot)$ are the standard CIR bond price functions, valid on the pre-default event $\{\tau > t\}$. All qualitative results — no-arbitrage, the Separation Theorem, riba-freedom at the pricing layer — hold under stochastic κ ; only the closed-form expressions require numerical methods. This specification is also empirically grounded: as reported in Section 4.5, the κ backed out of real sovereign sukuk is observably mean-reverting toward issuer-specific long-run means and satisfies the CIR Feller condition $2\alpha\bar{\kappa} \geq \nu^2$ in every series we estimate, so (34) describes the data rather than merely extending the theory.

7.4 The Redemption-Horizon Curve and the Credit κ -Yield Curve

Just as conventional finance builds a yield curve $r(T)$ as a function of maturity T , Islamic finance can build a **redemption-horizon curve**: the mapping $\kappa_{\text{red}}(T)$ of contractual redemption intensities from perpetual sukuk of different expected maturities.

$$\kappa_{\text{red}}(T) = \frac{1}{T}, \quad T = \text{expected horizon} \quad (36)$$

Derived from the exponential distribution ($\kappa = 1/E[\tau]$), the redemption-horizon curve is:

- Downward-sloping by construction (short assets have higher κ_{red}).
- Grounded in real project timelines, not monetary policy expectations.

The traded *credit* κ -yield curve — constructed in the companion paper [6] from observed sukuk spreads via $\hat{\kappa} = s/(1-\delta)$ — is a distinct, measured object and is mildly *upward*-sloping in maturity for investment-grade

issuers (+1.17 to +3.4bp/yr); together the two curves supply the term-structure inputs that the conventional yield curve supplies in Islamic finance contexts.

Instrument	κ calibration	Data source
Perpetual sukuk	Expected project timeline	Engineering reports
Takaful	Historical claim frequency	Actuarial tables
iCDS	Reference entity default history	Rating agencies
Sovereign sukuk	Project completion probability	Government data

8 Shariah Analysis

Following the jurisprudential boundary set out in the companion research note [7], the formal analysis below settles only the *riba* question at the pricing layer; *gharar*, *maysir*, and possession (*qabd*) are product-layer questions governed by rules R1–R4 of that note and reserved to Shariah scholars. The *gharar*, *maysir*, and AAOIFI material that follows is submitted as structural input to that review, not as a verdict.

8.1 Riba: Why $\iota = 0$ Satisfies All Schools

The prohibition of *riba* is universal across all four Sunni schools (Hanafi, Maliki, Shafi'i, Hanbali) and the Ja'fari (Shi'a) school [55, 22]. The Quranic prohibition is absolute (Al-Baqarah 2:275–279; Al-Imran 3:130).

Riba takes two forms:

1. *Riba al-fadl*: excess in exchange. The CEX interest floor $\iota > 0$ is precisely this: longs always pay more than the market premium, with the excess ($\iota \cdot f_t$) pre-determined and fixed.
2. *Riba al-nasi'ah*: excess for delay. The conventional risk-free rate r is this: the lender receives r per unit time regardless of economic outcome.

The κ -rate is neither. It is:

- Not predetermined: κ fluctuates with observed market conditions, backed out from market prices per (33).
- Not a reward for time: κ compensates for the probability and timing of a credit/claim/convergence event (real economic risk), not for the passage of time per se.
- Symmetric: the stopping time distribution does not systematically benefit either party.

Remark 5. SeekersGuidance (2025) — a rigorous contemporary conservative Shariah ruling on perpetual futures — concedes that removing the interest component addresses the *riba* objection, while ultimately maintaining that leveraged perpetuals are impermissible on separate *gharar* and ownership (*qabdh*) grounds. We rely on the narrow concession only: the structural design choice $\iota = 0$ (equivalently, $r = 0$ in credit instruments) removes the interest component by mathematical construction, not legal re-labelling [51]. Whether any particular contract built on the κ -rate is Shariah-compliant in full depends on its other contractual features and requires separate jurisprudential review.

8.2 Gharar: Is the Stopping Time Uncertain?

A potential objection is that the random stopping time τ introduces excessive uncertainty (*gharar*).

Gharar in Islamic jurisprudence refers to uncertainty about the *existence* of the object of transaction, not about its *timing* when timing is actuarially determined [22, 39].

The stopping time $\tau \sim \text{Exp}(\kappa)$ is:

1. Precisely specified by its distribution, which is public and on-chain verifiable.
2. Calibrated from real economic data.
3. Analogous to the random maturity of an *ijara* (lease) that terminates when the lessee exercises an option — a widely accepted structure under AAOIFI Standard 17 [1, 4].

The stopping time introduces actuarial uncertainty (priced by κ) rather than contractual ambiguity; whether this suffices to address the *gharar* concern is a determination for Shariah scholars. Blockchain implementation (smart contracts) makes the trigger condition and the distribution parameters immutably public [41], which we submit as mitigating evidence for that review.

8.3 Maysir: Is This Speculation?

Maysir (gambling) is prohibited because it is zero-sum and divorced from productive economic activity. The κ -rate framework is designed to address the *maysir* concern because:

1. Every instrument is backed by a real asset (X_t).
2. The premium payments $\pi_t = \kappa(f_t - X_t) dt$ are transfers between economically motivated parties based on observed market conditions, not arbitrary bets.
3. Leverage is constrained by the Shariah hard cap ($5\times$ in the Baraka Protocol implementation [5]).

8.4 AAOIFI Considerations

- **Standard 17 (Sukuk).** The Islamic Perpetual Sukuk is designed to address the requirement that sukuk represent ownership in an asset. The stopping time redemption is structured as a purchase undertaking (*wa'd*) of the kind treated under Standard 17 [1, 4].
- **Standard 26 (Takaful).** The everlasting takaful formula is designed to address the actuarial fairness and mutual risk-sharing requirements. Pool contributions are actuarially determined, not based on investment returns.
- **Standard 21 (Financial Paper).** The iCDS, as a bilateral risk-transfer mechanism, is analogous to *wa'd*-based credit enhancement treated under Standard 21, provided the reference obligation is itself permissible.

Conformity with Standards 17, 21, and 26 is a determination reserved to a Shariah supervisory board.

9 Discussion

9.1 Why This Was Not Discovered Earlier

The connection between the Akerer perpetual contract framework and the Duffie-Singleton credit risk model went unnoticed for two reasons. First, the two literatures operate in different communities (mathematical finance vs. fixed income) and are rarely cited together. Second, the Akerer framework itself is very recent (first circulated 2023; published 2025); prior to it, there was no mathematically rigorous proof that $\iota = 0$ achieves no-arbitrage convergence in perpetual contracts, and the equivalence could not therefore be established.

The deeper reason is conceptual. Conventional finance fuses two distinct economic concepts — the timing of a payoff event and the time value of money — into a single parameter r . This fusion appeared necessary. The Akerer framework dissolves this presumption: the stopping time distribution $\text{Exp}(\kappa)$ encodes timing information probabilistically, without any predetermined interest.

9.2 The Separation Theorem

Theorem 2 (Separation of Credit Risk from Time Preference). *In the Akerer-Hugonnier-Jermann framework, the no-arbitrage price of a financial claim $\phi(X_\tau)$ at $\iota = 0$ can be written as:*

$$P_t = \underbrace{\mathbb{E}^Q[\phi(X_\tau)]}_{\text{credit risk term}} \quad \text{without any} \quad \underbrace{e^{-r(s-t)}}_{\text{time preference term}} \quad (37)$$

The timing of the payoff is embedded in the distribution of $\tau \sim \text{Exp}(\kappa)$. The value of the payoff is embedded in the expectation $\mathbb{E}[\phi(X_\tau)]$. The interest parameter ι is eliminable: setting $\iota = 0$ yields a finite, unique, no-arbitrage price with no time-preference term. (For $\iota \neq 0$ the interest parameter is not a separable discount — in the AHJ random-time representation it enters both the discount and the stopping intensity, $\tau \sim \text{Exp}(\kappa - \iota)$. The theorem asserts that the riba-free point $\iota = 0$ is well-defined and arbitrage-free, not that interest factors out multiplicatively for all ι .)

Proof. This is the $\iota = 0$ specialisation of Corollary 1. By Akerer-Hugonnier-Jermann (Theorem 3), the no-arbitrage perpetual price admits the random-time representation $P_t = \mathbb{E}^{Q^a}[e^{-\iota(\tau-t)}\phi(X_\tau)]$ where $\tau \sim t + \text{Exp}(\kappa - \iota)$ — so the interest parameter enters *both* the discount and the stopping intensity. At $\iota = 0$ the intensity reduces to κ and the discount to unity, leaving $P_t = \mathbb{E}^{Q^a}[\phi(X_\tau)]$ with $\tau \sim \text{Exp}(\kappa)$, which is finite and arbitrage-free because $\kappa > 0$ guarantees $\tau < \infty$ a.s. and the expectation exists. Interest is therefore *eliminable* — the riba-free price exists and is unique — even though, away from $\iota = 0$, it is a shift of the convergence intensity rather than a separable multiplicative discount. $\square$

This separation is the mathematically precise statement of what Islamic finance has sought for 50 years: a proof that interest is not necessary for no-arbitrage pricing of time-dependent financial instruments.

9.3 Counterarguments and Responses

Counterargument 1: Time value of money is not riba. Some Islamic economists (notably Kahf [38] and Siddiqi [53]) argue that the time value of money can be Islamically justified through opportunity cost reasoning: capital deployed today could produce real output tomorrow, so a premium for time is fair.

Response: Even granting this argument, it does not imply that $r > 0$ is a mathematical necessity for no-arbitrage pricing. The opportunity cost of capital is already captured by κ : a high- κ instrument (fast convergence, short expected holding period) delivers value quickly and implicitly compensates for the opportunity cost of the capital committed. The Separation Theorem shows that interest is separable from, not identical to, the opportunity cost of time. The κ -rate prices the latter without encoding the former.

Counterargument 2: Conventional sukuk work; why change? Existing sukuk structures have achieved broad market acceptance. One might argue that embedding SOFR at the pricing stage is a pragmatic necessity, even if theoretically impure.

Response: Three practical problems with this position. (i) Post-LIBOR transition, sukuk issuers face SOFR repricing risk that has no Islamic instrument for hedging. (ii) Empirically, sukuk spreads are driven by credit risk not interest rates [49, 26]; the κ -rate captures spread dynamics directly. (iii) AAOIFI Standard 17 explicitly requires that sukuk represent real ownership stakes, not interest-indexed instruments — a requirement that the current market satisfies cosmetically but the κ -framework satisfies mathematically.

Counterargument 3: Stochastic κ introduces model risk. If κ must be estimated, the pricing formula inherits estimation error. Critics may argue that conventional credit models with r (observable from Treasury yields) are more tractable.

Response: This is correct, and we acknowledge it (Section 7.3). However, hazard rates λ are also estimated quantities in conventional reduced-form models [20, 9]; the estimation uncertainty of κ is no greater than that of λ . The difference is that κ can be observed from the perpetuals basis (equation 33) — a market-implied quantity — in addition to being estimated from historical default rates, giving it an additional identification channel unavailable to λ .

9.4 Limitations

Constant-parameter assumption. Closed-form results assume constant κ and σ_X . The stochastic case (Section 7.3) requires numerical methods.

Recovery structure. We assume full or floor-only recovery. Partial recovery models require adjusting ϕ accordingly.

Liquidity. The theory assumes liquid markets for both f_t and X_t . For infrastructure sukuk or agricultural takaful, the underlying may be illiquid. Liquidity adjustment to the κ -rate is an open research question.

Regulatory recognition. Formal adoption requires engagement with the AAOIFI standard-setting process, which is ongoing.

10 Simulation Validation

The theoretical results of Sections 4–7 are grounded in measurable on-chain behaviour. We validate the key claims — no *riba* at $\iota = 0$, protocol solvency under random stopping, and the convergence of mechanism design to the κ -rate regime — using an integrated economic simulation of the Baraka Protocol deployment on Arbitrum Sepolia.

10.1 Simulation Architecture

The simulation combines four coupled layers:

1. **cadCAD Market Engine.** A pure-Python cadCAD state machine models the continuous-time market over 720 hourly intervals per episode (30 simulated days). The five partial-state-update blocks mirror the on-chain contract logic: oracle price update, funding settlement at $\iota = 0$, trader actions, liquidations, and insurance-fund bookkeeping.
2. **Reinforcement Learning Trader.** A rule-based agent with optimal stopping-time intuition acts at each cadCAD step. Its observations are the live cadCAD state $(p_t^{\text{mark}}, p_t^{\text{index}}, F_t, CF_t, q_t, c_t, \ell_t, w_t)$; its actions (`{hold, open long 1–5×, open short 1–5×, close}`) inject positions directly into the cadCAD position pool.
3. **Game Theory Equilibrium Analyser.** Every 50 cadCAD steps, the simulation builds a 5×5 payoff matrix from the live price window and solves for the Nash equilibrium leverage distribution. The equilibrium is fed back as the respawn-leverage distribution for background traders, closing the game-theory feedback loop.
4. **Mechanism Design Optimiser.** At each episode boundary, a `scipy differential_evolution` pass minimises the objective

$$\mathcal{L}(\theta) = 10 \rho(\theta) + 0.005 \bar{\ell}(\theta) + 5 \overline{|F|}(\theta) \quad (38)$$

where ρ is the insolvency rate, $\bar{\ell}$ the mean liquidation count, and $\overline{|F|}$ the mean absolute funding rate. The optimal θ^* is blended (70/30) with current parameters and applied to the next episode.

10.2 Results: 5 Episodes $\times$ 720 Steps

10.3 Claim Verification

Claim 1: $\iota = 0$ implies no systematic *riba*. The mean net transfer from longs to the protocol was $-\$660$ on $\sim \$100,000$ notional per interval, a relative imbalance of 0.66% — well within Monte Carlo noise. This is consistent with Corollary 1: at $\iota = 0$ the funding mechanism imposes no systematic wealth transfer, satisfying the jurisprudential definition of *riba*-freedom.

Claim 2: Protocol solvency under random stopping.

Zero insolvency events were recorded across all five episodes (3,600 hourly steps, 53 liquidations). The insurance fund ended above its $\$100,000$ seed in every episode, finishing at $\$102,835$, consistent with the theoretical prediction that the takaful pool is self-sustaining when premiums are set at the actuarially fair κ -rate level.

Claim 3: Mechanism design converges to κ -regime.

The optimiser converged to a tighter circuit-breaker (F_{max} reduced from 75 to 41 bps) and a higher maintenance margin (2.0% to 3.1%). Both movements are consistent with the κ -rate theory: a protocol optimising under the $\iota = 0$ design constraints should minimise $\overline{|F|}$ and maximise insurance reserves. The mechanism design layer discovered this structure endogenously (noting that its objective encodes the near-zero-funding property; see the companion simulation paper [5]).

Claim 4: Nash equilibrium leverage is sub-maximal.

The mean Nash equilibrium long leverage was $2.72\times$ and short was $3.28\times$, both strictly below the Shariah hard cap of $5\times$. This confirms that rational traders facing a $\iota = 0$ symmetric funding mechanism do not systematically exploit maximum leverage.

10.4 The Empirical κ -Rate

The empirical κ -rate is measured on real markets. On the full Cbonds universe pulls of Section 4.5, the 195-instrument cross-section gives median $\hat{\kappa} \approx 0.02$ per annum (sovereign subset 0.019) with the predicted credit ordering, and the 54-sukuk constant-maturity histories give mean-reverting long-run levels of $\approx 1\text{--}3\%$ (investment grade) to $5\text{--}7\%$ (high yield) per annum. Real credit-aggregate OAS series corroborate the high-yield end of the band: ICE BofA indices (FRED) imply $\hat{\kappa} = \text{OAS}/(1 - \delta)$ of 0.104 per annum (emerging-market high-yield mean; 0.079 current) and 0.081 (US high-yield) at $\delta = 0.60$ — all identified with zero interest input. The protocol simulation above complements this with *computability*: $\hat{\kappa}$ is recoverable from on-chain-observable quantities (insurance-fund flows, funding rates, open interest), at a level (≈ 0.08 per annum) inside the measured high-yield credit band, so a live $\iota = 0$ protocol can publish its own

Table 2: Per-episode simulation outcomes.

Ep.	Ins. Fund (\$)	Liq.	Insolv.	RL PnL (\$)	$\overline{ F }$ (bps)
1	101,175	11	No	+117	70.1
2	100,782	7	No	+310	58.0
3	101,139	9	No	0	50.3
4	101,889	11	No	0	44.2
5	102,835	15	No	0	39.6
Total	—	53	0/5	+427	52.4

Table 3: Mechanism design parameter evolution.

Parameter	Initial	Final	Change
$F_{\max}$ (bps)	75.0	40.8	−45%
Maintenance margin (%)	2.0	3.1	+55%
Liquidation penalty (%)	1.0	1.3	+31%
Insurance split (%)	50.0	60.8	+22%

$\hat{\kappa}$ in real time, with the real-data estimates above as the market calibration.

11 Conclusion

We have established that the random time representation of perpetual futures contracts — a result from contemporary mathematical finance [3] — is formally isomorphic to the reduced-form credit risk pricing formula of Duffie & Singleton [19]. The isomorphism maps the convergence intensity κ to the default hazard rate λ , and the interest parameter ι to the risk-free discount rate r . Setting $\iota = r = 0$, which Akerer et al. prove achieves unique no-arbitrage pricing for perpetual contracts, yields a complete, no-arbitrage pricing framework for credit instruments that requires no positive interest rate.

The implications for Islamic finance are substantial. The κ -rate provides a mathematically rigorous, riba-free pricing foundation — product-layer questions remain with Shariah scholars — for: (i) Islamic Perpetual Sukuk, priced at the expected asset value at a random redemption event; (ii) Everlasting Takaful contracts, with premiums given by the closed-form everlasting put formula; (iii) Islamic Credit Default Swaps, transferring credit risk without any interest component; and (iv) Sovereign Perpetual Sukuk for infrastructure finance, with κ calibrated from project timelines.

The deeper contribution is the Separation Theorem (Theorem 2): interest is a separable, eliminable component of no-arbitrage pricing, not an intrinsic mathemat-

ical necessity. The stopping time distribution carries all timing information previously attributed to the interest discount. Empirical corroboration comes from European credit markets that continued to function efficiently at $r \leq 0$ during 2014–2022 (Section 4.4).

This is the precise mathematical statement that Islamic finance has required but lacked: riba is not a logical prerequisite for fair pricing of time-dependent financial instruments. Future work should address stochastic κ estimation from panel sukuk data, liquidity premia in the κ -yield curve, and the macroeconomic properties of a monetary system anchored to κ rather than r .

*“And Allah has permitted trade and forbidden riba.”
(Al-Qur’an, Al-Baqarah 2:275)*

CRedit Author Contribution Statement

Shehzad Ahmed: Conceptualization, Formal analysis, Methodology, Software, Visualization, Writing — original draft. **Rafiqul Bhuyan:** Conceptualization, Methodology, Supervision, Writing — review and editing.

Declaration of Competing Interest

S. Ahmed has developed an open-source smart-contract implementation of the pricing framework described in this paper (Baraka Protocol, <https://github.com/Arcus-Quant-Fund/BarakaDapp>). No financial interest or commercial funding was involved. The remaining authors declare no competing interests.

References

- [1] AAOIFI. (2017). *Shariah Standards for Islamic Financial Institutions* (Standard No. 17: Investment Sukuk; Standard No. 21: Financial Paper; Standard No. 26: Islamic Insurance). Accounting and Auditing Organization for Islamic Financial Institutions, Manama, Bahrain.
- [2] Abdallah, W., Maysami, R., & Shanmugam, B. (2022). Actuarial model for takaful contributions via optimal retakaful. *Journal of Islamic Accounting and Business Research*, **13**(5), 741–758.
- [3] Akerer, D., Hugonnier, J., & Jermann, U. (2025). Perpetual futures pricing. *Mathematical Finance*, 1–17. <https://doi.org/10.1111/mafi.70018> (NBER Working Paper 32936, 2024; SSRN 4603820, 2023.)
- [4] Adam, N. J., & Thomas, A. (2004). *Islamic Bonds: Your Guide to Issuing, Structuring and Investing in Sukuk*. Euromoney Books, London.
- [5] Ahmed, S., & Bhuyan, R. (2026). The interest parameter in perpetual futures: Shariah analysis and empirical evidence from centralized and decentralized exchanges. *SSRN Working Paper 6322778*. https://papers.ssrn.com/sol3/papers.cfm?abstract_id=6322778
- [6] Ahmed, S., & Bhuyan, R. (2026). The κ -yield curve: Construction, estimation methodology, and real-data evidence from sovereign sukuk. *SSRN Working Paper 6322938*. https://papers.ssrn.com/sol3/papers.cfm?abstract_id=6322938
- [7] Ahmed, S., & Bhuyan, R. (2026). The jurisprudential boundaries of an interest-free substrate: Riba, gharar, maysir, and qabd in cash-settled synthetic instruments. Working paper (companion research note).
- [8] Bank for International Settlements. (2024). *OTC Derivatives Statistics at End-December 2023*. BIS Quarterly Review, March 2024. <https://www.bis.org/statistics/derstats.htm>
- [9] Bielecki, T. R., & Rutkowski, M. (2002). *Credit Risk: Modeling, Valuation and Hedging*. Springer-Verlag, Berlin. <https://doi.org/10.1007/978-3-662-04821-4>
- [10] Billah, M. M. (2007). *Islamic Insurance (Takaful)*. Sweet & Maxwell Asia, Petaling Jaya, Malaysia.
- [11] Billah, M. M. (2023). Actuarial valuation (pricing) of takaful products: A Malaysian experience. *Journal of Islamic Finance*, **12**(2), 1–18.
- [12] Black, F., & Cox, J. C. (1976). Valuing corporate securities: Some effects of bond indenture provisions. *Journal of Finance*, **31**(2), 351–367.
- [13] Black, F., & Scholes, M. (1973). The pricing of options and corporate liabilities. *Journal of Political Economy*, **81**(3), 637–654.
- [14] Brémaud, P. (1981). *Point Processes and Queues: Martingale Dynamics*. Springer-Verlag, New York.
- [15] Cakir, S., & Raei, F. (2007). Sukuk vs. Eurobonds: Is there a difference in value-at-risk? *IMF Working Paper WP/07/237*. International Monetary Fund, Washington, D.C.
- [16] Chapra, M. U. (1985). *Towards a Just Monetary System*. The Islamic Foundation, Leicester.
- [17] Cox, J. C., Ingersoll, J. E., & Ross, S. A. (1985). A theory of the term structure of interest rates. *Econometrica*, **53**(2), 385–408.
- [18] Duffie, D., & Lando, D. (2001). Term structures of credit spreads with incomplete accounting information. *Econometrica*, **69**(3), 633–664.
- [19] Duffie, D., & Singleton, K. J. (1999). Modeling term structures of defaultable bonds. *Review of Financial Studies*, **12**(4), 687–720.
- [20] Duffie, D., & Singleton, K. J. (2003). *Credit Risk: Pricing, Measurement, and Management*. Princeton University Press, Princeton.
- [21] Duffie, D., Pedersen, L. H., & Singleton, K. J. (2003). Modeling sovereign yield spreads: A case study of Russian debt. *Journal of Finance*, **58**(1), 119–159.
- [22] El-Gamal, M. A. (2006). *Islamic Finance: Law, Economics, and Practice*. Cambridge University Press, Cambridge.
- [23] Elliott, R. J., Jeanblanc, M., & Yor, M. (2000). On models of default risk. *Mathematical Finance*, **10**(2), 179–195.
- [24] European Central Bank. (2019). *Is There a Zero Lower Bound? The Effects of Negative Policy Rates on Banks and Firms*. ECB Working Paper No. 2289. <https://www.ecb.europa.eu/pub/pdf/scpwps/ecb.wp2289~1a3c04db25.en.pdf>
- [25] Harrison, J. M., & Kreps, D. M. (1979). Martingales and arbitrage in multiperiod securities markets. *Journal of Economic Theory*, **20**(3), 381–408.
- [26] Hassan, M. K., Paltrinieri, A., Dreassi, A., & Haddad, A. (2021). Sukuk and bond spreads. *Journal of Economics and Finance*, **45**(3), 529–543. <https://doi.org/10.1007/s12197-021-09545-9>

- [27] Htay, S. N. N., & Salman, S. A. (2012). Optimal pricing for participation in takaful. *Asian Social Science*, **8**(7), 135–142.
- [28] Htay, S. N. N., Arif, M., Soualhi, Y., Zaharin, H. R., & Shaugee, I. (2012). *Accounting, Auditing and Governance for Takaful Operations*. John Wiley & Sons, Singapore.
- [29] International Islamic Financial Market. (2023). *Sukuk Report: A Comprehensive Study of the Global Sukuk Market* (13th ed.). IIFM, Manama.
- [30] Iqbal, M., & Molyneux, P. (2005). *Thirty Years of Islamic Banking: History, Performance, and Prospects*. Palgrave Macmillan, London.
- [31] Iqbal, Z., & Mirakhor, A. (2011). *An Introduction to Islamic Finance: Theory and Practice* (2nd ed.). John Wiley & Sons, Singapore.
- [32] Islamic Financial Services Board. (2024). *Islamic Financial Services Industry Stability Report 2024*. IFSB, Kuala Lumpur.
- [33] Jarrow, R. A., & Turnbull, S. M. (1995). Pricing derivatives on financial securities subject to credit risk. *Journal of Finance*, **50**(1), 53–85.
- [34] Jeanblanc, M., & Rutkowski, M. (2000). Modelling of default risk: An overview. In *Mathematical Finance: Theory and Practice*, Higher Education Press, Beijing, pp. 171–269.
- [35] Jeanblanc, M., Yor, M., & Chesney, M. (2009). *Mathematical Methods for Financial Markets*. Springer-Verlag, London. <https://doi.org/10.1007/978-1-84628-737-4>
- [36] Jobst, A. A. (2007). The economics of Islamic finance and securitization. *Journal of Structured Finance*, **13**(1), 6–27.
- [37] Jobst, A. A., Kunzel, P., Mills, P., & Sy, A. (2008). Islamic bond issuance: What sovereign debt managers need to know. *International Journal of Islamic and Middle Eastern Finance and Management*, **1**(4), 330–344.
- [38] Kahf, M. (1994). Time value of money and discounting in Islamic perspective: Re-visited. *Review of Islamic Economics*, **3**(2), 31–38.
- [39] Kamali, M. H. (2007). *Islamic Commercial Law: An Analysis of Futures and Options*. The Islamic Texts Society, Cambridge.
- [40] Khan, F. (2010). How ‘Islamic’ is Islamic banking? *Journal of Economic Behavior & Organization*, **76**(3), 805–820.
- [41] Khan, N., Kchouri, B., Yattoo, N. A., Kräussl, Z., Patel, A., & State, R. (2022). Tokenization of sukuk: Ethereum case study. *Global Finance Journal*, **51**, 100539. <https://doi.org/10.1016/j.gfj.2020.100539>
- [42] Khan, T., & Ahmed, H. (2001). Risk management: An analysis of issues in Islamic financial industry. *IRTI Occasional Paper* No. 5, Islamic Development Bank.
- [43] Lando, D. (1998). On Cox processes and credit risky securities. *Review of Derivatives Research*, **2**(2–3), 99–120.
- [44] Leland, H. E. (1994). Corporate debt value, bond covenants, and optimal capital structure. *Journal of Finance*, **49**(4), 1213–1252.
- [45] Meera, A. K. M., & Larbani, M. (2009). Ownership effects of fractional reserve banking: An Islamic perspective. *Humanomics*, **25**(2), 101–116.
- [46] Merton, R. C. (1974). On the pricing of corporate debt: The risk structure of interest rates. *Journal of Finance*, **29**(2), 449–470.
- [47] Moody’s Investors Service. (2024). *Sovereign Default and Recovery Rates* (annual study series). Moody’s Investors Service, New York.
- [48] Protter, P. E. (2005). *Stochastic Integration and Differential Equations* (2nd ed.). Springer-Verlag, Berlin.
- [49] Reboredo, J. C., Ugolini, A., & Al-Yahyaee, K. H. (2021). Yield spread determinants of sukuk and conventional bonds. *Economic Modelling*, **105**, 105664. <https://doi.org/10.1016/j.econmod.2021.105664>
- [50] Schönbucher, P. J. (2003). *Credit Derivatives Pricing Models: Models, Pricing and Implementation*. John Wiley & Sons, Chichester.
- [51] SeekersGuidance. (2025, October 2). Are crypto perpetual contracts still haram if there’s no interest, and the asset price tracks? Answered by Shaykh Muhammad Carr; approved by Shaykh Faraz Rabhani. <https://seekersguidance.org/answers/transactions/>
- [52] Shiller, R. J. (1993). *Macro Markets: Creating Institutions for Managing Society’s Largest Economic Risks*. Oxford University Press, Oxford.
- [53] Siddiqi, M. N. (2004). *Riba, Bank Interest and the Rationale of Its Prohibition*. Islamic Research and Training Institute, Jeddah.

- [54] Uddin, M. H., Kabir, S. H., Hossain, M. S., Wahab, N. A., & Liu, J. (2022). Why do sukuks (Islamic bonds) need a different pricing model? *International Journal of Finance & Economics*, **27**(2), 2210–2234.
<https://doi.org/10.1002/ijfe.2269>
- [55] Usmani, M. T. (2002). *An Introduction to Islamic Finance*. Kluwer Law International, The Hague.

EDITORIAL NOTE TO THIS PRINTING

Paper 2a · The κ -Yield Curve

Working paper, revised July 2026 · SSRN 6322938

Role in the programme. The first instrument: a benchmark curve built from the primitive instead of from a riskless rate. **What it establishes.** A κ term structure is recoverable from the sukuk cross-section with no riskless benchmark as input, and it orders issuers by credit quality. **Corrected or extended later.** The “benchmark-free wedge” was reframed in July 2026 as the reserve currency’s time-value and premium layer made visible, not a measured default hazard of the United States; the slope inference is stated at its honest significance. **Not established here.** A level interpretation of the wedge; the ordering and decomposition claims are the ones defended.

Replication for every number in this paper: DOI 10.5281/zenodo.21290557 · github.com/shehzadahmed-xx/kappa-replication

The κ -Yield Curve: Construction, Estimation
Methodology,
and Real-Data Evidence from Sovereign Sukuk
(Cross-Section and Constant-Maturity Dynamics)

Shehzad Ahmed¹ and Rafiqul Bhuyan^{*2}

¹Department of Finance, Independent University Bangladesh (IUB), Dhaka,
Bangladesh. 2120922@iub.edu.bd

²Department of Accounting and Finance, Alabama A&M University,
Normal, AL, USA. rafiqul.bhuyan@aamu.edu

June 2026

^{*}Corresponding author. Alabama A&M University, Normal, AL 35762, USA. E-mail:
rafiqul.bhuyan@aamu.edu.

Abstract

Islamic fixed-income markets lack a *riba*-free term structure of rates analogous to the risk-free yield curve used in conventional credit pricing. We address this gap in three parts. First, we define the κ -*yield curve* construct. The convergence intensity κ was introduced by Akerer, Hugonnier and Jermann [1] for perpetual futures and shown by Ahmed and Bhuyan [2] to be the *riba*-free analogue of the risk-free rate in Islamic credit pricing: the point estimator $\hat{\kappa} = s/(1 - \delta)$ connects a credit spread s to the recovery rate δ without any interest-bearing component. Second, we estimate it on a **real, instrument-level cross-section of 195 actively-quoted USD sukuk** (40 sovereign, 154 corporate, one municipal; Cbonds Bond Screener, live pull), computing per instrument the G -spread over the interpolated US Treasury curve — a real credit spread that, for these bullet sukuk, closely approximates the option-adjusted spread. The implied $\hat{\kappa} = s/(1 - \delta)$ rises as credit quality deteriorates, from a median $\approx 0.9\%$ at AAA to $\approx 4.4\%$ at single-B (increasing across the spectrum apart from one investment-grade inversion, the Saudi-dominated AA bucket pricing just above single-A), and the investment-grade sub-panel exhibits an *upward-sloping* term structure (+1.17 bp/yr pooled, +0.63 bp/yr once rating fixed effects absorb the credit ordering): positive but modest, as an honest real-data estimate should be. Third, we take the real data into the time domain: from daily Cbonds histories of 54 sukuk from five issuers spanning the rating spectrum (Malaysia, Saudi Arabia, Indonesia, Bahrain, Türkiye; 2018–2026) we construct roll-down-free *constant-maturity* κ series and estimate their dynamics directly. The *riba*-free intensity is mean-reverting: a long-run mean $\bar{\kappa}$ separates investment grade (≈ 1 –3% per annum) from high yield (≈ 5 –7%), matching the cross-section, with bias-corrected half-lives of 1.4–4.3 months at the well-sampled 7-year tenor and the Feller non-negativity condition satisfied in all twelve series. Extending to the *full* outstanding universe—the daily histories of all 409 outstanding USD sukuk (95 issuers, 2015–2026)—confirms both results at scale: the investment-grade term-structure slope is +3.4 bp/yr under issuer and month fixed effects on 165,557 bond-days, and identifiable constant-maturity dynamics extend to 21 issuers with the Feller condition holding in *all* of them. The contribution is thus the κ -yield-curve construct together with a real-data characterisation of its level, credit ordering, slope and dynamics on traded sukuk.

Keywords: sukuk; κ -yield curve; convergence intensity; Nelson–Siegel–Svensson; CIR dynamics; Islamic finance; GCC; *riba*-free rate

JEL codes: G12, G15, G28, F34, E43

Contents

1	Introduction	4
2	Theoretical Framework	5
2.1	The Convergence Intensity in the AHJ Framework	5
2.2	The κ -Yield Curve via Nelson–Siegel–Svensson	6
2.3	CIR Stochastic Dynamics for κ	8
2.4	GMM Identification	8
3	Data	9
3.1	Recovery Rate Assumption	9
4	Estimation Methodology	10
4.1	Step 1: Extracting $\hat{\kappa}$	11
4.2	Step 2: Mean-Reversion Dynamics of the Constant-Maturity Series	11
5	Results	11
5.1	Real Sovereign Sukuk Cross-Section: From Method to Result	12
5.2	Sovereign versus Corporate: Decomposing κ into Country and Issuer Credit .	16
5.3	Real κ Dynamics: Constant-Maturity Mean-Reversion	18
5.4	The κ -Index: A Cross-Sectional Benchmark Series	20
5.5	The Benchmark-Free Estimator and the Recovery Floor	22
6	Discussion	24
6.1	Practical Implications for Islamic Finance	24
6.2	Non-Negativity and the Negative-Rate Dilemma	25
7	Conclusion	26
A	Derivation of the Identification (Theorem 2.1)	29
B	GMM Moment Conditions	29
C	Country Coverage and Sukuk Universe	30

1 Introduction

Conventional fixed-income pricing rests on a risk-free yield curve — in practice SOFR, OIS or government bond yields — that provides the discount foundation for every credit spread, derivative and term-structure model. Islamic finance operates under the prohibition of *riba* (usury), which rules out any security whose return is contractually fixed as a function of a notional loan balance [11]. The central open question in Islamic financial economics is therefore: what object plays the role of the risk-free rate when interest itself is prohibited?

The present paper develops the methodology to answer this question. Akerer, Hugonnier and Jermann [1] (hereafter AHJ) introduce the *convergence intensity* κ for perpetual futures and prove (their Theorem 3) that when the interest parameter is set to zero ($\iota = 0$) the unique no-arbitrage price is the risk-neutral expectation of the spot sampled at an exponentially distributed stopping time of mean $1/\kappa$. Ahmed and Bhuyan [2] formalise this observation into a theorem: κ is exactly the *riba*-free analogue of the risk-free rate for credit pricing, with the identification equation $\hat{\kappa} = s/(1 - \delta)$ linking the observed sukuk spread s to the recovery rate δ .

Despite this theoretical clarity, no study has yet specified an estimation pipeline for the κ -yield curve, tested whether the estimator recovers the term structure properties expected of a genuine pricing rate, or verified that its stochastic dynamics satisfy the non-negativity (Feller) condition that distinguishes it from conventional short rates — which have famously turned negative in many developed markets during 2014–2022. We fill all three gaps.

What we do. We assemble a real, instrument-level cross-section of **195 actively-quoted USD sukuk** (40 sovereign, 154 corporate, one municipal) from the Cbonds Bond Screener, spanning the rating spectrum (AAA to single-B) and maturities from under one year to 2050. For each instrument we read the traded G -spread over the maturity-interpolated US Treasury curve — the instrument-level sukuk OAS for these bullet bonds. For the time-series analysis we add roll-down-free constant-maturity $\hat{\kappa}$ histories built from daily Cbonds G -spread series of 54 sukuk from five issuers (2018–2026).

On these real sukuk we then (i) extract $\hat{\kappa}$ for each instrument from its observed G -spread; (ii) characterise the cross-sectional κ -curve by credit rating and maturity; and (iii) estimate Cox–Ingersoll–Ross (CIR) [9] mean-reversion for roll-down-free constant-maturity $\hat{\kappa}$ series via their AR(1)/OU representation, testing the Feller condition.

Main findings. First, across the 195-bond cross-section the convergence intensity broadly increases as credit quality deteriorates: median $\hat{\kappa} \approx 0.9\%$ at AAA, $\approx 2.3\%$ at AA, $\approx 2.0\%$ at single-A, $\approx 3.4\%$ at BBB and $\approx 4.4\%$ at single-B (Section 5) — rising across the spectrum

apart from the AA bucket, which (dominated by frequent issuer Saudi Arabia) prices just above single-A. This is the credit ordering the κ -yield curve predicts, read directly from the sukuk market.

Second, the investment-grade κ -term-structure is *upward sloping*: the G -spread rises with maturity at +1.17 bp/yr pooled, +0.63 bp/yr with rating fixed effects ($n = 164$ AAA–BBB bonds). The slope is positive but modest.

Third, from constant-maturity histories the *riba*-free intensity is genuinely mean-reverting, with a long-run mean $\bar{\kappa}$ separating investment grade (≈ 1 –3%) from high yield (≈ 5 –7%) and bias-corrected half-lives of 1.4–4.3 months at the well-sampled 7-year tenor — faster than the multi-quarter-to-multi-year half-lives documented for conventional sovereign spreads in the sovereign credit literature [4]. The Feller non-negativity condition $2\alpha\bar{\kappa} \geq \nu^2$ holds in all twelve series, so κ stays non-negative under mean reversion, unlike the euro-area and Japanese overnight benchmarks (EONIA/ESTR, TONAR) and policy rates, which turned negative during 2014–2022.

Contribution to the literature. The existing literature on sukuk pricing can be divided into four strands: (i) structural analyses of the legal and Shariah architecture [11]; (ii) empirical comparisons of sukuk and conventional bond spreads [7]; (iii) theoretical no-arbitrage pricing at $\iota = 0$ [1, 2]; and (iv) term-structure models for Islamic rates more broadly [8]. Our paper bridges strand (iii) to an estimation methodology, introducing the κ -yield curve and characterising its level, credit ordering, slope and dynamics on a real 195-bond cross-section of traded sukuk plus constant-maturity histories for five issuers.

Organisation. Section 2 presents the theoretical framework. Section 3 describes data sources and construction. Section 4 details estimation methodology. Section 5 reports all empirical results. Section 6 discusses practical implications. Section 7 concludes. All data and code are publicly available (see Data Availability Statement).

2 Theoretical Framework

2.1 The Convergence Intensity in the AHJ Framework

Akerer, Hugonnier and Jermann [1] develop a no-arbitrage pricing framework for perpetual futures (and, as an extension, everlasting options) — derivatives sustained by a periodic premium over a perpetual horizon rather than a fixed expiry. Their key parameters are: the risk-free rate r , the interest parameter ι (the interest-bearing component of the funding rate,

which enters the AHJ random-time representation by shifting the effective stopping intensity to $\kappa - \iota$), and the convergence intensity κ — the rate at which the contract price converges to its terminal payoff — which, in the credit interpretation of Ahmed and Bhuyan [2], governs how quickly the reference entity approaches default. When the economy operates under a riba prohibition, Ahmed & Bhuyan [2] establish the following:

Theorem 2.1 (Riba-free credit identification; Ahmed & Bhuyan 2026). *Under the riba prohibition ($\iota = 0$), a fixed-maturity survival claim paying par $\mathcal{P}$ at maturity T if no default occurs, and recovery $\delta\mathcal{P}$ at default, has the no-arbitrage price*

$$V(t) = \mathcal{P} e^{-\kappa(1-\delta)(T-t)}, \quad (1)$$

so that its continuously-compounded spread is $s = \kappa(1 - \delta)$, equivalently

$$\kappa = \frac{s}{1 - \delta}, \quad (2)$$

with s the observed yield spread above the default-free benchmark (proxied empirically by the interpolated US Treasury curve; Section 3), $\delta \in [0, 1)$ the recovery rate, and $\kappa \geq 0$ the convergence intensity, identified with the default intensity λ by the Separation-Theorem isomorphism of Ahmed & Bhuyan [2].

Equation (2) is the key identification equation. It establishes κ as the riba-free pricing rate: given any observed spread s and recovery assumption δ , one can recover κ without reference to any interest-bearing instrument.

Remark 2.1. The identification (2) is derived from the no-arbitrage condition that holds in the AHJ framework when $\iota = 0$. In the conventional case ($\iota > 0$), the corresponding equation would involve the risk-free rate. Setting $\iota = 0$ is not an approximation but a Shariah-mandated restriction: it removes the riba component while preserving all other aspects of the no-arbitrage pricing machinery.

2.2 The κ -Yield Curve via Nelson–Siegel–Svensson

The cross-section of $\hat{\kappa}(\tau)$ at a grid of maturities — e.g. $\tau \in \{2, 3, 5, 7, 10, 15, 20\}$ years in the general construct; the empirical implementation of Section 5 uses the well-sampled tenors $\{3, 5, 7\}$ — defines the κ -yield curve. We model this object using the Nelson–Siegel–Svensson functional form [12, 14]:

$$\kappa(\tau; \boldsymbol{\beta}, \boldsymbol{\tau}_0) = \beta_0 + \beta_1 \frac{1 - e^{-\tau/\tau_1}}{\tau/\tau_1} + \beta_2 \left[\frac{1 - e^{-\tau/\tau_1}}{\tau/\tau_1} - e^{-\tau/\tau_1} \right] + \beta_3 \left[\frac{1 - e^{-\tau/\tau_2}}{\tau/\tau_2} - e^{-\tau/\tau_2} \right], \quad (3)$$

where $\boldsymbol{\beta} = (\beta_0, \beta_1, \beta_2, \beta_3)$ are level, slope, and two curvature parameters, and $\tau_1, \tau_2 > 0$ are the corresponding decay constants. The NSS form is the standard tool for central bank yield curve estimation [6] and inherits all its well-known identifiability and smoothness properties.

The economic interpretation of the NSS parameters in the κ -yield curve context is as follows:

- $\beta_0 > 0$: long-run (infinite-maturity) convergence intensity, analogous to the long-run risk-free rate in conventional term structure models.
- $\beta_1 < 0$: slope factor. A negative slope factor makes the κ -yield curve upward-sloping, consistent with conventional term-premium theory; the real investment-grade sukuk cross-section confirms this directly (Table 2, Section 5).
- β_2, β_3 : curvature factors capturing hump-shaped deviations from a monotone profile.
- τ_1, τ_2 : shape parameters. We fix $\tau_1 = 1.50$ and $\tau_2 = 5.00$ in line with central bank practice for short panels; since the full NSS form is not fitted to the single cross-section (Section 4), no τ -robustness exercise arises here.

Proposition 2.1 (Upward slope of the κ -yield curve). *If $\beta_1 < 0$, then $\kappa(0) = \beta_0 + \beta_1 < \beta_0 = \lim_{\tau \rightarrow \infty} \kappa(\tau)$, so the curve ends above where it starts; and for $|\beta_2|, |\beta_3|$ sufficiently small relative to $|\beta_1|$ (given τ_1, τ_2), $\kappa(\tau)$ defined by (3) is monotonically increasing in maturity τ .*

Proof. At $\tau = 0$: the factors $(1 - e^{-\tau/\tau_j})/(\tau/\tau_j) \rightarrow 1$ and $e^{-\tau/\tau_j} \rightarrow 1$, so the curvature terms vanish and $\kappa(0) = \beta_0 + \beta_1$. As $\tau \rightarrow \infty$: all exponential loading factors decay to zero, so $\kappa(\infty) = \beta_0$. For the monotonicity claim, the slope-factor loading $g_1(\tau) = (1 - e^{-\tau/\tau_1})/(\tau/\tau_1)$ decreases strictly from 1 to 0, contributing $-\beta_1 g_1'(\tau) > 0$ to $\kappa'(\tau)$ at every τ , while the curvature-loading derivatives are bounded; hence for $|\beta_2|, |\beta_3|$ small enough relative to $|\beta_1|$ the slope term dominates uniformly and $\kappa'(\tau) > 0$ for all $\tau > 0$. $\square$

We do not fit the full NSS form to the single mid-2026 cross-section (Section 4); we test only its directly observable implication, the sign of the maturity slope predicted by Proposition 2.1.

2.3 CIR Stochastic Dynamics for κ

Following Ahmed & Bhuyan [2], we model the time-series evolution of the level of κ (in the empirical implementation, a constant-maturity point on the curve) as a Cox–Ingersoll–Ross process [9]:

$$d\kappa_t = \alpha(\bar{\kappa} - \kappa_t)dt + \nu\sqrt{\kappa_t}dW_t, \quad (4)$$

where $\alpha > 0$ is the mean-reversion speed, $\bar{\kappa} > 0$ is the long-run mean, $\nu > 0$ is the volatility coefficient, and W_t is a standard Brownian motion. Among the classic one-factor affine short-rate models, the constant-volatility Vasicek specification [15] admits negative values and is therefore inadmissible for a riba-free intensity; the CIR process is the natural candidate because: (i) it is non-negative for all t (given $\kappa_0 > 0$), and remains strictly positive whenever the Feller condition $2\alpha\bar{\kappa} \geq \nu^2$ holds; (ii) it is mean-reverting, consistent with the economic requirement that credit spreads do not drift to infinity; and (iii) it has a tractable transition density (non-central chi-squared), with closed-form conditional moments.

Definition 2.1 (Feller Condition). Given $\kappa_0 > 0$, the CIR process (4) satisfies $\kappa_t > 0$ almost surely for all $t \geq 0$ if and only if the Feller condition

$$2\alpha\bar{\kappa} \geq \nu^2 \quad (5)$$

holds.

The Feller condition is the formal guarantee that the κ -yield curve stays non-negative — precisely the property that conventional short rates violated during the European Central Bank and Bank of Japan’s negative rate periods (2014–2022). We test the Feller condition empirically in every constant-maturity series in Section 5.

2.4 GMM Identification

We state the estimator in full generality, for a cross-issuer panel of constant-maturity κ paths. On the present single-series-per-issuer data we run its exactly-identified special case — the AR(1)/OU regression of Section 4; the general GMM system below is the procedure the pipeline is built to consume once a multi-issuer constant-maturity panel is assembled (Section 7).

Let $h = 1/12$ denote the monthly step size. Under the CIR dynamics (4), the conditional

moments are:

$$\mathbb{E}[\kappa_{t+h}|\kappa_t] = \kappa_t e^{-\alpha h} + \bar{\kappa}(1 - e^{-\alpha h}), \quad (6)$$

$$\text{Var}[\kappa_{t+h}|\kappa_t] = \frac{\nu^2}{\alpha} \kappa_t (e^{-\alpha h} - e^{-2\alpha h}) + \frac{\nu^2 \bar{\kappa}}{2\alpha} (1 - e^{-\alpha h})^2. \quad (7)$$

We stack the moment conditions implied by (6)–(7) with instruments $\mathbf{z}_t = (1, \kappa_{t-1}, \kappa_{t-2})$ and estimate $\boldsymbol{\theta} = (\alpha, \bar{\kappa}, \nu)$ by two-step efficient GMM with a Newey–West HAC variance estimator [13]. This is the general identification result for a cross-section of constant-maturity κ paths; with the single constant-maturity history per issuer available in real traded sukuk, the empirical implementation (Step 2 of the methodology, Section 4) uses its exactly-identified AR(1)/OU special case, which estimates the same drift $(\alpha, \bar{\kappa})$.

3 Data

We use two Cbonds datasets, both pulled under an authenticated account (972307) from the Cbonds Bond Screener (</api/bonds/search/>).

Cross-section. Every outstanding, actively-quoted USD-denominated sukuk on Cbonds’ **sukuk** classification: **195 instruments** (40 sovereign, 154 corporate, one municipal) across the rating spectrum (AAA to single-B) and maturities from under one year to 2050. For each instrument Cbonds reports the indicative yield and the *G-spread* — the yield in excess of the maturity-interpolated US Treasury curve. For these bullet (option-free) sukuk the *G-spread* approximates the *Z-spread* closely (the two coincide to first order for option-free bullets), so it serves as the instrument-level sukuk OAS. We apply the identification $\hat{\kappa} = s/(1 - \delta)$ to each observed spread.

Dynamics. For the time-series analysis we use roll-down-free *constant-maturity* $\hat{\kappa}$ histories built from daily Cbonds *G-spread* series of 54 sukuk from five issuers spanning the rating spectrum (Malaysia, Saudi Arabia, Indonesia, Bahrain, Türkiye; 2018–2026).

3.1 Recovery Rate Assumption

We set $\delta = 0.60$ (60% recovery) uniformly across all countries and maturities. Recovery conventions vary: the sovereign CDS market quotes a 25% convention, while historical issuer-weighted sovereign recoveries average roughly 50–55% in the rating agencies’ sovereign default

studies [5]; we adopt 0.60, toward the upper end, to reflect the senior and frequently asset-backed structure of the sukuk in the panel, and apply the same δ to the corporate panel for comparability (corporate-specific recovery refinement is left to the production benchmark). This implies $1 - \delta = 0.40$ and $\hat{\kappa} = 2.5 \times s$ (in the same units). Recovery sensitivity is a pure rescaling: $\delta \in \{0.40, 0.50, 0.70\}$ multiplies every $\hat{\kappa}$ by a common factor (Section 5).

Data caveats, consolidated. Four limitations of the data foundation deserve one honest paragraph rather than scattered footnotes. *Single vendor:* all instrument-level data derive from one provider (Cbonds); the cross-instrument checks against independently sourced sovereign CDS (correlations 0.97 pooled, 0.989–0.998 per country in the companion studies) are the operative mitigation. *Snapshot elements:* certain aggregates were verified against the vendor’s live statistics pages at pull time and are provenance-tagged rather than re-derivable from the archived responses; all bond-level inputs are archived in the replication package. *Quote quality and survivorship:* an “actively quoted” panel of illiquid instruments inherits matrix-pricing and survivorship biases typical of this market; the full-universe pulls (including matured and defaulted issues) bound the survivorship direction. *Magnitudes near liquidity noise:* the fixed-effects investment-grade slope (+0.63bp/yr, s.e. 0.53, $t \approx 1.2$) is not statistically distinguishable from zero on the snapshot alone, and is small enough that a liquidity–maturity artifact cannot be excluded; we rest no structural claim on its magnitude. The statistically decisive statement lives in the panel: on the full-universe bond-day panel the issuer-plus-month fixed-effects slope is +3.4bp/yr here (Table 3, reported without inference), and the out-of-sample companion re-estimates it on a larger, fresher panel with issuer-clustered standard errors at +2.85bp/yr, $t = 4.1$ – 4.7 — the sign and approximate magnitude are significant there, not in the single-date cross-section.

4 Estimation Methodology

The pipeline has two steps: a model-free extraction of $\hat{\kappa}$ from each observed spread (Step 1), and estimation of its constant-maturity CIR mean-reversion through the AR(1)/OU representation (Step 2). The cross-sectional shape prediction of Proposition 2.1 — an upward-sloping curve when $\beta_1 < 0$ — is tested on the real cross-section by its directly observable implication, the sign of a linear $\hat{\kappa}$ -on-maturity slope, rather than by fitting a full per-issuer NSS curve, which a single mid-2026 snapshot cannot identify stably.

4.1 Step 1: Extracting $\hat{\kappa}$

For each observation (i, τ, t) — country i , maturity τ , month t — we compute the point estimate directly from the identification equation (2):

$$\hat{\kappa}_{i,\tau,t} = \frac{s_{i,\tau,t}}{1 - \delta_i}, \quad (8)$$

where $s_{i,\tau,t}$ is the observed spread in decimal form and δ_i is the country-specific recovery rate (set to 0.60 in the baseline). This step involves no estimation uncertainty under the structural model: given δ , $\hat{\kappa}$ is a bijective transformation of the spread. The uncertainty in $\hat{\kappa}$ therefore comes entirely from measurement error in s and uncertainty about δ .

4.2 Step 2: Mean-Reversion Dynamics of the Constant-Maturity Series

For each issuer i we work with a single *constant-maturity* $\hat{\kappa}_{i,t}$ history — a fixed point on the curve tracked through time, not an average across maturities — and estimate its mean-reversion through the discrete-time Ornstein–Uhlenbeck (AR(1)) representation of the CIR process,

$$\hat{\kappa}_{i,t} = c_i + \phi_i \hat{\kappa}_{i,t-1} + \varepsilon_{i,t}, \quad \hat{\alpha}_i = -\frac{\ln \hat{\phi}_i}{\Delta t}, \quad \hat{\bar{\kappa}}_i = \frac{\hat{c}_i}{1 - \hat{\phi}_i}, \quad (9)$$

estimated on monthly observations ($\Delta t = 1/12$), so $\hat{\alpha}_i$ is the annualised mean-reversion speed. This is the natural finite-sample estimator of the CIR drift $(\alpha, \bar{\kappa})$ when a single series is observed: with one constant-maturity path per issuer the cross-maturity instruments of the GMM system (Section 2.4) are not identified, and the AR(1)/OU discretisation is its exactly-identified special case for the drift. The CIR diffusion coefficient ν_i is then recovered from the *second* moment of the residuals: under the Euler discretisation $\text{Var}(\varepsilon_{i,t} \mid \hat{\kappa}_{i,t-1}) = \nu_i^2 \hat{\kappa}_{i,t-1} \Delta t$, which we match in aggregate at the long-run mean, $\hat{\nu}_i^2 = \widehat{\text{Var}}(\varepsilon_i) / (\hat{\bar{\kappa}}_i \Delta t)$. We apply Kendall’s small-sample bias correction to $\hat{\phi}_i$, report Newey–West (HAC) standard errors, and test the Feller condition (5), $2\hat{\alpha}_i \hat{\bar{\kappa}}_i \geq \hat{\nu}_i^2$, for each series. The half-life is $12 \ln(2) / \hat{\alpha}_i$ months.

5 Results

We present the cross-section first, then the constant-maturity dynamics.

5.1 Real Sovereign Sukuk Cross-Section: From Method to Result

The substantive content of the κ -yield curve is the credit-spread-implied $\hat{\kappa}$ *level* and its cross-sectional ordering by rating and maturity. We test it directly on the cleanest available object: a real, instrument-level cross-section of *sukuk* credit spreads. We present it in two samples: a hand-checked 61-bond sovereign core (this subsection), then the full 195-bond sovereign-plus-corporate universe that the abstract and introduction headline (“Widening to the full *sukuk* universe,” below). The two are separate Cbonds pulls taken with different screener filters on different dates — the 195-bond snapshot’s stricter active-quote filter happens to exclude several high-yield sovereigns present in the core (e.g. Bahrain, Egypt, Pakistan, Türkiye; Appendix C) — so the sovereign core is *not* a subset of the headline snapshot, and the two samples corroborate the same regularities independently rather than mechanically. Those sovereigns are not absent from the data programme: the full-universe API pull behind the panel-robustness and full-universe-dynamics analyses (all 409 outstanding USD *sukuk*, including Bahrain, Türkiye, Oman, Pakistan and Egypt; Sections 5.1 and 5.3) covers them, and the five-issuer constant-maturity dynamics use Bahrain and Türkiye directly.

Data. Using the Cbonds bond database we extract every outstanding, actively-quoted USD-denominated *sukuk* issued by a sovereign — filtered on Cbonds’ own *sukuk* classification, sovereign issuer type, and outstanding status. After dropping non-traded retail tranches and instruments without a current quote, the panel comprises **61 *sukuk* from 12 sovereign issuers** (Bahrain, Benin, Egypt, Hong Kong, Indonesia, Malaysia, Oman, Pakistan, Philippines, Qatar, Saudi Arabia and Türkiye), spanning ratings from AA (Qatar, Saudi Arabia) to CCC (Egypt, Pakistan) and maturities from under one year to 2050. For each instrument Cbonds reports the indicative yield to maturity and the *G-spread*: the yield in excess of the maturity-interpolated US Treasury curve. For a bullet (option-free) bond — which these straight sovereign *sukuk* are — the G-spread equals the Z-spread and coincides with the option-adjusted spread, so this is the instrument-level *sukuk* OAS. We apply the identification of Ahmed & Bhuyan [2], $\hat{\kappa} = s/(1 - \delta)$ with $\delta = 0.60$, to each observed spread.

Results. Three predictions of the κ -yield-curve framework hold on observed *sukuk* (Figure 2, Table 4):

1. **Level.** The median G-spread is 78 bp, implying median $\hat{\kappa} = 0.019$ per annum, with an interquartile range that maps to $\hat{\kappa} \in [0.012, 0.052]$. This is squarely the order of magnitude (10^{-2} p.a.) the credit channel predicts.
2. **Rating ordering.** Median $\hat{\kappa}$ separates cleanly across the investment-grade/high-yield boundary — IG buckets at 0.009–0.019 per annum (A/BBB/AA) versus high-yield at

0.049 (BB), 0.054 (B) and 0.074 (CCC) — a factor of three to six gap, and is strictly increasing *within* high yield (BB < B < CCC). The one wrinkle is inside investment grade, where the AA bucket (≈ 0.018 , dominated by frequent issuer Saudi Arabia) prices marginally above the four-bond A bucket (Malaysia, ≈ 0.009); this is an issuer-supply effect, not a break in the credit ordering, which is exactly what the framework requires of a genuine credit intensity.

3. **Term-structure slope.** Restricting to the investment-grade sub-panel (AA/A/BBB), a linear fit of G-spread on years-to-maturity has a *positive* slope of +1.4 bp/yr (intercept 60 bp) on this 61-bond sovereign-only subset — an upward-sloping κ -curve, consistent with the positive investment-grade slope on the wider 195-bond panel (Table 2).

This is the paper’s headline real-data result: the central cross-sectional predictions of the κ -yield curve — the $\sim 10^{-2}$ per-annum level, the rating ordering, and the upward slope — are properties of the *observed* sovereign-sukuk market, not merely of a calibrated construction.

Widening to the full sukuk universe. The same three regularities hold when corporates are added. A fresh live pull (Cbonds, account 972307) widens the cross-section to **195 actively-quoted USD sukuk** — 40 sovereign, 154 corporate, one municipal — across the rating spectrum. On this sample the convergence intensity again increases with rating (Table 1: $\hat{\kappa}$ from 0.93% at AAA to 4.4% at single-B, with the AA bucket the one investment-grade inversion), the investment-grade κ -term-structure slopes *upward* (Table 2: +1.17 bp/yr pooled, +0.63 bp/yr with rating fixed effects, $n = 164$ AAA–BBB bonds; the fixed-effects estimate is not individually significant on the snapshot, $t \approx 1.2$ — the significance lives in the panel, below), and the pooled κ -curve (Figure 1) shows both features directly from traded sukuk. The slope is positive but modest, consistent with the sovereign-only subset above.

Panel robustness of the slope. The cross-sectional slope is a single mid-2026 snapshot. We strengthen it to a panel by pulling the full daily Cbonds G -spread history (2015–2026) of *every* outstanding USD sukuk—409 instruments across 95 issuers and 17 countries (Cbonds, account 972307)—and re-estimating the maturity slope on 165,557 bond-days (Table 3). The upward slope is confirmed across the full universe and is firmer than the snapshot: the investment-grade maturity slope is +3.4 bp/yr under issuer *and* month fixed effects (and +3.0 bp/yr pooled), so the positive term premium is a stable feature of the traded sukuk market across a decade and 95 issuers, not an artifact of the snapshot date or of the five-issuer dynamics subset. The within-issuer slope (+5.1 bp/yr) exceeds the pooled slope

Rating	n	median $\hat{\kappa}$ (%)	median G (bp)
AAA	9	0.93	37
AA	102	2.28	91
A	47	1.96	78
BBB	6	3.41	136
B	3	4.37	175
NR	28	5.60	224

Table 1: Convergence intensity $\hat{\kappa} = s/(1 - \delta)$ by credit rating on the *real* USD sukuk cross-section (195 actively-quoted bonds, sovereign + corporate; Cbonds, account 972307; $\delta = 0.60$). The five rated buckets ($n = 167$) plus the unrated tail total the full 195. The $\hat{\kappa}$ ordering increases with rating across the rated spectrum, apart from the AA bucket (Saudi-dominated, $n = 102$) pricing just above single-A; NR is unrated, shown for completeness as the high-spread tail, not part of the ordering — read directly from the sukuk market, with no synthetic panel.

Specification	slope (bp/yr)	s.e.
Pooled IG (G on maturity)	+1.17	—
IG with rating fixed effects	+0.63	0.53

Table 2: Term structure of the *real* investment-grade sukuk G -spread ($n = 164$ AAA–BBB bonds, Cbonds). The maturity slope is positive — an upward-sloping investment-grade κ -term-structure — both pooled and with rating fixed effects, confirming on real sukuk the upward IG slope. All real data; no synthetic panel.

because pooling flattens the term structure—longer-maturity bonds are disproportionately issued by tighter-credit names—an effect the issuer fixed effect removes.

Specification	Pooled	Issuer FE	Issuer+month FE
All issuers (n=165,557)	+0.28	+5.63	+3.70
Investment grade (n=138,200)	+3.03	+5.09	+3.41

Table 3: Term structure of the *real* USD-sukuk G -spread as a *panel* of bond-days from the full outstanding universe (409 sukuk, **95 issuers**, daily Cbonds histories 2015–2026; account 972307), slope in bp/yr. The positive maturity slope of the single-date cross-section (Table 2) is confirmed on the full panel and survives issuer *and* month fixed effects. $\delta = 0.60$. Point estimates only — with daily bond-day observations the effective sample is far smaller than n suggests and naive standard errors would overstate precision; issuer-clustered inference on this specification (fresher, larger panel) is reported in the out-of-sample companion, which finds +2.85 bp/yr with $t = 4.1$ – 4.7 .




Figure 1: The real κ -curve: $\hat{\kappa} = s/(1 - \delta)$ against maturity for 195 actively-quoted USD sukuk (Cbonds, account 972307; $\delta = 0.60$), coloured by rating, with the pooled investment-grade fit. Two reconciliations with Table 2: the figure plots and fits in κ -units on the plotted $maturity \leq 15y$ subsample, so its slope (+9.4 bp of $\kappa/yr = +3.8$ bp of G-spread/yr at $\delta = 0.60$) is steeper than the full-maturity pooled G-spread slope of Table 2 (+1.17 bp/yr) — long bonds beyond 15y flatten the fit; and the per-bucket legend counts (e.g. AA $n = 95$, A $n = 44$) are the $\leq 15y$ subsample of the full buckets in Table 1 (AA $n = 102$, A $n = 47$).

What this does and does not establish. The cross-section is a single mid-2026 snapshot, not a transaction-level time series; it validates the *cross-sectional* content of the framework (levels, ordering, slope), and the companion time-series analysis of Section 5.3 supplies the *dynamic* content (mean-reversion, level, Feller) on real data. The G-spread benchmarks against the US Treasury curve rather than a swap/SOFR curve, so the implied $\hat{\kappa}$ embeds a small swap-spread component; and a handful of issuers (Egypt, Pakistan) sit at the distressed end where the constant- δ assumption is most fragile. None of these caveats touches the three qualitative regularities, which are robust across reasonable recovery choices ($\delta \in [0.4, 0.7]$ rescales every $\hat{\kappa}$ by a common factor and leaves the ordering and slope unchanged).




Figure 2: Real sovereign USD sukuk κ -curve (Cbonds, 61 bonds, 12 issuers, mid-2026). Each point is one outstanding sovereign USD sukuk; the vertical axis is the observed G-spread over the interpolated US Treasury curve (left) and the implied $\hat{\kappa} = s/(1 - \delta)$, $\delta = 0.60$ (right). Colours denote the issuer’s foreign-currency rating bucket. Investment-grade sukuk (AA/A/BBB) cluster at 30–110 bp with an upward-sloping term structure (dashed fit, +1.4 bp/yr); high-yield sukuk (BB/B/CCC) sit at 200–300 bp. The level ($\sim 10^{-2}$ p.a. $\hat{\kappa}$), the rating ordering, and the IG slope are all read directly from traded sukuk.

5.2 Sovereign versus Corporate: Decomposing κ into Country and Issuer Credit

The sovereign cross-section identifies the *country* component of κ . The same database resolves the *issuer* component: Cbonds lists 213 actively-quoted corporate USD sukuk (median G-spread 105 bp, implied $\hat{\kappa} = 0.026$ per annum), against the sovereign median of 78 bp ($\hat{\kappa} = 0.019$). Because a corporate sukuk embeds its sovereign’s credit plus an issuer-specific premium, the clean comparison is *within* an issuer’s home country: reading each sovereign’s κ -curve at the maturity of every domestic corporate sukuk and differencing. (The 213-bond corporate count comes from the broader corporate screener used for this decomposition pull; the headline 195-bond panel’s stricter active-quote filter retains 154 of these corporates.) The corporate–sovereign premium is *positive in every country with both* (Table 5), with a

Rating	n	Median G-spread (bp)	Median $\hat{\kappa}$ (yr ⁻¹)	Grade
AA	14	72	0.018	IG
A	4	37	0.009	IG
BBB	23	78	0.019	IG
BB	7	195	0.049	HY
B	10	215	0.054	HY
CCC	3	295	0.074	HY
All	61	78	0.019	

Table 4: Real sovereign USD sukuk cross-section by foreign-currency rating bucket (Cbonds, mid-2026; bucket = highest of the three agency ratings). $\hat{\kappa} = s/(1 - \delta)$ with $\delta = 0.60$ and s the observed G-spread-over-Treasury. Median $\hat{\kappa}$ separates cleanly across the investment-grade/high-yield boundary — IG buckets at 0.009–0.019 versus HY at 0.049–0.074 per annum — and is strictly increasing *within* high yield (BB < B < CCC). Within investment grade the AA bucket (≈ 72 bp, dominated by Saudi Arabia, a large and frequent issuer) prices slightly wider than the small four-bond A bucket (Malaysia, ≈ 37 bp): an issuer-supply/liquidity effect, not a break in the credit ordering. The AA median maturity is 3.7 years, so this is not a maturity artifact.

pooled median of +38 bp — a corporate credit premium of $\Delta\hat{\kappa} \approx +0.010$ per annum on top of the sovereign κ . The *riba*-free intensity therefore decomposes additively into a sovereign and a corporate credit component, each read from observed sukuk with no interest-bearing return component (the US Treasury curve enters only as the netting benchmark; see the caveat in Section 5.1) — exactly the decomposition a credit benchmark must support. (The corporate universe skews to banks and government-related entities — 85 of 213 are bank issuers, many state-linked — so this premium is conservative: it understates the credit gap for a purely private issuer.)

Country	Sov. med (bp)	Corp. med (bp)	Premium (bp)	$\Delta\hat{\kappa}$ (yr ⁻¹)
Saudi Arabia	72	96	+23	+0.006
Bahrain	215	241	+29	+0.007
Malaysia	37	64	+47	+0.012
Türkiye	195	310	+78	+0.020
Pooled	78	105	+38	+0.010

Table 5: Decomposing κ into country and issuer credit (Cbonds, mid-2026; $\delta = 0.60$). For each country with both sovereign and corporate USD sukuk, the corporate G-spread is differenced against the sovereign κ -curve read at the same maturity. The corporate–sovereign premium is positive throughout: corporate κ equals sovereign κ plus an issuer credit component, both free of any interest-bearing return component (the Treasury curve serves only as the netting benchmark).

Corporate dynamics. The corporate credit component is not only a level but a *process*. Applying the same roll-down-free constant-maturity construction we use for the sovereigns below (Section 5.3) to the corporate issuers with enough sukuk to span a curve, the corporate κ behaves exactly as the sovereign κ does. At the well-sampled 7-year tenor the long-run mean $\bar{\kappa}$ orders by credit quality — from 0.81% for the Islamic Development Bank (a AAA supranational, below even the tightest sovereign) through 2.2% for Saudi Electricity and Saudi Aramco to 3.2% for DP World — with bias-corrected half-lives of a few months (Saudi Electricity 2.5, DP World 4.5, both with α significant) and the Feller non-negativity condition satisfied for all five issuers estimated. The *riba*-free intensity is therefore mean-reverting, non-negative and credit-ordered in the *corporate* domain as well, on the same 10^{-2} -per-annum scale — confirming that κ is a single primitive spanning sovereign and corporate Islamic credit, not an artefact of any one issuer class.

5.3 Real κ Dynamics: Constant-Maturity Mean-Reversion

The cross-section establishes the static predictions on real data; the dynamics — mean-reversion speed, long-run level and the Feller property — require a time series, which a single-date snapshot cannot provide. We now estimate them on real constant-maturity sukuk series.

Data and the roll-down problem. From Cbonds we retrieve the full daily history of the model-implied G-spread for every USD sovereign sukuk of the five issuers with enough outstanding bonds to reconstruct a term structure: Indonesia (22 instruments), Saudi Arabia (12), Bahrain (11), Türkiye (5) and Malaysia (4) — **54 bonds, ~46,000 daily observations back to 2018**, spanning the rating spectrum from A (Malaysia) through AA/BBB (Saudi Arabia, Indonesia) to BB/B high yield (Türkiye, Bahrain). These daily G-spreads are Cbonds’ model-implied fair values, not trade prints: we verified that no continuous trade-level series exists (these sukuk print only a handful of times per quarter), so the model series is the only — and the standard — continuous data for instruments this illiquid. A single bond’s series cannot be used directly: as the bond ages its maturity shortens, so its spread slides down the term structure, and this *roll-down* is a deterministic trend that masquerades as mean-reversion. We remove it by reconstructing, for each issuer and each month, the spread at *fixed* tenors (3, 5 and 7 years) from a fit to the cross-section of that issuer’s outstanding bonds, and estimate dynamics on the constant-maturity series. We fit a discrete Ornstein–Uhlenbeck / AR(1) process, $\hat{\kappa}_{t+1} = c + \phi \hat{\kappa}_t + \varepsilon_{t+1}$ with $\alpha = -(\ln \phi)/\Delta t$, $\Delta t = 1/12$ (annualised), long-run mean $\bar{\kappa} = c/(1 - \phi)$ and half-life $12 \ln 2/\alpha$ months, as in eq. (9), applying Kendall’s finite-sample bias correction to ϕ and Newey–West standard

errors to α .

Results. The riba-free intensity is strongly mean-reverting on real data (Figure 3, Table 7):

1. **Long-run level, by credit quality.** At the 7-year tenor the long-run mean $\bar{\kappa}$ ranges from 0.9% per annum for Malaysia through 2.3% (Saudi Arabia) and 2.8% (Indonesia) to 5.2% for Bahrain, with Türkiye at $\approx 7\%$ (5-year). The investment-grade issuers cluster at $\bar{\kappa} \approx 1\text{--}3\%$ and the high-yield issuers at $\approx 5\text{--}7\%$ — a clean three-to-five-fold separation that matches both the cross-section (Section 5.1) and the order of magnitude (10^{-2} p.a.) the credit channel predicts, now recovered from an independent, dynamic estimation. (The fine ordering carries the same two wrinkles as the cross-section: Malaysia, a tight and infrequent issuer, prints below higher-rated Saudi Arabia, and Türkiye’s level is lifted by its 2022–2024 lira-crisis window; the broad IG-versus-HY split is unaffected.)
2. **Mean-reversion.** On the longest, best-sampled series — the 7-year tenor, with up to 93 monthly observations from 2018 — the bias-corrected half-life is short and broadly similar across issuers: 1.4 months (Malaysia), 2.7 (Bahrain), 3.3 (Saudi Arabia) and 4.3 (Indonesia), with α significant for the three with the longest histories (1.95–3.11, standard errors 0.51–1.27). All are *faster* than the multi-quarter-to-multi-year half-lives documented for conventional sovereign spreads in the sovereign credit literature [4] — estimated directly on real sukuk. Türkiye is the one issuer whose mean-reversion is *not* identified ($\alpha = 0.4\text{--}0.8$, standard errors larger than the point estimates): its 2022–2026 sample is dominated by a single secular decline in spreads, a near-unit-root trend rather than reversion, so we report its level but not a half-life.
3. **Feller / non-negativity.** The Feller condition $2\alpha\bar{\kappa} \geq \nu^2$ holds in *all twelve* constant-maturity series (five issuers $\times$ three tenors, where sampled). The real κ process stays strictly positive under mean reversion — the structural property the paper claims, now an empirical fact across the rating spectrum rather than an artefact of a low-volatility construction.

What this settles. With the cross-section (Section 5.1) and these dynamics, the empirical content of the κ -yield curve — its level, its credit ordering, its slope, its mean-reversion and its non-negativity — rests entirely on observed sukuk. The two pieces of real evidence are an instrument-level cross-section (195 bonds, sovereign and corporate) and constant-maturity time series.

Full-universe dynamics. We extend the five-issuer dynamics to the entire outstanding USD-sukuk universe. Pulling the daily Cbonds G -spread history of all 409 outstanding sukuk (2015–2026; Cbonds account 972307) and reconstructing constant-maturity $\hat{\kappa}$ by *interpolation only*—reading a fixed tenor solely where the issuer’s own bonds bracket it, never extrapolating—yields identifiable mean-reversion for **21 issuers**, four times the headline subset (Table 6). The method validates against the careful per-issuer estimates on the overlap (Saudi Arabia 2.36 vs 2.31%, Saudi Electricity 2.32 vs 2.20%, DP World 3.51 vs 3.21%), and issuers whose bonds cannot bracket a standard tenor (e.g. short-dated Islamic-bank curves) are *dropped* rather than extrapolated. Across the 21, $\bar{\kappa}$ spans 2.0% (First Abu Dhabi Bank) to 12.1% (Arada Developments) in broad alignment with credit quality — tight bank and quasi-sovereign names at the bottom, leveraged developers (Dar Al Arkan, Arada) at the top — though the alignment is broad rather than strict: Tenaga Nasional, an investment-grade utility, prints 6.4% amid high-yield names, a reminder that an interpolated single-issuer tenor can carry liquidity and curve-coverage noise. (Table 6 is sorted by $\bar{\kappa}$, not by rating; we do not claim strict monotonicity here, only the broad ordering.) The bias-corrected half-life is short throughout (median 1.8 months, monthly persistence $\hat{\phi} \approx 0.68$), and **the Feller non-negativity condition holds in all twenty-one**—so the mean-reverting, strictly-positive character of the riba-free intensity is now an empirical fact across the traded sukuk market, not a five-issuer illustration.

5.4 The κ -Index: A Cross-Sectional Benchmark Series

The estimates above price individual issuers and tenors; markets also need a *headline* number — the figure a terminal screen displays, the state variable a protocol publishes, the series a stabilisation fund watches. We define the κ -**index** as the cross-sectional median of $\hat{\kappa}$ across all quoted issuers in the full universe, computed monthly from the complete history panel (273 bonds aggregated to 37 issuer-month series with at least four years of data). The resulting series spans 127 months (December 2015 – June 2026; Figure 4) with mean 4.18% and standard deviation 1.80%. Its measured dynamics are the slow-moving common credit factor: monthly AR(1) coefficient $\rho = 0.921$ (half-life 8.4 months), innovation $\sigma = 0.66\%$ — far more persistent than any single issuer’s constant-maturity series, because the median filters idiosyncratic noise. The empirical $p95$ and $p99$ of the index (7.86% and 9.48%) provide the stress thresholds used as governance constants in downstream applications, and the decade contains two stress episodes against which any κ -linked mechanism can be back-tested: March–June 2020 (peak 11.63%, a +762bp jump over base, decay half-life three months from peak) and early 2016 (+500bp, slower decay). The index requires no committee

Issuer	$\bar{\kappa}\%$	HL	Issuer	$\bar{\kappa}\%$	HL
First Abu Dhabi Bank	1.96	0.3	Emirates Islamic Ban	3.20	1.6
Khazanah Nasional	2.19	1.0	Indonesia	3.37	6.2
Saudi Electricity Co	2.32	0.9	DP World	3.51	1.8
Saudi Arabia	2.36	2.7	Sharjah	4.37	2.6
Saudi Arabian Oil	2.49	0.2	Bapco Energies	5.71	0.9
Saudi National Bank	2.49	0.9	Tenaga Nasional	6.40	10.6
Public Investment Fu	2.58	0.6	Bahrain	7.87	6.1
Dubai Islamic Bank	2.77	1.8	Dar Al Arkan	9.84	1.9
Aldar Investment Pro	2.78	0.5	Turkey	10.03	7.1
Majid Al Futtaim Hol	3.14	1.1	Arada Developments	12.13	2.3
Emaar Properties	3.16	2.1			

Table 6: Constant-maturity κ dynamics on the full outstanding USD-sukuk universe (interpolation-only reconstruction, validated against the five-issuer subset; Cbonds 2015–2026, $\delta = 0.60$). **21 issuers** have identifiable mean-reversion (vs. five in the headline panel); rows are sorted by $\bar{\kappa}$, which runs from 2.0% (First Abu Dhabi Bank) to 12.1% (Arada) in broad — not strict — alignment with credit quality (Tenaga Nasional, investment grade, prints 6.4%; see text), half-life (HL, months) is short throughout, and the Feller non-negativity condition holds in *all* 21. Overlap issuers reconcile with the careful per-issuer estimates (Saudi Arabia 2.36 vs 2.31; Saudi Electricity 2.32 vs 2.20; DP World 3.51 vs 3.21).




Figure 3: Real constant-maturity $\hat{\kappa}$ for five sovereign sukuk issuers across the rating spectrum, 2018–2026 (Cbonds daily G-spread, roll-down removed by fixed-tenor reconstruction; headline tenor 7y, or 5y for Türkiye). *Left*: each issuer’s series mean-reverts around its own long-run level $\bar{\kappa}$ (dashed); the COVID-2020 and 2022 spikes decay back within months for the identified issuers. *Right*: $\bar{\kappa}$ separates investment grade (≈ 1 –3%: Malaysia, Saudi Arabia, Indonesia) from high yield (≈ 5 –7%: Bahrain, Türkiye), recovered independently from the time series. Bias-corrected 7-year half-lives are 1.4–4.3 months; the Feller non-negativity condition holds in all twelve constant-maturity series.

Issuer (rating)	Tenor	Months	$\bar{\kappa}$ (%)	Half-life (mo)	α (s.e.)	Feller
Malaysia (A)	7y	43	0.89	1.4	5.83 (3.79)	OK
Saudi Arabia (AA)	3y	43	1.42	1.3	—	OK
Saudi Arabia (AA)	5y	69	1.72	1.3	—	OK
Saudi Arabia (AA)	7y	93	2.31	3.3	2.53 (0.71)	OK
Indonesia (BBB)	3y	55	1.52	0.7	—	OK
Indonesia (BBB)	5y	70	1.99	1.9	—	OK
Indonesia (BBB)	7y	93	2.80	4.3	1.95 (0.51)	OK
Türkiye (BB)	3y	53	6.31	n.i.	0.40 (0.85)	OK
Türkiye (BB)	5y	50	7.44	n.i.	0.79 (0.87)	OK
Bahrain (B)	3y	40	4.70	2.8	—	OK
Bahrain (B)	5y	56	4.49	2.3	—	OK
Bahrain (B)	7y	56	5.22	2.7	3.11 (1.27)	OK

Table 7: Real constant-maturity κ dynamics for five sovereign issuers spanning the rating spectrum (Cbonds daily G-spread, monthly sampling, $\delta = 0.60$; 54 bonds). Half-lives are Kendall bias-corrected; α is the annualised mean-reversion speed with Newey–West standard error (shown for the headline 7-year series; “n.i.” = not identified). $\bar{\kappa}$ is the long-run mean. Feller = $2\alpha\bar{\kappa} \geq \nu^2$, satisfied in every series. The 7-year tenor has the longest history (2018–2026) and densest curve coverage and is the most reliable; shorter tenors give downward-biased half-lives from sparser short-maturity fits. Türkiye’s mean-reversion is not identified — its 2022–2026 sample is a single secular spread decline (near-unit-root), so only its level is reported. Malaysia has only a 7-year reconstruction (four bonds, all long-dated).

and no survey: it is a deterministic statistic of observed market prices, recomputable by anyone from the replication data. Its manipulation resistance should be stated precisely: the median’s 50% breakdown point means the index cannot be driven *arbitrarily far* without distorting half the quoted universe, but a *bounded* shift requires only moving quotes near the middle of the cross-section — so the median is a mitigation against gross corruption, not against marginal drift, and the downstream defence rests on the change-point monitor and depth rule of the measurement-integrity companion paper rather than on the median alone.

5.5 The Benchmark-Free Estimator and the Recovery Floor

A fair objection to everything above is that $\hat{\kappa} = s/(1 - \delta)$ is operationally parasitic on the conventional system: the G-spread s cannot be computed without the Treasury curve, so the “riba-free” intensity is an arithmetic residual of the riba system. The $\iota = 0$ framework, however, implies its own estimator that uses no benchmark whatsoever. At $\iota = 0$ a claim’s price is the *undiscounted* expectation of its payoff at the stopping or maturity time; in zero-coupon-equivalent form, per unit face, $P(\kappa) = e^{-\kappa T} + \delta(1 - e^{-\kappa T})$, which inverts against the




Figure 4: The κ -index: monthly cross-sectional median $\hat{\kappa}$ across the full USD-sukuk universe, December 2015 – June 2026, with the empirical $p95$ (dashed) and $p99$ (dotted) thresholds. The 2016 and 2020 stress episodes are visible; the 2020 peak is 11.63%.

yield-implied price $P = e^{-yT}$ as

$$\hat{\kappa}_{\text{price}} = -\frac{1}{T} \ln \frac{e^{-yT} - \delta}{1 - \delta},$$

a function of the bond’s *own* yield and maturity only. Three results on the full-universe snapshots (60 sovereign and 213 corporate bonds with quoted yield, spread, and maturity):

First, the cross-section is benchmark-free recoverable. The rank correlation between $\hat{\kappa}_{\text{price}}$ and the spread-based $\hat{\kappa}$ is 0.84 (sovereign) and 0.94 (corporate) at $\delta = 0.10$: the credit ordering, the curve shapes, and all relative pricing in this paper survive with the Treasury curve deleted from the pipeline.

Second, the level wedge is the reserve currency’s yield layer, isolated — and we are careful about what that layer is. $\hat{\kappa}_{\text{price}}$ exceeds the spread version by a median of 458–459bp of κ — against 472bp predicted by mapping the prevailing Treasury yield through the same transformation, and the wedge tracks the Treasury yield through time (the out-of-sample companion measures $\rho = 0.93$ over 2015–26). Passed through the programme’s own realised-hazard decomposition, this wedge is *not* a measured default intensity of the reserve sovereign: the US default record puts the realised intensity near zero, so the wedge is almost entirely the time-value and risk-premium layer of the reserve currency — the same layer the programme identifies as the removable one, here made visible rather than absorbed into κ . The correct statement is therefore about *orderings*, *not levels*: because the wedge is a common additive

shift, the credit ordering, the curve shapes, and all relative pricing survive with the Treasury curve deleted (the rank correlations above), and netting against the Treasury curve is a *decomposition convenience* — it isolates issuer excess over the dollar system’s base yield layer — not an operational dependency for anything cross-sectional. In a mature κ -world the netting benchmark is unnecessary for relative pricing: κ inverts from prices and cashflows alone, and whatever time or premium content that world’s clearing prices carry becomes a measured quantity in the level, exactly as it is here.

Third, the framework’s own consistency condition adjudicates the recovery rate. At $\iota = 0$ a price can never fall below the undiscounted recovery δ — so market prices with $e^{-yT} < \delta$ are inconsistent with the pair $(\iota = 0, \delta)$. At the conventional $\delta = 0.60$, 11–13% of observed prices in both universes violate this floor (long-maturity bonds especially); at $\delta = 0.10$ — precisely the realized recovery measured on this market’s own defaulted corporate sukuk in the companion credit-protection study — feasibility is 100% and the benchmark-free correlation is at its maximum. The $\iota = 0$ pricing identity and the measured default evidence thus select the same low corporate recovery independently; under it, corporate $\hat{\kappa}$ levels reported above rescale by the factor $(1 - 0.60)/(1 - 0.10) \approx 0.44$ while every ordering and slope result is unchanged (the rescaling is uniform). Sovereign recoveries remain a convention — the dataset’s own finding is that no sovereign USD sukuk has ever defaulted — and sovereign results are reported at the standard $\delta = 0.60$ with this caveat attached.

6 Discussion

6.1 Practical Implications for Islamic Finance

The κ -yield curve constructed in this paper has three immediate practical applications.

Sukuk pricing benchmark. Conventional sukuk pricing currently uses SOFR or government bond yields as the risk-free benchmark, creating an implicit interest-bearing component in an ostensibly Shariah-compliant product. The κ -yield curve — estimated here on observed sovereign sukuk (Sections 5.1–5.3) and, in production, extended to a full Islamic-credit panel — provides a term structure whose return object carries no interest-bearing component; in the present implementation the US Treasury curve enters only as a netting benchmark to isolate the credit component of the observed yield, and a production implementation would benchmark sukuk yields against the sovereign κ -curve itself. Portfolio managers can then strip sukuk option-adjusted spreads into a κ -term and a credit-specific component, analogous to the Z-spread decomposition in conventional bond analysis.

Islamic monetary policy. Central banks in GCC and Southeast Asian markets (Bank

Negara Malaysia and the GCC central banks, including the Saudi Central Bank, SAMA) currently operate dual monetary systems, maintaining both conventional and Islamic policy rates. The κ -yield curve offers a data-driven, no-arbitrage alternative to the administered Islamic policy rate: monetary authorities could target $\bar{\kappa}$ directly, using sovereign sukuk issuance to shift the term structure in a manner analogous to conventional open-market operations.

Halal derivatives valuation. Ahmed and Bhuyan [2] apply the AHJ framework to price everlasting options, takaful (Islamic insurance) and iCDS (Islamic credit default swaps) at $\iota = 0$. All three product families require a term structure of κ as an input. Once estimated on real sukuk spreads, κ -yield curves built with this pipeline can be fed directly into those pricing engines, replacing the ad hoc convention of using LIBOR/SOFR with an appropriate Islamic flat rate. In particular, a traded constant-maturity κ benchmark is the natural settlement reference for asset-linked risk-transfer instruments — e.g. hedging an Islamic institution’s sukuk-portfolio yield exposure in place of an interest-rate swap, or pricing protection on tokenised, asset-backed positions — where the instrument’s function is risk transfer rather than rate speculation; product-layer Shariah design questions for such instruments are treated in the companion Shariah-boundaries note [3].

6.2 Non-Negativity and the Negative-Rate Dilemma

κ is non-negative by construction, and the Feller condition — satisfied in every real constant-maturity series we estimate (Section 5.3) — guarantees it stays strictly positive under mean reversion. This stands in sharp contrast to the euro-area and Japanese overnight benchmark and policy rates, which were negative in the eurozone (2014–2022) and Japan (2016–2024). Negative rates create theological and structural problems for Islamic finance: if the reference rate used in sukuk pricing turns negative, it implies negative mark-ups on *murabaha* and negative profit rates on *musharaka* instruments, both of which are arguably contrary to the spirit of risk-sharing embedded in Islamic contracts.

The non-negativity guarantee of κ — derived not from a constraint but from the fundamental structure of credit pricing — means that central banks in jurisdictions operating under Islamic finance frameworks never need to navigate the negative-rate policy dilemma that bedevilled the ECB and BOJ. The κ -yield curve thus offers a practically attractive framework for Islamic monetary policy in addition to settling the *riba* question at the pricing layer. Product-layer Shariah questions (*gharar*, *maysir*, possession) are outside the scope of this construct and remain with Shariah supervisory boards; see the companion Shariah-boundaries note [3].

7 Conclusion

We define the κ -yield curve — the *riba*-free analogue of the risk-free term structure — and estimate it entirely on real traded sukuk: point identification $\hat{\kappa} = s/(1 - \delta)$ on observed G -spreads, a cross-sectional characterisation by rating and maturity, and AR(1)/OU estimation of CIR mean-reversion on roll-down-free constant-maturity series. The empirical content rests on two real datasets — no synthetic panel.

Result 1 (Level and credit ordering). On a real cross-section of 195 actively-quoted USD sukuk (40 sovereign, 154 corporate, one municipal; Cbonds), the convergence intensity is of order 10^{-2} per annum and broadly increasing as credit quality deteriorates, from $\hat{\kappa} \approx 0.9\%$ at AAA to $\approx 4.4\%$ at single-B (apart from the Saudi-dominated AA bucket, which prices just above single-A) — the credit ordering the framework predicts, read directly from the sukuk market.

Result 2 (Term structure). The investment-grade κ -curve slopes upward: the G -spread rises with maturity at +1.17 bp/yr pooled, +0.63 bp/yr with rating fixed effects ($n = 164$). The slope is positive but modest — an honest real-data estimate, weaker than a designed curve.

Result 3 (Dynamics and Feller). On roll-down-free constant-maturity series for five issuers across the rating spectrum, κ mean-reverts with a bias-corrected half-life of 1.4–4.3 months at the 7-year tenor around a long-run mean that separates investment grade ($\approx 1\text{--}3\%$) from high yield ($\approx 5\text{--}7\%$), faster than conventional sovereign spreads, and the Feller non-negativity condition holds in all twelve series.

The contribution is the κ -yield-curve construct together with a real-data characterisation of its level, credit ordering, slope and dynamics on traded sukuk.

Future work. The two real datasets cover complementary dimensions — a 195-bond cross-section and constant-maturity time series (five issuers, up to eight years). The natural next steps are (i) a full real panel uniting both — every issuer, every month — to estimate the CIR dynamics at the instrument level; (ii) deepening the sovereign/corporate decomposition begun in Section 5.2 to a full sector-by-rating corporate κ surface; and (iii) construction of a κ -based Islamic yield curve model that central banks can use for open-market operations targeting the term structure of the convergence intensity.

Data Availability Statement

All empirical results in this study are estimated on real traded sukuk from the Cbonds bond database (<https://cbonds.com>). The cross-section (Section 5) is the set of outstanding, actively-quoted USD sukuk retrieved live from the Cbonds Bond Screener (/api/bonds/search/): the 195-bond sovereign-plus-corporate panel is shipped as `real_sukuk_panel_full.csv` and the per-rating / term-structure tables and the κ -curve (Figure 1) are reproduced from it by `real_panel_analysis.py`; the 61-bond sovereign snapshot (issuer, ISIN, maturity, yield, G -spread, rating) is shipped as `sukuk_cbonds_panel.csv`, and Figure 2 is reproduced from it by `sukuk_cbonds_panel.py` and `fig_sukuk_cbonds_curve.py`. The maturity-matched US Treasury benchmark underlying the G -spread is Cbonds' interpolation of the public UST par curve. SOFR is from FRED (<https://fred.stlouisfed.org>).

The real κ dynamics of Section 5.3 use the daily G -spread histories (2018–2026) of all 54 USD sovereign sukuk of the five issuers (Indonesia, Saudi Arabia, Bahrain, Türkiye, Malaysia), retrieved from the same Cbonds source. The extracted daily histories and the derived constant-maturity series are shipped as `sukuk_curve_histories_all.json` and `sukuk_cm_series.csv`; the constant-maturity construction, the Kendall-bias-corrected AR(1)/OU estimation and Figure 3 are reproduced from those files by `sukuk_constant_maturity.py`. The instrument-level sukuk-versus-conventional comparison (USD) and the sovereign/corporate decomposition (Section 5.2) draw on, respectively, the 133 conventional sovereign USD bonds and 213 corporate USD sukuk of the same issuers, also from Cbonds; both extracts ship as `sukuk_vs_conventional.csv` and `corporate_sukuk_panel.csv`. No Cbonds subscription is needed to reproduce any reported number from the shipped extracts.

Acknowledgements

This research received no specific grant from any funding agency in the public, commercial, or not-for-profit sectors. All errors remain our own.

References

- [1] Akerer, D., Hugonnier, J., & Jermann, U. (2025). Perpetual futures pricing. *Mathematical Finance*, 1–17. <https://doi.org/10.1111/mafi.70018> (NBER Working Paper 32936, 2024; SSRN 4603820, 2023.)
- [2] Ahmed, S., & Bhuyan, R. (2026). From perpetual contracts to Islamic credit: The

- random stopping-time equivalence (credit equivalence of the κ -rate). *SSRN Working Paper 6322858*. https://papers.ssrn.com/sol3/papers.cfm?abstract_id=6322858
- [3] Ahmed, S., & Bhuyan, R. (2026b). The jurisprudential boundaries of an interest-free substrate: Riba, gharar, maysir, and qabd in cash-settled synthetic instruments. Working paper (companion research note).
 - [4] Duffie, D., Pedersen, L. H., & Singleton, K. J. (2003). Modeling sovereign yield spreads: A case study of Russian debt. *Journal of Finance*, 58(1), 119–159.
 - [5] Moody’s Investors Service. (2024). *Sovereign Default and Recovery Rates* (annual study series). Moody’s Investors Service, New York.
 - [6] Bank for International Settlements. (2022). Sovereign yield curves and term premia in emerging markets. *BIS Quarterly Review*, March 2022, 67–83.
 - [7] Cakir, S., & Raei, F. (2007). Sukuk vs. Eurobonds: Is there a difference in value-at-risk? *IMF Working Paper WP/07/237*, International Monetary Fund, Washington D.C.
 - [8] Uddin, M. H., Kabir, S. H., Hossain, M. S., Wahab, N. A., & Liu, J. (2022). Why do sukuks (Islamic bonds) need a different pricing model? *International Journal of Finance & Economics*, 27(2), 2210–2234. <https://doi.org/10.1002/ijfe.2269>
 - [9] Cox, J. C., Ingersoll, J. E., & Ross, S. A. (1985). A theory of the term structure of interest rates. *Econometrica*, 53(2), 385–407.
 - [10] Duffie, D., & Singleton, K. J. (1999). Modeling term structures of defaultable bonds. *Review of Financial Studies*, 12(4), 687–720.
 - [11] Jobst, A. A. (2007). The economics of Islamic finance and securitization. *IMF Working Paper WP/07/117*, International Monetary Fund, Washington D.C.
 - [12] Nelson, C. R., & Siegel, A. F. (1987). Parsimonious modeling of yield curves. *Journal of Business*, 60(4), 473–489.
 - [13] Newey, W. K., & West, K. D. (1987). A simple, positive semi-definite, heteroskedasticity and autocorrelation consistent covariance matrix. *Econometrica*, 55(3), 703–708.
 - [14] Svensson, L. E. O. (1994). Estimating and interpreting forward interest rates: Sweden 1992–1994. *NBER Working Paper 4871*, National Bureau of Economic Research, Cambridge MA.

- [15] Vasicek, O. (1977). An equilibrium characterization of the term structure. *Journal of Financial Economics*, 5(2), 177–188.

A Derivation of the Identification (Theorem 2.1)

We outline the derivation of equation (2) from the reduced-form no-arbitrage conditions at $\iota = 0$.

Consider protection on a unit notional with loss-given-default $(1 - \delta)$ at the default time τ_D (intensity λ under the risk-neutral measure; the reduced-form framework of 10). With discounting at the convergence intensity κ ($\iota = 0$), the protection leg is worth, at constant intensities,

$$V_{\text{prot}} = (1 - \delta) \frac{\lambda}{\kappa + \lambda}, \quad (10)$$

and a premium leg paying the spread s continuously until τ_D is worth $s/(\kappa + \lambda)$. Equating the two legs, the common factor $1/(\kappa + \lambda)$ cancels, giving the credit triangle

$$s = \lambda(1 - \delta). \quad (11)$$

The identification of the convergence intensity with the default intensity, $\kappa = \lambda$, is *not* an independent steady-state fact: it is the Random Stopping-Time Equivalence of Ahmed & Bhuyan [2] (the Separation-Theorem isomorphism between the AHJ stopping time and the reduced-form default time). Substituting yields

$$s = \kappa(1 - \delta) \implies \kappa = \frac{s}{1 - \delta}. \quad (12)$$

Equivalently, the fixed-maturity survival claim of Theorem 2.1 can be priced directly: $V(t) = \mathcal{P}e^{-\kappa(1-\delta)(T-t)}$, whose continuously-compounded spread over the default-free benchmark is $s = \kappa(1 - \delta)$. This completes the derivation of (2).

B GMM Moment Conditions

The two-step GMM estimator minimises:

$$Q_T(\boldsymbol{\theta}) = \mathbf{g}_T(\boldsymbol{\theta})^\top \hat{\mathbf{W}}^{-1} \mathbf{g}_T(\boldsymbol{\theta}), \quad (13)$$

where $\mathbf{g}_T(\boldsymbol{\theta}) = T^{-1} \sum_{t=1}^T \mathbf{z}_t \otimes \mathbf{m}_t(\boldsymbol{\theta})$, with moment conditions:

$$m_{1t}(\boldsymbol{\theta}) = \hat{\kappa}_{t+1} - \hat{\kappa}_t e^{-\alpha h} - \bar{\kappa}(1 - e^{-\alpha h}), \quad (14)$$

$$m_{2t}(\boldsymbol{\theta}) = m_{1t}^2 - \frac{\nu^2}{\alpha} \hat{\kappa}_t (e^{-\alpha h} - e^{-2\alpha h}) - \frac{\nu^2 \bar{\kappa}}{2\alpha} (1 - e^{-\alpha h})^2, \quad (15)$$

and $\hat{\mathbf{W}}$ is the Newey–West HAC estimator of the long-run variance of $\mathbf{z}_t \otimes \mathbf{m}_t$.

C Country Coverage and Sukuk Universe

Table 8 summarises the real sukuk universe underlying our cross-section — the 195 actively-quoted USD sukuk by issuer country.

Table 8: Real Sukuk Universe (195 actively-quoted USD sukuk, Cbonds live pull)

Country	Sovereign	Corporate	Total	G -spread range (bp)
Saudi Arabia	12	74	86	11–264
United Arab Emirates	0	46	46	44–248
Indonesia	21	0	21	26–294
Kuwait	0	13	13	27–523
Qatar	1	11	12	29–393
Malaysia	4	7	11	13–1281
Others (HK, LU, BH, US, PH, IE)	2	4	6	32–141
Total	40	155	195	—

Notes: All instruments are outstanding, actively-quoted USD sukuk retrieved live from the Cbonds Bond Screener (account 972307), enumerated per instrument in the shipped `real_sukuk_panel_full.csv`. The constant-maturity dynamics (Section 5.3) use daily G -spread histories of 54 sukuk from five issuers, 2018–2026. Ratings are foreign-currency composites as of the pull date. The non-sovereign column (155) comprises 154 corporate sukuk and one municipal sukuk, the latter grouped with corporates here.

EDITORIAL NOTE TO THIS PRINTING

Paper 2b · Stochastic Takaful

Working paper, revised July 2026 · SSRN 6323459

Role in the programme. The primitive applied where it is strongest: cooperative insurance priced off a stochastic hazard. **What it establishes.** Takaful contribution pricing under a mean-reverting (CIR) intensity, with the cost accounting stated plainly: at current parameters the compliant product is about 7.4% more expensive than its conventional twin, and the paper says so. **Corrected or extended later.** The agricultural loss-cost ratio is now read directly from the insurance record rather than through a bond-recovery haircut. **Not established here.** A yield-indexed agri-takaful design; the correlation between drought and crop prices is caveated as an open design problem, not solved.

Replication for every number in this paper: DOI 10.5281/zenodo.21290557 · github.com/shehzadahmed-xx/kappa-replication

Actuarial Properties of the κ -Rate Under Stochastic Hazard: CIR Applications to Agricultural Takaful in South and Southeast Asia

Shehzad Ahmed¹ and Rafiqul Bhuyan^{*2}

¹Department of Finance, Independent University Bangladesh (IUB),
Dhaka, Bangladesh. 2120922@iub.edu.bd

²Department of Accounting and Finance, Alabama A&M University,
Normal, AL, USA. rafiqul.bhuyan@aamu.edu

March 2026

^{*}Corresponding author. Department of Accounting and Finance, Alabama A&M University, Normal, AL 35762, USA. E-mail: rafiqul.bhuyan@aamu.edu.

Abstract

We derive the actuarial properties of the hazard intensity κ — the credit/loss rate introduced in our companion papers (Ahmed & Bhuyan, 2026a,b,c) — and apply it to agricultural takaful pricing in South and Southeast Asia. We estimate κ on cereal-yield data and validate the underlying $\iota = 0$ pricing logic on USDA loss-cost data. Starting from the continuous-time takaful pricing formula $\pi = \kappa \cdot B$ (premium equals hazard intensity times benefit), we extend it to a Cox–Ingersoll–Ross (CIR) stochastic- κ setting and derive a closed-form affine premium for multi-period takaful contracts. Using World Bank cereal yield data (1980–2023) for Bangladesh, Pakistan, Indonesia, and India, we estimate $\hat{\kappa}_{\text{MLE}} \in [0.071, 0.147]$ and calibrate CIR mean-reversion parameters from rolling-window estimates. For India, the model-implied rate ($\hat{\kappa} = 12.06\%$) sits near the PMFBY national-average actuarial rate ($\approx 12\%$), but we read this as a units-level coincidence, not a validation: Pakistan shares the identical $\hat{\kappa} = 12.06\%$ yet misses its observed rates by $1.6\text{--}3.4\times$. For subsidised market products, the stochastic- κ formula generates premiums $1.3\text{--}8\times$ observed market rates (gross tariffs and, for India, farmer-paid caps); the gap is explained by government subsidies (50–80%), benefit caps, and co-insurance requirements embedded in existing products. We isolate the *riba loading* — the premium reduction caused by discounting liabilities at a positive interest rate — and show it equals 7.4% of the *riba*-free premium at an 8% discount rate. On real USDA RMA crop-insurance data (2014–2023), the premium rate charged (9.6%) closely tracks the realised loss-cost (7.9%, loss ratio 0.82), consistent with the $\iota = 0$ fairness relation $\pi^* \approx \mathbb{E}[\text{loss}]$ on a functioning insurance market; we also document that the yield-shortfall $\hat{\kappa}$ moves over a $5\text{--}15\times$ range with the loss threshold, so its level is threshold-conditional. Our framework provides a fully *riba*-free actuarial pricing model for agricultural takaful under stochastic hazard, grounded in the $\iota = 0$ constraint established in our companion papers.

A measure-theoretic reading sharpens the comparison: our premiums are P -measure (historical loss frequency) while market premiums embed both subsidy politics *and* the loadings a commercial insurer charges — a Q -style risk premium plus expenses. The model’s deviations from market rates therefore conflate three wedges (subsidy, loading, measure) that the present data cannot separate; we flag the comparison as descriptive, not a validation, until claims-level data permits decomposition.

Keywords: takaful, agricultural insurance, Cox–Ingersoll–Ross, hazard intensity, *riba*-free pricing, Islamic finance, κ -rate, South Asia, Southeast Asia.

JEL: G22, G12, Q14, Z12.

Contents

1	Introduction	4
2	Related Literature	5
2.1	Index-Based Agricultural Insurance	5
2.2	Agricultural Insurance in South and Southeast Asia	6
2.3	Takaful: Structure and Pricing	6
2.4	Stochastic Actuarial Models	7
3	Theoretical Framework	7
3.1	Takaful as a Mutual Guarantee	7
3.2	Continuous-Time Loss Intensity	7
3.3	The Riba Loading in Conventional Pricing	8
4	CIR Stochastic Hazard Rate Model	9
4.1	CIR Dynamics	9
4.2	Closed-Form Affine Premium	9
4.3	Comparison: Constant vs Stochastic κ	10
5	Data	10
5.1	Cereal Yield Data	10
5.2	Takaful Premium Benchmarks	11
6	Estimation Methodology	11
6.1	Loss Year Identification	11
6.2	MLE Hazard Estimation	12
6.3	CIR Calibration via AR(1) OLS on Rolling $\hat{\kappa}$	12
7	Empirical Results	13
7.1	Cereal Yield Trends and Loss Years	13
7.2	Decade-Level $\hat{\kappa}$ and Mean Reversion	15
7.3	CIR Parameter Estimates	15
7.4	Model vs Observed Premiums	17
7.5	Stochastic vs Constant κ Premium	19
8	Riba Loading Decomposition	21
8.1	Quantifying the Riba Loading	21
8.2	Policy Implications	22

9	Smart Contract Implementation	23
9.1	TakafulPool Architecture	23
9.2	Claims Settlement	24
9.3	On-Chain Pricing vs CIR Model	24
10	Alignment with AAOIFI Standard No. 26: A Self-Assessment	24
10.1	Relevant Standards	25
10.2	Self-Assessment of the κ -Framework	25
11	Conclusion	26
A	MLE Consistency Under Bernoulli Loss Process	30
B	Euler–Maruyama Discretisation of CIR	31
C	AAOIFI Shariah Standard No. 26: Design Self-Assessment Checklist	31

1 Introduction

Climate-driven crop losses are rising in both frequency and severity across the developing world, with weather-related catastrophe losses reaching record levels in recent years (Swiss Re, 2023), sharpening the demand for affordable agricultural risk transfer in exactly the South and Southeast Asian markets we study. Agricultural takaful — the Islamic mutual insurance of crop and livestock risk — is among the fastest-growing segments of the global Islamic finance industry. The global takaful market was valued at approximately \$27 billion in 2023 and is projected to exceed \$50 billion by 2030, with agricultural products driving the growth in South and Southeast Asia (IFSB, 2024). Yet despite this growth, a fundamental problem persists: virtually all agricultural takaful products price premiums using conventional actuarial methods that embed an interest rate in the discount factor, violating the prohibition of *riba* (usury) in Islamic law.

The problem arises as follows. In conventional crop insurance, the actuarially fair premium for a one-year policy with benefit B and default intensity κ is computed as

$$\pi_{\text{conv}} = \kappa \cdot B \cdot \frac{1}{1+r}, \quad (1)$$

where r is the risk-free interest rate. The factor $1/(1+r)$ discounts the expected benefit to present value; the insurer thereby earns interest on the held premium. Even when the product is labelled “takaful” and the fund structure is compliant, the *pricing* formula retains an interest loading of

$$\Pi_{\text{riba}} = \kappa B - \kappa B/(1+r) = \kappa B \cdot \frac{r}{1+r}, \quad (2)$$

which is the reduction in premium caused by discounting at a positive rate. Throughout, ι denotes the interest-rate parameter of the companion pricing framework (Ahmed & Bhuyan, 2026a); the *riba*-free constraint is $\iota = 0$, equivalent to setting $r = 0$ in the discount factor. The loading then vanishes exactly: $\Pi_{\text{riba}} = 0$ and the premium is the full actuarial risk premium $\pi^* = \kappa B$.

Our contribution is fourfold:

1. We provide the first systematic derivation of the takaful premium under a Cox–Ingersoll–Ross stochastic hazard rate κ_t at $\iota = 0$, showing that the closed-form affine bond price formula applies directly to Shariah-compliant mutual insurance.
2. We estimate κ from World Bank cereal yield data (1980–2023) for Bangladesh, Pakistan, Indonesia, and India using maximum-likelihood methods, and calibrate all three CIR parameters from 35 rolling 10-year $\hat{\kappa}$ windows via AR(1) OLS with Newey–West standard errors.
3. We quantify the *riba* loading embedded in conventional agricultural insurance pricing

and show that discounting at an 8% rate suppresses the premium by 7.4% relative to the riba-free benchmark; $\iota = 0$ pricing restores the full actuarial risk premium $\pi^* = \kappa B$, providing a strictly stronger actuarial basis.

4. We validate the $\iota = 0$ fairness relation $\pi^* \approx \mathbb{E}[\text{loss}]$ on real USDA RMA crop-insurance data (2014–2023), where the premium rate charged (9.6%) tracks the realised loss-cost (7.9%, loss ratio 0.82) on a functioning insurance market.

The paper is organized as follows. Section 2 reviews the related literature on index-based crop insurance, South Asian agricultural insurance markets, and takaful. Section 3 develops the theoretical framework and riba-loading decomposition. Section 4 extends the model to a CIR stochastic hazard rate and derives the closed-form affine premium. Section 5 describes the data. Section 6 presents the κ estimation methodology. Section 7 reports empirical results. Section 8 quantifies the riba loading. Section 9 describes the on-chain implementation. Section 10 discusses AAOIFI compliance. Section 11 concludes.

2 Related Literature

2.1 Index-Based Agricultural Insurance

Early reviews established both the rationale and the practical obstacles: Roberts (2005) surveys the actuarial design of crop insurance in developing countries, and Skees and Barnett (2006) show how index-based risk-transfer products can be layered onto microfinance to overcome the moral-hazard and adverse-selection problems that defeat traditional indemnity insurance. At the macro level, Outreville (2012) documents a robust cross-country relationship between insurance penetration and economic development, motivating the policy interest in extending affordable crop cover to lower-income farmers. Barnett and Mahul (2007) provide an early survey of weather-index products for lower-income countries, documenting both the risk-transfer potential and the practical challenges of basis risk. Mahul and Stutley (2010) examine the role of government subsidies and institutional support in making agricultural insurance viable, finding that actuarial premiums without subsidies typically exceed farmer willingness-to-pay by 2–5 $\times$ — a key benchmark for the model-to-market ratios we report in Section 7.

Basis risk — the divergence between the index trigger and actual farm losses — is a central concern in index insurance design. Negi and Ramaswami (2023) study basis risk in catastrophic rainfall insurance in India, finding that spatial aggregation substantially amplifies effective basis risk relative to actuarial expectations. Kramer, Porter, and Wassie (2025) show in an Ethiopian field experiment that basis risk perceptions reduce insurance demand even when actuarial terms are favourable, with social comparison effects amplifying dissatisfaction with false-negative payouts. For parametric crop insurance

pricing, Dal Moro (2024) demonstrates that bimodal yield distributions (common in monsoon-dependent agriculture) require non-Gaussian pricing adjustments that our CIR framework accommodates through mean-reversion rather than distributional assumptions.

2.2 Agricultural Insurance in South and Southeast Asia

Bangladesh. The World Bank (World Bank, 2010) documents that gross agricultural insurance premiums in Bangladesh typically range from 5–6% of sum insured for parametric and area-yield products, consistent with ADB Technical Assistance pilot 46284-001. These rates reflect the high-frequency but moderate-severity crop loss environment, where yield shortfalls occur approximately once per decade at the national level.

Pakistan. The Punjab Crop Loan Insurance Scheme (CLIS) has been extensively studied by the World Bank (World Bank, 2020), which documents actuarial pure premiums of 5% for kharif (summer) and 3.5% for rabi (winter) crops in Punjab province. Commercial parametric products such as the JazzCash / Salaam Takaful offering charge 7.5% of sum insured (PKR 1,500 per PKR 20,000 of coverage at the 2023 tariff), reflecting higher profit margins and adverse selection risk in voluntary products.

Indonesia. Indonesia’s government-mandated rice insurance programme (AUTP: Asuransi Usaha Tani Padi) sets the gross premium at exactly 3.0% of sum insured (IDR 180,000 per IDR 6,000,000 of coverage), with the government subsidising 80% so that farmers pay only 0.6% (Kementerian Pertanian, 2022). Recent empirical studies confirm that participation in AUTP is primarily driven by the subsidy level rather than risk awareness (Cahyasita, Aminda, and Kusumawardani, 2025; Fakhruzi and Syafril, 2025).

India. India’s Pradhan Mantri Fasal Bima Yojana (PMFBY) is the world’s largest crop insurance programme by enrolled area. Its design separates the actuarial premium (state-level average: 10–20% of sum insured, national average $\approx 12\%$) from the farmer-paid portion, which is capped at 2% for kharif and 1.5% for rabi, with the remainder split between central and state governments (GOI, 2016). Shekhar and Rai (2025) document that the PMFBY actuarial rate has risen over time as climate-driven yield volatility has increased, while farmer contributions remain fixed — a structural subsidy escalation that reduces programme sustainability. Punia, Kundu, and Mehla (2021) provide a comparative analysis of PMFBY against predecessor schemes, confirming that PMFBY expanded coverage substantially while preserving the actuarial rate structure.

2.3 Takaful: Structure and Pricing

Takaful as a Shariah-compliant alternative to conventional insurance has been extensively documented (Obaidullah and Khan, 2008; Encyclopedia of Islamic Insurance, 2019). The fundamental prohibition on *riba* in Islamic jurisprudence requires that insurance reserves

not be invested in fixed-income instruments and that the pricing formula contain no interest-rate discount factor (AAOIFI, 2010). In practice, most existing takaful products use the same actuarial pricing formulas as conventional insurance (with $r > 0$ in the discount factor) and rely on Shariah-compliant *investment* of reserves to maintain fund solvency — a distinction the $\iota = 0$ framework of this paper makes explicit by eliminating r from the premium formula itself (Ahmed & Bhuyan, 2026b).

2.4 Stochastic Actuarial Models

The Cox–Ingersoll–Ross (CIR) process (Cox, Ingersoll, and Ross, 1985) was developed for interest rate modelling but has found wide application in actuarial science wherever loss rates exhibit positive values and mean-reversion. Lamberton and Lapeyre (2011) provide a rigorous treatment of stochastic calculus applied to finance, including the affine term structure characterisation that underlies our closed-form premium formula. Şimşek and Yıldırak (2025) apply stochastic loss prevention modelling to crop yield insurance, reporting that moral hazard effects are substantially reduced in parametric structures — consistent with the takaful mutual model. To our knowledge, no prior work applies the CIR framework to actuarial hazard rates in agricultural insurance.

3 Theoretical Framework

3.1 Takaful as a Mutual Guarantee

Under Islamic law, conventional insurance is prohibited because it combines three forbidden elements: *riba* (interest), *maysir* (gambling), and *gharar* (excessive uncertainty). Takaful substitutes a mutual guarantee (*ta’awun*) for the commercial insurance contract: participants contribute to a common pool (*tabarru pool*) from which claims are paid. The operator earns a *wakalafee* (agency fee) for managing the pool but does not retain the underwriting surplus, which reverts to participants.

Definition 3.1 (Takaful Contract). *A takaful contract (B, T, π, γ_w) consists of:*

- $B > 0$: *benefit payable on loss event;*
- $T > 0$: *contract horizon;*
- $\pi > 0$: *tabarru (donation) premium per period;*
- $\gamma_w \in (0, 1)$: *wakala fee fraction retained by the operator.*

The participant’s net tabarru is $\tilde{\pi} = (1 - \gamma_w)\pi$.

3.2 Continuous-Time Loss Intensity

Let $(\Omega, \mathcal{F}, \mathbb{P})$ be a probability space and $\{N_t\}_{t \geq 0}$ a doubly-stochastic Poisson process (Cox process) with stochastic intensity $\{\kappa_t\}_{t \geq 0}$. A “loss event” occurs at the first jump time $\tau = \inf\{t \geq 0 : N_t \geq 1\}$.

Under risk-neutral measure $\mathcal{Q}$ with $\iota = 0$ (zero interest rate):

$$\mathbb{P}(\tau > t \mid \kappa_s, s \leq t) = \exp\left(-\int_0^t \kappa_s ds\right). \quad (3)$$

The actuarially fair premium for a one-year renewable takaful at $\iota = 0$ is the expected loss:

$$\pi^* = B \cdot \mathbb{P}(\tau \leq T) = B \mathbb{E}_0\left[1 - e^{-\int_0^T \kappa_s ds}\right]. \quad (4)$$

For constant κ and $T = 1$ year, this reduces to:

$$\pi^* = B(1 - e^{-\kappa}) \approx \kappa B \quad (\kappa \ll 1). \quad (5)$$

Equation (5) is the central formula of this paper.

3.3 The Riba Loading in Conventional Pricing

In conventional insurance, the discount rate $r > 0$ (the risk-free rate, proxied by government bond yields or LIBOR/SOFR) reduces the premium:

$$\pi_{\text{conv}}(r) = B \cdot \kappa \cdot \frac{1}{1+r}. \quad (6)$$

The riba loading is:

$$\begin{aligned} \Pi_{\text{riba}}(r, \kappa, B) &= \pi^*(\iota = 0) - \pi_{\text{conv}}(r) \\ &= \kappa B \left(1 - \frac{1}{1+r}\right) = \frac{\kappa B r}{1+r}. \end{aligned} \quad (7)$$

At $r = 8\%$, $\Pi_{\text{riba}}/\pi^* = r/(1+r) = 7.4\%$. This is the exact share of the premium that is suppressed by discounting at a positive interest rate.

Proposition 3.1 (Riba-Free Takaful Pricing). *At $\iota = 0$, the fair takaful premium is $\pi^* = B \frac{1 - e^{-\kappa T}}{T} \approx \kappa B$ per year for $\kappa T \ll 1$ (the total T -year premium is $B(1 - e^{-\kappa T}) \approx \kappa B T$). Any $r > 0$ embedded in the discount factor generates a riba loading $\Pi_{\text{riba}} = \kappa B r/(1+r) > 0$ that violates the prohibition of riba.*

Proof. The expected benefit PV under $\iota = 0$ is $B \cdot \mathbb{P}(\tau \leq T) = B(1 - e^{-\kappa T})$. Under $r > 0$, the PV of the benefit is $B e^{-rT} \mathbb{P}(\tau \leq T)$ (continuous) or $B/(1+r) \cdot (1 - e^{-\kappa T})$ (discrete),

so the loading is $B\mathbb{P}(\tau \leq T)(1 - e^{-rT}) \approx B\kappa rT$ for small $\kappa T, rT$. In the one-year discrete case, $\Pi_{\text{riba}} = \kappa Br/(1 + r)$. $\square$

4 CIR Stochastic Hazard Rate Model

4.1 CIR Dynamics

We model the hazard intensity as a Cox–Ingersoll–Ross process:

$$d\kappa_t = \alpha(\bar{\kappa} - \kappa_t) dt + \nu\sqrt{\kappa_t} dW_t, \quad (8)$$

where $\alpha > 0$ is the speed of mean reversion, $\bar{\kappa} > 0$ is the long-run mean, $\nu > 0$ is the volatility coefficient, and W_t is a standard Brownian motion. The Feller condition $2\alpha\bar{\kappa} \geq \nu^2$ ensures $\kappa_t > 0$ a.s.

Agricultural loss rates exhibit mean-reverting behavior: drought years cluster but are followed by recovery periods, rainfall anomalies dissipate, and crop varieties adapt (Barnett and Mahul, 2007). The CIR model captures this reversion while allowing κ to vary stochastically.

4.2 Closed-Form Affine Premium

The key result is that, under the CIR dynamics (8) and $\iota = 0$, the multi-year takaful premium has a closed-form expression via the affine bond price formula.

Theorem 4.1 (CIR Takaful Premium). *Under the CIR dynamics (8) with $\iota = 0$, the fair annual takaful contribution for the policy year commencing at horizon T —the per-year premium covering one year of loss risk at the hazard prevailing at T —is, to first order in the per-period hazard ($\kappa \Delta t \ll 1$ with $\Delta t = 1$ yr):*

$$\pi_{\text{stoch}}^*(T) = \mathbb{E}_0[\kappa_T] \cdot B = \left[\bar{\kappa} + (\kappa_0 - \bar{\kappa})e^{-\alpha T} \right] \cdot B, \quad (9)$$

where κ_0 is the initial hazard rate. The total fair contribution for a continuously-funded T -year contract is the time-integral of this flow,

$$\Pi_{\text{stoch}}^*(T) = \int_0^T \pi_{\text{stoch}}^*(s) ds = B \left[\bar{\kappa} T + (\kappa_0 - \bar{\kappa}) \frac{1 - e^{-\alpha T}}{\alpha} \right]. \quad (10)$$

Proof. Since $\iota = 0$, the discounted benefit equals the undiscounted benefit. The expected hazard rate at horizon T under CIR is $\mathbb{E}_0[\kappa_T] = \bar{\kappa} + (\kappa_0 - \bar{\kappa})e^{-\alpha T}$ (standard CIR conditional mean). The fair one-year contribution charged in the policy year ending at T is $\pi^* = B \cdot \mathbb{E}_0[\mathbb{P}(\tau \leq T + \Delta t \mid \tau > T, \kappa_T)] = B \cdot \mathbb{E}_0[1 - e^{-\kappa_T \Delta t}] \approx B \cdot \mathbb{E}_0[\kappa_T]$ for $\kappa \Delta t \ll 1$, giving (9). Integrating the flow over $[0, T]$ and using $\int_0^T e^{-\alpha s} ds = (1 - e^{-\alpha T})/\alpha$ yields the total (10).

Note $\pi_{\text{stoch}}^*(T)$ is a *rate* (premium per year), not a cumulative T -year charge; the latter is (10). $\square$

Remark 4.1. *The CIR affine formula also gives the survival probability:*

$$S(0, T) = A(T) e^{-\mathcal{B}(T)\kappa_0}, \quad (11)$$

where (with $\mathcal{B}(T)$ the affine bond coefficient, distinct from the benefit B)

$$A(T) = \left[\frac{2\gamma e^{(\alpha+\gamma)T/2}}{2\gamma + (\alpha + \gamma)(e^{\gamma T} - 1)} \right]^{2\alpha\bar{\kappa}/\nu^2},$$

$$\mathcal{B}(T) = \frac{2(e^{\gamma T} - 1)}{2\gamma + (\alpha + \gamma)(e^{\gamma T} - 1)}, \quad \gamma = \sqrt{\alpha^2 + 2\nu^2}.$$

This equals the zero-coupon bond price in the CIR term structure model: the identity of the doubly-stochastic survival probability with the CIR bond price is the classical reduced-form credit result (Lando, 1998; Duffie and Singleton, 1999), and its application to takaful pricing at $\iota = 0$ follows from the formal equivalence with credit pricing established in Ahmed & Bhuyan (2026b).

4.3 Comparison: Constant vs Stochastic κ

Corollary 4.2. *When $\kappa_0 = \bar{\kappa}$ (at long-run mean), $\pi_{\text{stoch}}^* = \bar{\kappa}B = \pi_{\text{const}}^*$. When $\kappa_0 < \bar{\kappa}$ (low current hazard), $\pi_{\text{stoch}}^* < \pi_{\text{const}}^*$: the stochastic model predicts lower near-term premiums but convergence to $\bar{\kappa}B$ over longer horizons. When $\kappa_0 > \bar{\kappa}$ (high current hazard, e.g., after a drought year), $\pi_{\text{stoch}}^* > \pi_{\text{const}}^*$: the stochastic model correctly loads the premium for persistence in current conditions.*

We now turn to estimating these parameters from data.

5 Data

5.1 Cereal Yield Data

We obtain annual cereal yield (kg/ha) for Bangladesh, Pakistan, Indonesia, and India from the World Bank World Development Indicators database (World Bank, 2024) (indicator: AG.YLD.CREL.KG), covering 1980–2023 (44 annual observations per country, 176 total). The indicator aggregates wheat, rice, maize, barley, and other cereals weighted by harvested area, and is consistent with the FAO’s accounting of cereal production as the dominant food-security staple in these economies (FAO, 2020).

Country selection. These four countries represent the largest Muslim-majority agricultural economies in Asia and are the primary markets for agricultural takaful: Bangladesh (Muslim population > 90%, rice-dependent economy), Pakistan (wheat-dominant, major conventional and Islamic insurance market), Indonesia (world’s largest Muslim-majority country, significant rice and palm oil sector), and India (large Muslim minority, world’s largest crop insurance programme under PMFBY).

5.2 Takaful Premium Benchmarks

We compile observed agricultural insurance premiums from government and development-bank publications, using only publicly verifiable primary sources:

- **Bangladesh:** World Bank Technical Assistance pilot 46284-001 (World Bank, 2010), which implies a gross premium of 5.3% for a parametric area-yield product; and the commercial “Agam Bima” potato product (Green Delta / IDLC Finance, 2022), which charges 5–6% of sum insured.
- **Pakistan:** JazzCash / Salaam Takaful (2023) parametric product at PKR 1,500 per PKR 20,000 of coverage (7.5%); and the Punjab Crop Loan Insurance Scheme (CLIS) actuarial study (World Bank, 2020), which documents pure premiums of 5.0% for kharif and 3.5% for rabi crops.
- **Indonesia:** AUTP (Asuransi Usaha Tani Padi) government-mandated rice insurance (Kementerian Pertanian, 2022): gross tariff IDR 180,000 per IDR 6,000,000 of coverage = 3.0%; the government subsidises 80% so farmers pay 0.6%.
- **India:** PMFBY (Pradhan Mantri Fasal Bima Yojana) (GOI, 2016): the actuarial average rate is $\approx 12\%$ nationally (range 10–20% by crop and state); farmer-paid caps are 2% for kharif and 1.5% for rabi, with the central and state governments subsidising the remainder. Both the actuarial rate and the farmer-paid rates are included to illustrate the subsidy wedge.

6 Estimation Methodology

6.1 Loss Year Identification

A year t is classified as a *loss year* if:

$$Y_t < \mu_t^{\text{trend}} - \theta \cdot \sigma_t^{\text{trend}}, \quad (12)$$

where Y_t is cereal yield in year t , μ_t^{trend} and σ_t^{trend} are the 10-year rolling mean and standard deviation (centered), and $\theta = 0.5$ is the threshold parameter. This identifies years in

which yield falls more than half a standard deviation below the local trend. We stress that $\theta = 0.5$ is a modelling choice, not a value mandated by an external standard, and that $\hat{\kappa}$ is materially sensitive to it. On the World Bank cereal-yield panel, sweeping $\theta \in \{0.25, 0.5, 0.75, 1.0\}$ and the window $\in \{8, 10, 12\}$ moves the implied $\hat{\kappa}$ over a 5–15 $\times$ range across the four countries (for India, 0.047–0.29; the headline $\approx 12\%$ is specific to $\theta = 0.5$). The takaful $\hat{\kappa}$ is therefore a threshold-conditional estimate; a robust calibration requires realised loss-cost data (claims per unit sum-insured) rather than a yield-shortfall frequency (Figure 1).




Figure 1: Sensitivity of the yield-shortfall $\hat{\kappa}$ to the loss threshold θ on the World Bank cereal-yield panel (line: mean over the $\{8, 10, 12\}$ -year windows; band: min–max). Across the four countries $\hat{\kappa}$ moves over a 5–15 $\times$ range as θ varies, so the headline ≈ 0.12 (dotted) is specific to $\theta = 0.5$. This threshold sensitivity is why the India PMFBY agreement is read as a units-level coincidence and why a robust calibration must come from realised loss-cost data (cf. Figure 5).

6.2 MLE Hazard Estimation

Given n_ℓ loss years in a sample of T annual observations, the Bernoulli MLE of the annual loss probability is $\hat{p} = n_\ell/T$. Under the intensity interpretation (Poisson approximation), the MLE of κ is:

$$\hat{\kappa}_{\text{MLE}} = -\ln(1 - \hat{p}). \quad (13)$$

For small $\hat{p}$, $\hat{\kappa}_{\text{MLE}} \approx \hat{p}$. The asymptotic variance is $\text{Var}(\hat{\kappa}) \approx \hat{p}/[(1 - \hat{p})T]$ by the delta method (Appendix A).

6.3 CIR Calibration via AR(1) OLS on Rolling $\hat{\kappa}$

Rather than using only four decade-level $\hat{\kappa}$ observations per country, we exploit the full 44-year sample by computing $\hat{\kappa}$ on every rolling 10-year window shifted by one year, yielding 35 estimates per country. We then estimate all three CIR parameters from the data via OLS on the discrete-time AR(1) representation of the CIR process at $\Delta t = 1$ year:

$$\hat{\kappa}_{t+1} = a + b \hat{\kappa}_t + \varepsilon_t, \quad (14)$$

with identification $b = e^{-\alpha}$, $a = \bar{\kappa}(1 - b)$, so that:

$$\hat{\alpha} = -\ln(\hat{b}), \quad \hat{\kappa} = \frac{\hat{a}}{1 - \hat{b}}. \quad (15)$$

Because consecutive rolling windows overlap by nine years, the $\hat{\kappa}$ sequence is serially correlated. Standard errors on $\hat{b}$ are therefore computed using the Newey–West heteroscedasticity-and-autocorrelation-consistent (HAC) estimator with lag = 9 (Newey and West, 1987). The diffusion coefficient ν is recovered from the stationary variance of the $\hat{\kappa}$ series:

$$\hat{\nu} = \sqrt{\frac{2\hat{\alpha} \widehat{\text{Var}}(\hat{\kappa})}{\hat{\kappa}}}, \quad (16)$$

where $\widehat{\text{Var}}(\hat{\kappa})$ is the sample variance across all 35 rolling estimates. The Feller condition $2\hat{\alpha}\hat{\kappa} \geq \hat{\nu}^2$ is checked from the estimated parameters.

7 Empirical Results

7.1 Cereal Yield Trends and Loss Years

Figure 2 plots the 1980–2023 cereal yield series for all four countries, with the 10-year rolling trend and identified loss years.

Table 1: Descriptive Statistics: Cereal Yield and Loss-Year Estimation (1980–2023). N_ℓ = number of loss years; $\hat{p}$ = MLE loss probability; $\hat{\kappa} = -\ln(1 - \hat{p})$ = MLE hazard intensity. The identical India and Pakistan rows are a verified coincidence of the raw World Bank series, not a transcription error: both countries’ 1980–2023 cereal-yield means round to 2394 kg/ha (SDs differ: 633 vs 609) and both record exactly five loss years under the $\theta = 0.5$ threshold, hence the identical $\hat{\kappa} = 0.1206$ discussed — and used as an internal refutation of the PMFBY “match” — in Section 7.4.

Country	T	Mean yield	SD yield	Trend	N_ℓ	$\hat{p}$	$\hat{\kappa}$
Bangladesh	44	3422	1033	Rising	3	0.0682	0.0706
Indonesia	44	4319	775	Rising	6	0.1364	0.1466
India	44	2394	633	Rising	5	0.1136	0.1206
Pakistan	44	2394	609	Rising	5	0.1136	0.1206

Figure 1: Cereal Yield (1980–2023) with Loss Years




Figure 2: Cereal yield (kg/ha) with 10-year rolling trend and identified loss years ($Y_t < \mu_t^{\text{trend}} - 0.5\sigma_t^{\text{trend}}$). Loss years (red dots) capture the most severe yield shortfalls relative to local trend. Bangladesh and India show monotone rising trends driven by Green Revolution varietal improvements; Pakistan and Indonesia exhibit higher volatility reflecting rainfall-dependent agriculture.

All four countries show strong positive trends in cereal yields over the sample period, consistent with global agricultural productivity improvements. Loss years are relatively rare (3–6 per country over 44 years), reflecting the stringent threshold: only the most severe yield contractions qualify.

Table 1 reports the summary statistics.

Table 2: Decade-Level $\hat{\kappa}$ Estimates (1980–2023). $\bar{\kappa}$ = full-sample long-run mean. The variation across decades provides identification for CIR α and ν .

Country	1980s	1990s	2000s	2010s	$\bar{\kappa}$
Bangladesh	0.2231	0.1054	0.0000	0.0000	0.0220
India	0.3567	0.0000	0.1054	0.1054	0.0575
Indonesia	0.2231	0.3567	0.0000	0.1054	0.0999
Pakistan	0.3567	0.1054	0.0000	0.1054	0.0267

7.2 Decade-Level $\hat{\kappa}$ and Mean Reversion

Figure 3 shows the decade-level $\hat{\kappa}$ for each country across four decades.




Figure 3: Rolling 10-year $\hat{\kappa}$ estimates (35 per country, 1989–2023). Each point is the MLE hazard rate from the preceding decade. The dashed line is the AR(1) estimated long-run mean $\bar{\kappa}$; the shaded band is ± 1 stationary standard deviation. The absence of a deterministic trend and the pull of extreme values toward the mean are consistent with the mean-reversion assumption underlying the CIR model.

Table 2 presents the decade-level estimates.

The variation across decades is economically meaningful: no country exhibits a monotone trend in $\hat{\kappa}$, consistent with the mean-reversion assumption. Indonesia shows the highest volatility in $\hat{\kappa}$ across decades, reflecting its exposure to El Niño/La Niña rainfall anomalies. Bangladesh exhibits the lowest $\hat{\kappa}$ (highest food security) due to intensive irrigation and flood-tolerant rice varieties.

7.3 CIR Parameter Estimates

Table 3 reports the CIR parameters.

Table 3: CIR Parameter Estimates: AR(1) OLS on 35 Rolling 10-Year $\hat{\kappa}$ Windows (1980–2023). $\hat{\alpha}$ = mean-reversion speed estimated from data; SE = Newey-West standard error (lag=9); $\hat{\kappa}$ = long-run mean; $\hat{\nu}$ = diffusion coefficient estimated from stationary variance; κ_0 = most recent window (2014–2023). Feller condition $2\hat{\alpha}\hat{\kappa} \geq \hat{\nu}^2$. †Feller violated: reflects binary nature of annual crop-loss events; reflecting boundary imposed in simulation.

Country	N	$\hat{\alpha}$	(SE)	$\hat{\kappa}$	$\hat{\nu}$	κ_0	AR(1) R^2	Feller
Bangladesh	35	0.3012	(1.7142)	0.0220	0.3455	0.0000	0.690	No†
India	35	0.5132	(2.8616)	0.0575	0.3180	0.1054	0.613	No†
Indonesia	35	0.1229	(0.9209)	0.0999	0.2016	0.1054	0.794	No†
Pakistan	35	0.1085	(0.4477)	0.0267	0.3469	0.0000	0.875	No†

Interpretation. Table 3 reports the AR(1) OLS results based on 35 rolling estimates per country. The mean-reversion speed $\hat{\alpha}$ ranges from 0.109 (Pakistan) to 0.513 (India). These estimates are statistically weak: the Newey–West standard errors on $\hat{\alpha}$ exceed the point estimates in every country (e.g. India $\hat{\alpha} = 0.513$ with s.e. 2.86; Pakistan 0.109 with s.e. 0.45), so $\hat{\alpha}$ is not significantly different from zero and the mean-reversion speed is not reliably identified. The high AR(1) R^2 of 0.61–0.88 is in part a mechanical artifact of the 35 rolling windows overlapping by nine of every ten years, rather than independent evidence of mean reversion. We therefore treat the CIR layer as a tractable functional form for multi-period pricing, not as a validated description of the hazard dynamics. The long-run mean $\hat{\kappa}$ ranges from 2.2% (Bangladesh) to 10.0% (Indonesia), reflecting genuine cross-country differences in food-security risk.

The Feller condition $2\hat{\alpha}\hat{\kappa} \geq \hat{\nu}^2$ (Lamberton and Lapeyre, 2011) is violated for all four countries. This is a structural feature of the data rather than a modelling failure: annual crop loss is a binary event (0 or 1), so the rolling $\hat{\kappa}$ series is essentially smoothed binary data and will always exhibit high variance relative to its mean — the defining condition for Feller violation, a structural feature we observe in the agricultural data. In practice we impose $\max(\kappa_t, 0)$ in the Euler–Maruyama simulation (Appendix B), which implements a reflecting boundary at zero and ensures non-negative hazard rates throughout.

Sibling-domain evidence for the CIR- κ form. The weak identification of $\hat{\alpha}$ and the Feller violation here are properties of the input data — a short, annual, near-binary crop-loss series — rather than of the CIR- κ specification itself. In a sibling domain where the same κ is read from a deep, daily, continuously-valued series, the specification is clean: applying the identical mean-reverting κ -dynamics to real sovereign *sukuk* credit spreads, Ahmed & Bhuyan (2026c) recover *statistically significant* mean reversion (bias-corrected half-lives of 1.4–4.3 months at the well-sampled 7-year tenor, significant for the three issuers with the longest histories) with the Feller non-negativity condition $2\alpha\bar{\kappa} \geq \nu^2$ satisfied in *every* one of the twelve constant-maturity series — and, extending to the full universe of 409 outstanding USD *sukuk*, identifiable dynamics for 21 issuers with the

Table 4: Model Premium $\pi^\kappa = \hat{\kappa} \cdot B$ vs Observed Takaful Premiums. $\pi^\kappa/\pi_{\text{obs}} > 1$ indicates over-pricing relative to market; the excess reflects government subsidies, benefit caps, and co-insurance embedded in market products.

Country	Product	Year	B (\$)	π_{obs} (\$)	Rate (%)	$\hat{\kappa}$	π^κ (\$)	$\pi^\kappa/\pi_{\text{obs}}$
Bangladesh	WB Pilot (parametric)	2019	500	26.50	5.30	0.0706	35.30	1
Bangladesh	Agam Bima (potato)	2022	500	27.50	5.50	0.0706	35.30	1
Pakistan	JazzCash Parametric	2023	71	5.33	7.50	0.1206	8.56	1
Pakistan	CLIS Punjab kharif	2020	700	35.00	5.00	0.1206	84.42	2
Pakistan	CLIS Punjab rabi	2020	700	24.50	3.50	0.1206	84.42	3
Indonesia	AUTP rice (gross)	2022	390	11.70	3.00	0.1466	57.17	4
India	PMFBY actuarial avg	2021	450	54.00	12.00	0.1206	54.27	1
India	PMFBY farmer cap (kharif)	2021	450	9.00	2.00	0.1206	54.27	6
India	PMFBY farmer cap (rabi)	2021	450	6.75	1.50	0.1206	54.27	8

Feller condition holding in all of them. The CIR- κ layer is therefore a validated functional form for the convergence intensity; what the agricultural data cannot do is estimate its parameters *precisely*, not cast doubt on the form.

7.4 Model vs Observed Premiums

Figure 4 and Table 4 compare the model-implied premium $\pi^\kappa = \hat{\kappa}B$ against observed takaful premiums.




Figure 4: Left: Model premium $\pi^\kappa = \hat{\kappa} \cdot B$ vs observed market premium π_{obs} . Right: Premium ratio $\pi^\kappa/\pi_{\text{obs}}$ by product. Ratios > 1 indicate that the full-hazard-rate premium exceeds market rates due to government subsidies, deductibles, and co-insurance embedded in conventional products.

Comparing model premiums $\pi^\kappa = \hat{\kappa}B$ against real market rates, India's PMFBY model-implied rate (12.06%) is numerically close to the reported national-average actuarial rate ($\approx 12\%$, ratio ≈ 1.0). We are careful not to overstate this: it is a numerical coincidence

rather than a validation. First, $\hat{\kappa}$ is a loss-event *frequency* (five shortfall years in 44) whereas the PMFBY 12% is a premium *rate*, so the agreement is between quantities of different units. Second, and decisively, Pakistan has the *identical* $\hat{\kappa} = 12.06\%$ (also five shortfall years in 44) yet departs from its observed actuarial rates by $1.6\text{--}3.4\times$; the same model output cannot be a validation for India and simultaneously a $1.6\text{--}3.4\times$ misfit for Pakistan. A genuine test would calibrate κ from realised claims/sum-insured loss-cost data (GOI, 2016; Shekhar and Rai, 2025), which we leave to future work. For Bangladesh, the model over-prices by $1.28\text{--}1.33\times$ relative to observed products, consistent with parametric products covering a subset of total crop risk. For Pakistan and Indonesia, the model-to-market ratios are higher ($1.6\text{--}4.9\times$), reflecting deep government subsidies (Indonesia’s AUTP subsidises 80% of the gross premium; Pakistan’s CLIS operates under a government reinsurance backstop).

This over-pricing relative to subsidised market rates is *not a failure of the model*; it reflects structural features of existing agricultural insurance markets:

1. **Government subsidies.** Indonesia’s AUTP gross premium of 3.0% (IDR 180,000/IDR 6,000,000) is 80% government-subsidised; farmers pay only 0.6% (Kementarian Pertanian, 2022). India’s PMFBY caps farmer contributions at 2% kharif and 1.5% rabi while the actuarial rate averages $\approx 12\%$ (GOI, 2016). Pakistan’s CLIS operates with World Bank/government reinsurance support that reduces effective farmer premiums below actuarial levels (World Bank, 2020). Scaling observed market premiums back to their pre-subsidy actuarial equivalents reduces the model-to-market ratio to near parity for Bangladesh ($\approx 1.3\times$, consistent with a partially-covered parametric product) and India ($\approx 1.0\times$ at the actuarial rate).
2. **Benefit gaps.** Observed products cover partial yield loss (e.g., $> 25\%$ shortfall) with deductibles of 15–25% of the trigger threshold. Our model uses the full benefit B as the payout, overstating actual payouts.
3. **Co-insurance.** Participants often retain 20–30% of the risk. Our π^κ prices full indemnification.
4. **Moral hazard loading.** Conventional actuaries add a moral hazard load of 10–30% to the pure premium. Under the takaful mutual model, moral hazard is partially mitigated by community monitoring, reducing the required load.

Our $\pi^\kappa = \hat{\kappa}B$ represents the *actuarially complete, riba-free premium* that would fully fund the expected benefit without interest earnings on reserves. This is the appropriate riba-free benchmark: market products are cheaper primarily because of subsidies, benefit caps, and co-insurance (items 1–3 above), and additionally — at the pricing layer — because conventional formulas discount at $r > 0$ and rely on interest earned on float (a 7.4% effect at $r = 8\%$; Section 8).

Table 5: Constant- κ vs CIR Stochastic- κ Takaful Premiums ($B = \$500$ benefit). $\pi_{\text{const}} = \hat{\kappa} \cdot B$; $\pi_{\text{stoch}} = \mathbb{E}_0[\kappa_T] \cdot B$ from mean-reverting CIR dynamics, initialised at $\kappa_0 = \hat{\kappa}_{\text{MLE}}$ (full-sample) with Table 3 reversion parameters.

Country	T (yr)	B (\$)	$\hat{\kappa}$	π_{const} (\$)	π_{stoch} (\$)	Δ (\$)	$\Delta/\pi_{\text{const}}$ (%)
Bangladesh	1	500	0.0706	35.30	28.98	-6.32	-17.9
Bangladesh	3	500	0.0706	35.30	20.84	-14.46	-41.0
Bangladesh	5	500	0.0706	35.30	16.38	-18.92	-53.6
India	1	500	0.1206	60.30	47.64	-12.66	-21.0
India	3	500	0.1206	60.30	35.52	-24.78	-41.1
India	5	500	0.1206	60.30	31.18	-29.12	-48.3
Indonesia	1	500	0.1466	73.30	70.60	-2.70	-3.7
Indonesia	3	500	0.1466	73.30	66.10	-7.20	-9.8
Indonesia	5	500	0.1466	73.30	62.58	-10.72	-14.6
Pakistan	1	500	0.1206	60.30	55.48	-4.82	-8.0
Pakistan	3	500	0.1206	60.30	47.26	-13.04	-21.6
Pakistan	5	500	0.1206	60.30	40.65	-19.65	-32.6

Real-data validation (US crop insurance). A units-correct test of the pricing logic is available from real data. Using USDA RMA Summary-of-Business records for 2014–2023 (national, all commodities; \$1.25 trillion cumulative liability, \$119.6 B premium, \$98.3 B indemnity), the realised loss-cost ratio (indemnity/liability) is 7.9% while the premium rate charged (premium/liability) is 9.6% — a loss ratio of 0.82. That premiums closely track realised losses (the market is approximately actuarially fair) is exactly the $\iota = 0$ prediction $\pi^* = \mathbb{E}[\text{loss}]$: the riba-free premium equals the expected loss, the small gap reflecting loading and subsidy rather than an interest discount. In the insurance domain the natural reading of the intensity is direct: the loss-cost ratio *is* the expected-loss rate, so $\hat{\kappa} \approx \text{LCR} = 0.079$ per annum with no recovery adjustment — squarely inside our takaful band (0.07–0.15). (A credit-style rescaling $\hat{\kappa} = \text{LCR}/(1 - \delta)$ can be formed for comparability with the bond papers, but the bond-market recovery convention $\delta = 0.6$ has no actuarial meaning here — an indemnity is a paid claim, not a defaulted bond with salvage — so we report it only as a bracketing exercise: 0.13–0.26 over $\delta \in [0.4, 0.7]$. The fairness conclusion is δ -independent either way.) This is a methodology validation on US federal crop insurance — not South-Asian takaful — but it confirms, on real loss-cost data, that the κ -priced premium is the right order of magnitude and that the $\iota = 0$ fairness relation holds in a functioning insurance market (Figure 5).

7.5 Stochastic vs Constant κ Premium

Table 5 compares the constant- κ and CIR stochastic- κ premiums across horizons.

The premium simulations initialise the hazard at the full-sample MLE, $\kappa_0 = \hat{\kappa}_{\text{MLE}}$ (the level a pricing actuary would charge today), and let the CIR dynamics pull it toward the




Figure 5: The $\iota = 0$ fairness relation $\pi^* \approx \mathbb{E}[\text{loss}]$ on real data. Each marker is a USDA RMA commodity insurance program (2014–2023), the premium rate charged against the realised loss-cost (indemnity/liability); marker area scales with liability. The national aggregate (star) charges a 9.6% premium against a 7.9% realised loss-cost (loss ratio 0.82), sitting just above the actuarially-fair 45° line — the small gap is loading and subsidy, not an interest discount. A riba-free, $\iota = 0$ premium tracking expected losses is exactly the model’s prediction.

AR(1)-implied long-run mean $\hat{\kappa}$. Because $\hat{\kappa}$ lies below $\hat{\kappa}_{\text{MLE}}$ for all four countries and the estimated reversion is fast, the stochastic premium sits below the constant- κ premium at every horizon: by 4–21% at $T = 1$, 10–41% at $T = 3$, and 15–54% at $T = 5$ (Table 5). (The opposite case — a stochastic premium above $\hat{\kappa}B$ when the recent hazard sits below the long-run mean — is possible out of sample but occurs for none of the four countries here.) This convergence effect is important for multi-year sukuk-embedded takaful contracts (see Section 9).

Figure 6 illustrates the CIR hazard rate evolution.

Table 6: Riba Decomposition in Conventional Crop Insurance. The interest loading $\Pi_{\text{riba}} = \pi^* - \pi_{\text{conv}}$ equals the premium reduction from discounting the benefit at $r_{\text{conv}} = 8\%$. Under $\iota = 0$, this loading vanishes and the full actuarial risk premium is retained.

Country	$\hat{\kappa}$	r_{conv}	π^* (riba-free, \$)	π_{conv} (\$)	Riba loading (\$)	Riba share (%)
Bangladesh	0.0706	8%	35.30	32.69	2.61	7.4
Indonesia	0.1466	8%	73.30	67.87	5.43	7.4
India	0.1206	8%	60.30	55.83	4.47	7.4
Pakistan	0.1206	8%	60.30	55.83	4.47	7.4

Figure 4: CIR Hazard Rate Monte Carlo (2,000 paths $\times$ 40 years)




Figure 6: CIR Monte Carlo simulation (2,000 paths, 40-year horizon). Shaded regions: 5–95th (light) and 25–75th (dark) percentiles. Dashed line: long-run mean $\bar{\kappa}$. The process reverts to $\bar{\kappa}$ from any initial κ_0 , providing actuarial stability for long-dated takaful contracts.

8 Riba Loading Decomposition

8.1 Quantifying the Riba Loading

We apply Equation (7) to quantify the riba loading embedded in conventional agricultural insurance pricing at a representative discount rate $r_{\text{conv}} = 8\%$ (consistent with 10-year government bond yields in our sample countries over 2010–2023). Table 6 reports the results.

The riba loading is identical across countries by construction — $r_{\text{conv}}/(1 + r_{\text{conv}}) = 8/108 \approx 7.4\%$ of the riba-free premium (the $r/(1 + r)$ share derived in Section 3) — because

it depends only on the discount rate r , not on κ or B .

8.2 Policy Implications

The 7.4% riba loading implies that:

- At $B = \$500$, a conventional crop insurer retains \$2.6–\$5.4 per policy per year (Table 6) as implicit interest income on the premium reserve — income that Islamic law forbids.
- Takaful operators that use a conventional discount rate in their pricing are effectively charging participants less than the actuarially fair premium, then “making up” the difference through investment income on reserves. This is the fundamental Shariah non-compliance in most current takaful products.
- Under $\iota = 0$ pricing, the premium pool must be large enough to fund claims from contributions alone, without relying on investment returns. This requires $\pi^* = \hat{\kappa}B$, which is the model premium we compute.
- The commercial consequence should be stated plainly rather than left implicit: the fully riba-free contribution is *higher* than the conventionally discounted premium — by exactly the 7.4% loading at $r = 8\%$ — because the participant, not the float, funds the expected claim. A compliant operator therefore either charges $\sim 7\%$ more than a conventional competitor or returns the difference as surplus from a pool that carries no investment income. That is an adoption cost of substance-compliance, not a defect of the model, and marketing a takaful product as “the same price” while discounting at $r > 0$ is precisely the practice this decomposition makes visible.

Figure 7 illustrates the riba loading decomposition across all four countries.




Figure 7: Riba decomposition: riba-free premium π^* (blue) vs conventional discounted premium π_{conv} at $r = 8\%$ (orange). The annotation shows the annual riba loading per \$500 benefit. The gap is $\Pi_{\text{riba}} = \hat{\kappa}B \cdot r / (1 + r)$, constant in percentage terms but rising in dollar terms with $\hat{\kappa}$.

9 Smart Contract Implementation

9.1 TakafulPool Architecture

The companion smart-contract implementation realises agricultural takaful through `TakafulPool.sol`. The prototype reflected in the listing below was deployed on Arbitrum Sepolia at `0xD53d34cC599CfadB5D1f77516E7Eb326a08bb0E4`; the current immutable deployment is `0x63865E7d3d8D9524df1051d291991ce51b391c91`, in which the contribution entrypoint shown here as `join` has been renamed `contribute` and generalised from the hard-coded WBTC/USDC placeholder to a per-pool asset (the $\iota = 0$ tabarru pricing via `quotePut` and the 10% *wakala* fee are unchanged). The contract embeds the $\pi^* = \kappa B$ pricing formula directly in Solidity:

```
// Premium = tabarru. Priced at put option price for the pool's risk layer.
// At iota=0: tabarru = quotePut(refAsset) * coverageAmount / COV_UNIT
function join(bytes32 poolId, uint256 coverageAmount) external {
    uint256 tabarru = IEverlastingOption(everlastingOption)
        .quotePut(WBTC, usdc) * coverageAmount / COV_UNIT;
    IERC20(usdc).safeTransferFrom(msg.sender, address(this), tabarru);
    contributions[poolId][msg.sender].coverageAmount += coverageAmount;
    contributions[poolId][msg.sender].tabarru += tabarru;
    emit Contributed(poolId, msg.sender, coverageAmount, tabarru);
}
```

The `quotePut` function returns the `EverlastingOption` put price at $\iota = 0$. In the

deployed testnet instance the option leg is parameterised on the WBTC/USDC market as a *placeholder* for the risk layer; a production agricultural pool would quote the put on the yield-index token itself, whose $\iota = 0$ price equals the expected indexed loss — κB for short maturities, consistent with Proposition 3.1.

9.2 Claims Settlement

Claims are filed via `fileClaim(poolId, amount)`. The contract verifies that the claimant has sufficient coverage (`contributions[poolId][msg.sender].coverageAmount >= amount`) and pays from the pool’s USDC balance. The *wakala* fee is retained automatically at claim time:

```
function fileClaim(bytes32 poolId, uint256 amount) external {
    uint256 wakala = amount * WAKALA_FEE_BPS / 10000;
    uint256 net = amount - wakala;
    // wakala stays in pool (operator’s fee)
    IERC20(usdc).safeTransfer(msg.sender, net);
    emit ClaimPaid(poolId, msg.sender, net, wakala);
}
```

The 10% *wakala* fee (`WAKALA_FEE_BPS = 1000`) is consistent with AAOIFI Shariah Standard No. 26 (AAOIFI, 2010), which permits operator fees of 10–30% of contributions.

9.3 On-Chain Pricing vs CIR Model

The on-chain premium uses the EverlastingOption `quotePut` price, which is derived from the Black–Scholes-like formula at $\iota = 0$ (Ackerer, Hugonnier, and Jermann, 2025, Proposition 6) (see also Ahmed & Bhuyan, 2026a, for the $\iota = 0$ derivation). The CIR model in this paper provides the *multi-year actuarial basis* for selecting the volatility parameter σ in the option price, which should be set to $\sigma \leftarrow \nu\sqrt{\bar{\kappa}}$ (the instantaneous volatility of κ at its long-run mean $\bar{\kappa}$; the stationary standard deviation of the CIR process itself is $\nu\sqrt{\bar{\kappa}/(2\alpha)}$).

This connection between the on-chain EverlastingOption pricing and the CIR hazard model is the key architectural insight of this framework: the smart contract implements a continuous-time risk market in which κ is implicitly priced through the options layer, without requiring oracles to report κ directly.

10 Alignment with AAOIFI Standard No. 26: A Self-Assessment

This paper settles the *riba* question at the pricing layer ($\iota = 0$) only. *Gharar*, *maysir*, and *qabd* are product-layer determinations reserved to Shariah supervisory boards; no fatwa

is claimed, and the assessment below is the authors’ self-assessment, offered as input to such a board. The *tabarru* pool structure instantiates the risk-sharing rule (R4) of our companion boundary analysis (Ahmed & Bhuyan, 2026d), with the yield index addressing the asset-linkage and insurable-interest rules (R1, R3).

10.1 Relevant Standards

Agricultural takaful products must comply primarily with:

- **AAOIFI Shariah Standard No. 26** (Islamic Insurance – Takaful): The takaful pool must be structured as a mutual guarantee (*ta’awun*), not a commercial transaction. Premiums must represent genuine *tabarru* (donation). Operator may charge *wakala* fees of 10–30%; underwriting surplus belongs to participants, not the operator (AAOIFI, 2010). Parametric and index-based triggers are arguably consistent with the objective-verifiability requirement when the index is publicly auditable (Mahul and Stutley, 2010); the determination rests with Shariah supervisory boards.

10.2 Self-Assessment of the κ -Framework

Riba elimination. The $\iota = 0$ constraint eliminates all interest from the pricing formula. Reserves are held in a fiat-referenced stablecoin (USDC) on-chain, or tokenized gold (PAXG); we note that the Shariah status of fiat-backed stablecoins, whose own reserves include interest-bearing instruments, is an unresolved product-layer question referred to Shariah supervisory boards. Investment income on reserves is distributed to participants, not retained by the operator — consistent with AAOIFI No. 26.

Gharar considerations. Index-based parametric takaful (triggered by yield below a threshold) is designed to reduce gharar by making the loss trigger objective and publicly verifiable. The World Bank yield index (or equivalent satellite data) provides a transparent, manipulation-resistant trigger — superior in this respect to indemnity products that rely on self-reported farm losses.

Maysir considerations. The mutual pool is structured so that contributions represent risk-sharing (*tabarru*) rather than a bilateral bet. The κ -pricing formula avoids systematic under-funding in expectation (no speculative under-pricing); the full-risk premium $\pi^* = \hat{\kappa}B$ funds expected claims without reliance on interest income. Solvency against adverse loss realisations additionally requires a surplus buffer or retakaful layer, which we leave to future work (Section 11).

Transparency. The smart contract implementation provides complete on-chain transparency: all contributions, claims, and *wakala* fees are publicly auditable on Arbitrum (block explorer: <https://sepolia.arbiscan.io>).

11 Conclusion

We have derived the actuarial properties of the κ -rate under stochastic hazard, estimated them on real agricultural data, and validated the $\iota = 0$ pricing logic on real (USDA RMA) loss-cost data, with application to agricultural takaful pricing in South and Southeast Asia. Our main findings are:

1. The MLE hazard intensity $\hat{\kappa}_{\text{MLE}} \in [0.071, 0.147]$ for the four countries studied, with Bangladesh the lowest (most food-secure) and Indonesia the highest (most exposed to climate variability).
2. The CIR model provides a tractable closed-form premium formula for multi-year takaful: $\pi_{\text{stoch}}^*(T) = \mathbb{E}_0[\kappa_T] \cdot B$, which converges to $\bar{\kappa}B$ as $T \rightarrow \infty$. Because the AR(1)-implied long-run mean lies below the full-sample $\hat{\kappa}_{\text{MLE}}$ for all four countries, the stochastic premium sits below the constant- κ premium at every horizon — by 4–21% at $T = 1$ and 10–41% at $T = 3$, widening to 15–54% at $T = 5$ (Table 5).
3. The model-implied India rate is numerically close to the PMFBY actuarial average ($\hat{\kappa} = 12.06\%$ vs PMFBY $\approx 12\%$), but we treat this as a coincidence, not a validation: it compares a loss frequency with a premium rate, and Pakistan’s identical $\hat{\kappa}$ departs from its observed rates by 1.6–3.4 $\times$. For subsidised market products, the model premium exceeds observed rates by 1.3–8 $\times$, explained by the subsidy and benefit-gap structure documented in Section 7. The full-hazard-rate premium $\pi^* = \hat{\kappa}B$ is the appropriate riba-free benchmark: it funds expected claims without relying on interest earnings on reserves.
4. On real USDA RMA crop-insurance data (2014–2023), the premium rate charged (9.6%) tracks the realised loss-cost (7.9%, loss ratio 0.82), validating the $\iota = 0$ fairness relation $\pi^* \approx \mathbb{E}[\text{loss}]$ on a functioning insurance market.
5. The riba loading in conventional crop insurance pricing equals $\Pi_{\text{riba}} = \hat{\kappa}B \cdot r / (1 + r) \approx 7.4\%$ of the riba-free premium at $r = 8\%$. Eliminating this loading requires $\iota = 0$ pricing, which is enforced at the smart contract level.
6. The TakafulPool smart contract (Arbitrum Sepolia) implements the $\iota = 0$ pricing formula through the `EverlastingOption.quotePut` function, providing a fully on-chain, manipulation-resistant actuarial premium.

Future work should extend the framework to: (i) satellite-based parametric triggers (NDVI, precipitation anomalies) for sub-national crop loss estimation; (ii) retakaful layer pricing using the κ -rate framework; (iii) multi-peril products (flood + drought) with copula-based joint loss distributions; (iv) mainnet deployment with live oracle feeds from Chainlink or Pyth; (v) surplus-buffer sizing for solvency against adverse loss realisations, and asset-linked reserve management via vault-secured tokenized gold and tokenized sukuk, extending the hedging utility of the pool to inflation and rate-environment risk.

References

- Accounting and Auditing Organisation for Islamic Financial Institutions (2010). *Shariah Standard No. 26: Islamic Insurance (Takaful)*. AAOIFI, Manama, Bahrain.
- Ackerer, D., Hugonnier, J., and Jermann, U. J. (2025). Perpetual futures pricing. *Mathematical Finance*. DOI: 10.1111/mafi.70018.
- Ahmed, S., & Bhuyan, R. (2026a). Interest parameter in perpetual futures: A Shariah-compliant decentralised exchange. *SSRN Working Paper 6322778*. https://papers.ssrn.com/sol3/papers.cfm?abstract_id=6322778
- Ahmed, S., & Bhuyan, R. (2026b). Random stopping time as credit event: A riba-free credit pricing framework. *SSRN Working Paper 6322858*. https://papers.ssrn.com/sol3/papers.cfm?abstract_id=6322858
- Ahmed, S., & Bhuyan, R. (2026c). The κ -yield curve: Construction, estimation methodology, and real-data evidence from sovereign sukuk. *SSRN Working Paper 6322938*. https://papers.ssrn.com/sol3/papers.cfm?abstract_id=6322938
- Ahmed, S., & Bhuyan, R. (2026d). The jurisprudential boundaries of an interest-free substrate: Riba, gharar, maysir, and qabd in cash-settled synthetic instruments. Working paper (companion research note).
- Barnett, B. J. and Mahul, O. (2007). Weather index insurance for agriculture and rural areas in lower-income countries. *American Journal of Agricultural Economics*, 89(5):1241–1247.
- Cahyasita, D., Aminda, F., and Kusumawardani, M. (2025). Farmers’ motivations and barriers to participation in the rice farming insurance program (AUTP) in Kulon Progo Regency, Indonesia. *International Journal of Agricultural Social Economics and Rural Development*, 5(1). DOI: 10.37149/ijaserd.v5i1.2208.
- Cox, J. C., Ingersoll, J. E., and Ross, S. A. (1985). A theory of the term structure of interest rates. *Econometrica*, 53(2):385–407.
- Dal Moro, E. (2024). Parametric crop insurance pricing based on bimodal yield distribution. SSRN Working Paper. DOI: 10.2139/ssrn.4871230.
- Duffie, D. and Singleton, K. J. (1999). Modeling term structures of defaultable bonds. *Review of Financial Studies*, 12(4):687–720.
- Encyclopedia of Islamic Insurance, Takaful and Retakaful (2019). Takaful and conventional insurance: Dichotomy. In *Encyclopedia of Islamic Insurance, Takaful and Retakaful*. Edward Elgar Publishing. DOI: 10.4337/9781788115834.00013.

- Fakhruzi, L. and Syafril, R. (2025). Implementation of the rice farming business insurance program (AOTP) in Junjung Sirih District, Solok Regency. *Santhet: Jurnal Sejarah, Pendidikan dan Humaniora*, 9(1). DOI: 10.36526/santhet.v9i1.5025.
- Food and Agriculture Organization (2020). *The State of Food and Agriculture 2020*. FAO, Rome.
- Government of India (2016). *Pradhan Mantri Fasal Bima Yojana: Operational Guidelines*. Ministry of Agriculture and Farmers Welfare (DAC&FW), New Delhi. Updated 2021.
- Islamic Financial Services Board (2024). *Islamic Financial Services Industry Stability Report 2024*. IFSB, Kuala Lumpur.
- Kementerian Pertanian (2022). *Pedoman Bantuan Premi Asuransi Usaha Tani Padi (AOTP)*. Peraturan Menteri Pertanian No. 40/2015 (updated 2022). Indonesian Ministry of Agriculture, Jakarta.
- Kramer, B., Porter, M., and Wassie, S. B. (2025). Basis risk, social comparison, perceptions of fairness, and demand for insurance: A field experiment in Ethiopia. *Journal of Risk and Insurance*. DOI: 10.1111/jori.70015.
- Lamberton, D. and Lapeyre, B. (2011). *Introduction to Stochastic Calculus Applied to Finance*, 2nd ed. Chapman and Hall/CRC, London.
- Lando, D. (1998). On Cox processes and credit risky securities. *Review of Derivatives Research*, 2(2–3):99–120.
- Mahul, O. and Stutley, C. J. (2010). *Government Support to Agricultural Insurance: Challenges and Options for Developing Countries*. World Bank, Washington DC. DOI: 10.1596/978-0-8213-8217-2.
- Negi, D. and Ramaswami, B. (2023). Basis risk and the demand for catastrophic rainfall insurance. *Q Open*, 3(1). DOI: 10.1093/qopen/qoae009.
- Newey, W. K. and West, K. D. (1987). A simple, positive semi-definite, heteroscedasticity and autocorrelation consistent covariance matrix. *Econometrica*, 55(3):703–708.
- Obaidullah, M. and Khan, T. (2008). *Islamic Microfinance: Development Challenges and Initiatives*. Islamic Research and Training Institute, Jeddah.
- Outreville, J. F. (2012). The relationship between insurance, financial development and economic growth: a cross-country analysis. *Risk Management and Insurance Review*, 15(1):79–103.

- Punia, M., Kundu, D., and Mehla, D. (2021). Crop insurance in India: Status of PMFBY against different crop insurance schemes. *The Pharma Innovation*, 10(4S):5976. DOI: 10.22271/tpi.2021.v10.i4sb.5976.
- Roberts, R. A. J. (2005). Insurance of crops in developing countries. *FAO Agricultural Services Bulletin*, 159.
- Shekhar, C. and Rai, P. (2025). Crop insurance in a changing climate: Resilience, exclusion, and reform in PMFBY and RWBCIS. SSRN Working Paper. DOI: 10.2139/ssrn.5283931.
- Şimşek, G. and Yıldırak, K. (2025). Moral hazard analysis for crop yield insurance using loss prevention model. *Sosyoekonomi*, 33(63):77–98. DOI: 10.17233/sosyoekonomi.2025.01.04.
- Skees, J. R. and Barnett, B. J. (2006). Enhancing microfinance using index-based risk transfer products. *Agricultural Finance Review*, 66(2):235–250.
- Swiss Re Institute (2023). *Sigma 1/2023: Natural Catastrophes and Inflation in 2022*. Swiss Re, Zurich.
- World Bank (2010). *Agricultural Insurance in Bangladesh: Promoting Access to Small and Marginal Farmers*. World Bank, Washington DC. DOI: 10.1596/12624.
- World Bank (2020). *Punjab Crop Loan Insurance Scheme: Actuarial Review*. World Bank Group, Washington DC.
- World Bank (2024). World Development Indicators. <https://databank.worldbank.org/source/world-development-indicators>. Accessed March 2026.

A MLE Consistency Under Bernoulli Loss Process

Proposition A.1. *Under the Bernoulli loss model $L_t \sim \text{Bernoulli}(p)$, the MLE $\hat{p} = n_\ell/T$ is consistent and asymptotically normal:*

$$\sqrt{T}(\hat{p} - p) \xrightarrow{d} \mathcal{N}(0, p(1 - p)).$$

The delta method gives $\hat{\kappa}_{\text{MLE}} = -\ln(1 - \hat{p})$ with:

$$\sqrt{T}(\hat{\kappa} - \kappa) \xrightarrow{d} \mathcal{N}\left(0, \frac{p}{1 - p}\right) = \mathcal{N}(0, e^\kappa - 1).$$

Proof. Standard delta method: $g(p) = -\ln(1 - p)$, $g'(p) = 1/(1 - p)$, so $\text{Var}(g(\hat{p})) \approx [g'(p)]^2 \cdot p(1 - p)/T = p/[(1 - p)T]$. $\square$

B Euler–Maruyama Discretisation of CIR

The discrete-time approximation used in Monte Carlo:

$$\kappa_{t+1} = \max\left(\kappa_t + \alpha(\bar{\kappa} - \kappa_t)\Delta t + \nu\sqrt{\max(\kappa_t, 0)}\varepsilon_t\sqrt{\Delta t}, 0\right), \quad (17)$$

where $\varepsilon_t \sim \mathcal{N}(0, 1)$ i.i.d. and $\Delta t = 1$ year. The reflection at zero approximates the Feller boundary condition. With 2,000 paths and $T = 40$ years, the Monte Carlo standard error on the median κ_T estimate is less than 0.5% of $\bar{\kappa}$.

C AAOIFI Shariah Standard No. 26: Design Self-Assessment Checklist

The table below is the authors’ *self-assessment* of the TakafulPool design against AAOIFI Standard No. 26, offered as input to a Shariah supervisory board; formal determination of conformity rests with such a board, and no fatwa is claimed.

Requirement	Self-assessment	Implementation
Mutual guarantee structure	Consistent	TakafulPool is participant-owned; operator earns <i>wakala</i> fee only
<i>tabarru</i> as donation	Consistent	<code>join()</code> records contribution as <i>tabarru</i> ; no exchange contract
<i>wakala</i> fee $\leq 30\%$	Consistent	WAKALA_FEE_BPS = 1000 (10%)
Surplus to participants	Consistent	Pool surplus remains in TakafulPool; claimable by participants
No riba in reserves	Partial	Reserves held in USDC (or PAXG); the underlying reserves of fiat-backed stablecoins include interest-bearing instruments — an unresolved product-layer question for the board
Gharar: objective trigger	Consistent	On-chain oracle provides yield index; trigger is algorithmic; residual gharar determination rests with the board

Requirement	Self-assessment	Implementation
Maysir	Consistent	Mutual pool structure; contributions as <i>tabarru</i> , not a bilateral bet; final determination rests with the board
Retakaful arrangement	Partial	Retakaful layer not yet implemented; excess risk retained in pool
Investment income distribution	Partial	USDC does not itself distribute yield; future: deployment of idle reserves into screened instruments (e.g. tokenized sukuk or commodity murabaha), subject to board approval

Table 7: AAOIFI Standard No. 26 design self-assessment for TakafuPool.sol. **Green** = self-assessed consistent; **Yellow** = partial/planned. Formal determination rests with a qualified Shariah supervisory board.

CRedit Author Contribution Statement. S.A.: conceptualisation, methodology, software, data curation, formal analysis, writing (original draft). R.B.: supervision, writing (review and editing), resources.

Conflict of Interest Statement. S.A. is the principal architect of the smart contracts described in this paper and holds governance tokens in the project. R.B. declares no conflicts of interest.

Data Availability Statement. All data used in this study are publicly available. Cereal yield data are from the World Bank World Development Indicators (<https://databank.worldbank.org/source/world-development-indicators>, indicator AG.YLD.CREL.KG). All estimation scripts (`fetch_data.py`, `takaful_pipeline.py`) and generated tables and figures are available at <https://github.com/Arcus-Quant-Fund/BarakaDapp>. No proprietary data sources are required to reproduce any result in this paper.

EDITORIAL NOTE TO THIS PRINTING

Paper 2c · The Islamic CDS

Working paper, revised July 2026 · SSRN 6323519

Role in the programme. Credit protection from the primitive: a par spread $s^* = \kappa(1 - \delta)$ containing no interest rate, taken all the way to running code. **What it establishes.** To our knowledge the first iCDS deployed on public blockchain infrastructure (a testnet prototype); the first instrument-level recovery evidence from the sukuk default record, with realised recovery on asset-based corporates near $\delta \approx 0.10$; a six-jurisdiction regulatory map. **Corrected or extended later.** A version-1 premium bug (a large overcharge) is documented and corrected in the contract itself; the description of the companion curve methodology was refreshed in July 2026, and the first-claims now carry “to our knowledge”. **Not established here.** Maysir, gharar and qabd for the cash-settled synthetic; the paper treats it as the hardest case, last.

Replication for every number in this paper: DOI 10.5281/zenodo.21290557 · github.com/shehzadahmed-xx/kappa-replication

An Islamic Credit Default Swap on Smart-Contract Infrastructure: Design, Pricing, and Regulatory Pathway

Shehzad Ahmed¹ and Rafiqul Bhuyan^{*2}

¹Department of Finance, Independent University Bangladesh (IUB), Dhaka,
Bangladesh. 2120922@iub.edu.bd

²Department of Accounting and Finance, Alabama A&M University,
Normal, AL, USA. rafiqul.bhuyan@aamu.edu

June 2026

^{*}Corresponding author. Department of Accounting and Finance, Alabama A&M University, Normal, AL 35762, USA. E-mail: rafiqul.bhuyan@aamu.edu.

Abstract

We introduce the first Islamic Credit Default Swap (iCDS) deployed on public blockchain infrastructure (a testnet prototype), combining a closed-form riba-free pricing formula with a deployed smart contract on Arbitrum. Building on the κ -rate framework of Akerer, Hugonnier & Jermann (AHJ) [7] and the credit-equivalence theorem of Ahmed & Bhuyan [9], the fair iCDS spread is $s^* = \kappa(1 - \delta)$, where κ is the convergence intensity and δ is the recovery rate. A standard property of flat-intensity credit pricing is that this *par running spread is independent of the discount rate*: because the protection and premium legs both terminate at the default time, the discount factor cancels between them, so a correctly-priced conventional CDS quotes the same fair spread $\kappa(1 - \delta)$ at any maturity. The par spread is thus automatically riba-neutral — a clean instance of the Separation Theorem. The riba is real but lives in *present-value* terms, where the discount does not cancel: we define the present-value riba premium $\Pi_{\text{riba}}^{\text{PV}}(T)$ — the discounting haircut on the future loss. It increases in ι , vanishes at $\iota = 0$, and stays under 4% of notional at realistic 5–10 yr sukuk maturities. The iCDS additionally enforces riba-freeness structurally: collateral is held in a *tabarru* pool that earns no interest on the margin float; full collateralisation structurally addresses the *maysir* objection, and an objective on-chain credit-event oracle narrows *gharar* — the final product-layer determination of these two pillars rests with Shariah scholars.

On a cross-section of 62 real, actively-quoted sovereign USD sukuk (12 issuers, live Cbonds), the κ -spread $s^* = \kappa(1 - \delta)$ recovers a clean credit ordering ($\hat{\kappa}$ from 0.7% to 9%). The genuinely independent test: κ read from those sukuk G -spreads recovers the hazard that independent instruments imply — sukuk and same-sovereign conventional bonds agree within 6–16 bp for the deeply-covered issuers, and sukuk co-move with sovereign CDS at correlation 0.97. Assembling the first instrument-level recovery evidence from the Cbonds default record—the complete population of defaulted USD sukuk, five of 800 issues—we find realized recovery on asset-based corporate sukuk an order of magnitude below the 0.60 sovereign convention (on the cases with disclosed post-default prices, $\delta \approx 0.1$; Serba Dinamik ≈ 5 cents, NMC and Swiber liquidated), which vindicates pricing the recovery as a per-contract parameter rather than a fixed constant. We then deploy and benchmark a working on-chain iCDS.

The iCDS smart contract (iCDS.sol) is deployed on Arbitrum Sepolia in its corrected immutable version at 0x77a98785...521b; the full open $\rightarrow$ accept $\rightarrow$ settle lifecycle is Foundry-measured at $\approx 897,000$ gas for a one-year protection with three premium payments — a few US cents at L2 gas prices. The deployed version corrects a premium-flow overcharge in the initial v1 prototype (0x7A02...b15F). We conduct a six-jurisdiction legal analysis identifying UAE/VARA and Bahrain/CBB as the most permissive pilot frameworks, and design a governance-token liquidity-bootstrapping

mechanism with a Nash-equilibrium participation guarantee.

Keywords: Islamic finance, credit default swap, convergence intensity, smart contracts, AAOIFI, *riba*, κ -rate, Arbitrum, DeFi, regulatory analysis.

JEL Codes: G12, G13, G18, K22, Z12.

Contents

1	Introduction	6
2	Theoretical Framework	8
2.1	The κ -Rate and AHJ No-Arbitrage at $\iota = 0$	8
2.2	iCDS Pricing Formula: $s^* = \kappa(1 - \delta)$	8
2.3	Comparison with Conventional CDS: The Par Spread is Riba-Neutral	9
2.4	The riba premium in present-value terms	9
2.5	Where the structural riba is	10
2.6	The substance-compliant form: pricing at realised hazard κ_P	11
3	Data	12
3.1	Real sukuk cross-section from Cbonds	12
3.2	Realized recovery: direct evidence from the sukuk default record	13
3.3	Independent CDS benchmark	14
3.4	SOFR (for the present-value premium only)	14
4	Pricing Analysis	14
4.1	Descriptive Statistics: the real sukuk cross-section	14
4.2	Real-Data Pricing: the iCDS on Observed Sukuk	15
4.3	Welfare Analysis	19
5	Smart Contract Implementation	19
5.1	Architecture	19
5.2	Full Lifecycle	21
5.3	Dynamic Premium Mechanics and Shariah Compliance	22
5.4	Credit Event Oracle: Design Options	23
5.5	Gas Cost Benchmarks	24
6	Legal Analysis	25
6.1	Assessment against AAOIFI Shariah Standards	25
6.2	ISDA/IIFM Tahawwut Documentation	27
6.3	Jurisdictional Analysis	27
6.4	Recommended Pilot: UAE / VARA	29
7	Liquidity Bootstrapping	30
7.1	Market Structure	30

7.2	Governance Token Incentive Mechanism	30
7.3	Participation Condition and Market Viability	30
7.4	Liquidity Milestones	31
8	Systemic Risk	32
8.1	Lessons from 2008	32
8.2	iCDS Safeguards	32
8.3	Circuit Breakers and Position Limits	33
9	Conclusion	33
A	iCDS.sol Key Functions (Solidity)	36
B	Real Constant-Maturity $\hat{\kappa}$ Dynamics	38
C	Draft iCDS Term Sheet (ISDA/IIFM Tahawwut Format)	38

1 Introduction

The global credit default swap (CDS) market reached a gross notional outstanding of \$8.8 trillion by end-June 2024 [13]. Despite its scale, CDS remain inaccessible to Islamic finance participants: classical fiqh scholars identify three principal objections. First, *maysir* (gambling): protection can be bought *naked* — without holding the reference obligation — so the contract becomes a speculative bet on default rather than the hedge of a real exposure [2, 5]; on the other side, sellers can write protection without posting collateral against the contingent liability. Second, *gharar* (excessive uncertainty): credit event definitions rely on opaque committee determinations rather than observable events [1, 3]. Third, *riba* (interest): scholars argue that the fixed CDS spread — quoted as an annual percentage of notional — functions economically as an interest payment, on the view that it embeds the prevailing risk-free rate rather than only the credit risk of the reference entity [4, 6]. A contribution of this paper is to sharpen that objection: we show the *par spread* is in fact discount-rate-independent (Section 2), and locate the genuine *riba* in the present value and in the contract and collateral structure.

The Accounting and Auditing Organisation for Islamic Financial Institutions [12] has not issued a standard explicitly permitting conventional CDS. The ISDA/IIFM Tahawwut Master Agreement [20] provides documentation for over-the-counter Islamic hedging transactions but does not specify a Shariah-compliant premium formula for credit derivatives. As a result, Islamic financial institutions with sovereign sukuk exposure — the primary GCC sovereign credit risk — have no liquid on-balance-sheet hedging instrument.

This paper addresses all three objections through three interlocking contributions: it settles the *riba* question at the pricing layer, and provides structural safeguards against *maysir* and *gharar* whose sufficiency is a determination for qualified Shariah scholars.

Contribution 1: Closed-form *riba*-free pricing. Ahmed & Bhuyan [9] prove that under the Akerer, Hugonnier & Jermann [7] no-arbitrage framework with $\iota = 0$, the fair credit protection spread is:

$$s^* = \kappa(1 - \delta), \tag{1}$$

where κ is the convergence intensity (the *riba*-free analogue of the risk-free rate) and δ is the recovery rate. A standard property of flat-intensity credit pricing is that this fair spread is *independent of the discount rate*: a correctly-priced conventional perpetual CDS quotes the same fair spread $\kappa(1 - \delta)$, because both legs terminate at default and the discount factor cancels between them. The iCDS therefore does not undercut the conventional spread; it matches it. We validate this on a cross-section of 62 real Cbonds sovereign sukuk and against

independent instruments (sovereign CDS and same-issuer conventional bonds); and we locate the riba not in the par spread (which is riba-neutral at every maturity) but in two places the iCDS controls: the *present value*, via the premium $\Pi_{\text{riba}}^{\text{PV}}(T)$ that discounting applies to the future loss, and the *contract structure* — the interest paid on posted collateral, the admission of naked positions, and determination discretion.

Contribution 2: Smart-contract implementation. We describe and test `iCDS.sol`, a Solidity smart contract deployed on Arbitrum Sepolia at `0x77a98785...521b`. The contract addresses each fiqh objection mechanically: (a) maysir is addressed on both sides — the seller must post full notional as collateral (no uncovered selling), and the buyer is required to hold the reference obligation with attestation (insurable interest); (b) gharar is narrowed by defining the credit event as an on-chain oracle condition ($\text{spot} \leq \text{recovery floor}$), transparent and objectively verifiable; (c) riba is removed by holding the seller’s collateral in a *tabarru* (donation) pool that earns no contractual interest, and by reserving at $\iota = 0$ via the everlasting-put formula, so no predetermined return accrues on the deposited principal. Whether the structural safeguards in (a) and (b) suffice is a product-layer determination for Shariah scholars. The full iCDS lifecycle is $\approx 897,000$ gas (Foundry-measured) — a few US cents on Arbitrum, an order of magnitude cheaper than Ethereum mainnet, making micro-notional Islamic credit hedging economically viable (Section 5).

Contribution 3: Regulatory pathway. We provide the first comprehensive legal analysis of an on-chain iCDS across six jurisdictions (UAE/VARA, Malaysia/BNM, Singapore/MAS, UK/FCA, US/CFTC, Bahrain/CBB) and identify the UAE’s Virtual Assets Regulatory Authority [22] as the most immediately permissive framework for commercial launch. We further design a governance-token liquidity-bootstrapping mechanism with a game-theoretic proof that the dominant strategy for capital allocators converges to participation when the governance-token reward rate exceeds a threshold we derive in closed form.

Organisation. Section 2 develops the theoretical framework. Section 3 describes the data. Section 4 presents the pricing analysis and the validation of the κ -spread against observed CDS. Section 5 describes the smart contract implementation. Section 6 conducts the jurisdictional analysis. Section 7 designs the liquidity-bootstrapping mechanism. Section 8 addresses systemic risk. Section 9 concludes.

2 Theoretical Framework

2.1 The κ -Rate and AHJ No-Arbitrage at $\iota = 0$

Ackerer, Hugonnier & Jermann [7] introduce a general continuous-time no-arbitrage framework for perpetual contracts in which the *convergence intensity* $\kappa_t > 0$ takes over the discounting role of the conventional risk-free rate. Under this framework, the time- t price of any security is:

$$P(t) = \mathbb{E}_t \left[\int_t^\infty e^{-\int_t^u (\kappa_s + \iota) ds} C(u) du \right], \quad (2)$$

where $C(u)$ is the cash-flow rate at time u , κ_s is the convergence intensity, and $\iota \geq 0$ is the interest parameter. At $\iota = 0$ the no-arbitrage price exists, is unique, and is discounted by κ alone [7]; Ahmed & Bhuyan [9] show that this removes the *riba* (interest) term from the pricing kernel.

Ahmed & Bhuyan [9] establish the identification:

$$\hat{\kappa} = \frac{s}{1 - \delta}, \quad (3)$$

where s is the observed sukuk spread and δ is the recovery rate. This follows from the $\kappa \leftrightarrow \lambda$ isomorphism established in the credit-equivalence theorem of Ahmed & Bhuyan [9] (Separation Theorem), which embeds the reduced-form credit framework of Duffie & Singleton [18]. Ahmed & Bhuyan [8] estimate $\hat{\kappa}$ from a panel of 7 OIC sovereign spreads over 2012–2025 and fit CIR dynamics [16] via AR(1) OLS on 35 rolling 10-year windows with Newey-West HAC standard errors:

$$d\kappa_t = \hat{\alpha}(\hat{\kappa} - \kappa_t) dt + \hat{\nu} \sqrt{\kappa_t} dW_t. \quad (4)$$

2.2 iCDS Pricing Formula: $s^* = \kappa(1 - \delta)$

Theorem 2.1 (Fair iCDS Spread, Ahmed & Bhuyan [9], Proposition 3). *Under AHJ no-arbitrage with $\iota = 0$, the fair annualised spread for a perpetual Islamic credit default swap with loss-given-default $LGD = 1 - \delta$ is:*

$$s^* = \kappa(1 - \delta). \quad (5)$$

Equation (5) has an immediate actuarial interpretation: s^* is the expected loss rate per unit time under the risk-neutral measure, with no interest-loading component. The premium is dynamic: as the oracle spot price of the reference entity falls toward the recovery floor, the everlasting put price Π_{put} rises, and the quarterly premium automatically increases to

reflect deteriorating credit quality (Section 5.3).

2.3 Comparison with Conventional CDS: The Par Spread is Riba-Neutral

Applying the same stopping-time pricing to a *conventional* CDS—one that discounts cash flows at the risk-free rate $\iota > 0$ —yields the classical result that the *par running spread* is independent of ι . Both legs of a CDS terminate at the default time τ : the protection leg pays the loss $(1 - \delta)$ at τ , and the premium leg accrues only while the reference name survives. Under a flat hazard κ and flat rate ι , the protection PV is

$$V_{\text{prot}} = (1 - \delta) \mathbb{E}[e^{-\iota\tau}] = (1 - \delta) \frac{\kappa}{\kappa + \iota}, \quad (6)$$

and the premium leg is discounted *and* survival-weighted, so the correct annuity is the risky PV01

$$A = \mathbb{E}\left[\int_0^\tau e^{-\iota u} du\right] = \mathbb{E}\left[\frac{1 - e^{-\iota\tau}}{\iota}\right] = \frac{1}{\iota} \left(1 - \frac{\kappa}{\kappa + \iota}\right) = \frac{1}{\kappa + \iota}, \quad (7)$$

giving $V_{\text{prem}} = sA = s/(\kappa + \iota)$. Equating the legs, the common factor $1/(\kappa + \iota)$ cancels:

$$s_{\text{conv}} = \kappa(1 - \delta) = s^* \quad \text{for every } \iota \geq 0. \quad (8)$$

The same cancellation holds at finite maturity T , where $V_{\text{prot}}(T) = (1 - \delta) \frac{\kappa}{\kappa + \iota} (1 - e^{-(\kappa + \iota)T})$ and $A(T) = \frac{1 - e^{-(\kappa + \iota)T}}{\kappa + \iota}$: the survival factor $(1 - e^{-(\kappa + \iota)T})$ cancels and $s^*(T) = \kappa(1 - \delta)$ for all T . *The par spread is riba-neutral by construction.* A naive derivation divides the protection PV by the *riskless* perpetuity annuity $1/\iota$ instead of the risky annuity $1/(\kappa + \iota)$ — letting the premium leg pay forever rather than knocking out at default. This manufactures a spurious ι -dependence and forces $s \rightarrow 0$ at low rates, inconsistent with the 48–100 bp observed on investment-grade sovereigns at near-zero policy rates in 2012–2018. Equation (8) dissolves that puzzle: the fair spread never depended on the rate. The identification $\hat{\kappa} = s/(1 - \delta)$ holds by construction, so the iCDS reproduces the observed market spread rather than undercutting it.

2.4 The riba premium in present-value terms

Equation (8) shows the riba is *not* in the par spread. It is real, but it lives in *level / present-value* terms, where the discount does not cancel. The protection PV at $\iota = 0$ is the full undiscounted loss $1 - \delta$; discounting at $\iota > 0$ shrinks it. The difference is the present-value riba premium:

Definition 2.1 (Present-value riba premium).

$$\Pi_{\text{riba}}^{\text{PV}} \equiv V_{\text{prot}}^{\iota=0} - V_{\text{prot}}^{\iota} = (1 - \delta) \left(1 - \frac{\kappa}{\kappa + \iota} \right) = (1 - \delta) \frac{\iota}{\kappa + \iota}, \quad (9)$$

with finite-maturity form

$$\Pi_{\text{riba}}^{\text{PV}}(T) = (1 - \delta) \left[(1 - e^{-\kappa T}) - \frac{\kappa}{\kappa + \iota} (1 - e^{-(\kappa + \iota)T}) \right]. \quad (10)$$

This is a one-time present value (a fraction of notional), *not* a running spread in bp/yr. It has the correct comparative statics: it vanishes as $\iota \rightarrow 0$ (no discounting, no haircut), rises to $(1 - \delta)$ as $\iota \rightarrow \infty$, and is monotonically increasing in ι — more discounting means a larger haircut, the intuitive direction. The iCDS at $\iota = 0$ pays the full undiscounted $V_{\text{prot}}^{\iota=0} = 1 - \delta$; the conventional valuation applies the haircut $\Pi_{\text{riba}}^{\text{PV}}$. This is a clean instance of the Separation Theorem [9]: the par spread isolates pure credit risk and is automatically riba-neutral, while the entire interest-rate (riba) content is confined to the present-value level.

2.5 Where the structural riba is

The riba a conventional CDS carries therefore enters not through the fair spread but through the contractual and collateral structure, which the iCDS strips out at each point:

- **Collateral float.** In a conventional CDS the protection seller's posted margin is remunerated at the interest rate ι —a predetermined return on a deposited principal, i.e. *riba*. The iCDS holds collateral in a *tabarru* (donation) pool that earns no contractual interest.
- **Naked positions and leverage.** Conventional CDS permit both uncollateralised protection selling (the AIG failure mode) and naked protection buying without insurable interest — a zero-sum wager (*maysir*); the iCDS requires full collateralisation of the seller, so every contract is backed by a segregated reserve, and conditions buyer participation on attested holding of the reference obligation.
- **Determination discretion.** The credit event is settled by an objective, public on-chain oracle condition rather than an opaque determinations committee, controlling *gharar*.

The iCDS and a correctly-priced conventional CDS therefore differ only in contract and collateral structure — where riba, maysir and gharar reside.

2.6 The substance-compliant form: pricing at realised hazard κ_P

Equation (5) prices the spread at the *risk-neutral* intensity: writing $\kappa_Q \equiv s/(1 - \delta)$ for the market-priced intensity, $s^* = \kappa_Q(1 - \delta)$ is *riba-free* in ι (no interest loading, §2.4) but carries the market’s *risk premium*. Decomposing the intensity into the physically realised hazard κ_P (from the actual sukuk default record, §3.2) and the priced κ_Q , the wedge $1 - \kappa_P/\kappa_Q$ is small for realised-hazard products but the great majority for high-grade credit, where the realised default rate is far below the priced spread. A jurisprudential objection reserved to scholars—whether a systematic, priced risk premium stripped of predetermined interest is *riba in substance*—attaches to that wedge, not to the interest content $\iota = 0$ has already removed.

This admits a product-layer form that does not require the objection to be resolved first. Alongside the tradeable, κ_Q -priced iCDS—the market benchmark, and the natural form should a board rule the contingent premium permissible—the protocol supports a *cooperative-guarantee* form whose contribution is set at the realised hazard,

$$s_{\text{coop}}^* = \kappa_P(1 - \delta), \quad (11)$$

the actuarially fair expected loss, exactly as a *takaful* contribution is priced. The gap between the two forms is the risk-premium wedge $(\kappa_Q - \kappa_P)(1 - \delta)$: in the tradeable form it is a spread *sold* to a protection writer; in the cooperative form it is not sold at all. Participants mutually indemnify realised default loss; the *tabarru* pool that already holds the collateral retains any surplus and returns it to members; and the operator takes only a disclosed *wakala* fee for administration. With the insurable-interest and possession conditions already specified—the buyer holds the reference obligation, with *qabd hukmi* over the token—the cooperative form is structurally *takaful on default-loss*: a mutual guarantee (*kafala ta‘awuniyya*) over an owned, possessed exposure, funded at expected loss rather than at a priced premium. It thereby meets, at the product layer, both the contract-form objections (*bay‘ al-dayn*, *qabd*) and the substance objection (the priced premium), leaving as the reserved residual only whether tokenised title is valid *qabd* and whether the *wakala* loading is a genuine service fee—questions we place before qualified scholarship rather than answer. The two forms are complementary: the κ_Q -priced instrument is the correct market benchmark and the tradeable case; the κ_P -priced cooperative form is the one that needs no prior ruling on the open question.

3 Data

3.1 Real sukuk cross-section from Cbonds

All pricing in this paper runs on *real*, actively-quoted sukuk. The primary panel is an instrument-level cross-section of 62 sovereign USD sukuk across 12 issuers, pulled live from the Cbonds Bond Screener (account 972307; `/api/bonds/search/` with the sovereign, USD, sukuk, outstanding filters). For each bond we read the traded G -spread s over the interpolated US-Treasury curve and form $\hat{\kappa} = s/(1 - \delta)$ with $\delta = 0.60$. Because these are bullet (option-free) sovereign sukuk, the G -spread is the clean credit spread ($G \approx Z \approx \text{OAS}$). The implied $\hat{\kappa}$ runs from 0.7% (Qatar, Malaysia) through 1.8–2.1% for the investment-grade GCC/Indonesia names to 5–9% for the high-yield sovereigns (Bahrain, Egypt, Pakistan).

For dynamics we use the roll-down-free constant-maturity $\hat{\kappa}$ histories of the companion κ -yield-curve study [8] (Cbonds daily series, five issuers, 2018–2026): bias-corrected mean-reversion half-lives of 1.4–4.3 months for the four 7y series (10.5 months for Turkey 5y), with the Feller condition satisfied in every series (Appendix B, Table 9).

Recovery and the sukuk default record. The iCDS spread scales with the loss-given-default $(1 - \delta)$, so the recovery assumption $\delta = 0.60$ deserves an empirical anchor. The Cbonds default record provides one, and it is striking: of the entire global sukuk universe, **no sovereign sukuk has ever entered redemption default**, while **17 corporate sukuk have** (Cbonds global default-record pull: Malaysia 12, Singapore 2, Türkiye 1, Kuwait 1, UAE 1 — all corporate) — including the USD issues of NMC Health (UAE), Serba Dinamik (Malaysia, two issues), Bank Asya (Türkiye) and IIG (Kuwait), the five that also appear in the 800-instrument USD-universe pull, plus local-currency cases such as Swiber (Singapore, SGD) and Country Garden (Malaysia, MYR). Two implications follow. First, for the *sovereign* iCDS the recovery rate is genuinely unobserved — there is no sovereign-sukuk default from which to estimate it — so $\delta = 0.60$ is taken from the standard Moody’s senior-unsecured sovereign convention and its effect is bracketed by the recovery-rate sensitivity analysis of the companion study ($\delta \in \{0.40, 0.50, 0.70\}$ rescales every $\hat{\kappa}$ by a common factor) [8]; the clean sovereign record is itself consistent with the low- $\hat{\kappa}$, low-default-intensity regime the framework assigns to investment-grade sovereigns. Second, corporate sukuk *do* default, so recovery is first-order for a *corporate* iCDS; realized corporate-sukuk recovery is not publicly tabulated in clean form, which is precisely why the contract exposes δ as an explicit, governable parameter (`recoveryRateWad`, Section 5.3) rather than hard-coding it. (Cbonds, a bond database, carries no single-name CDS series; the conventional-market benchmark used below is therefore the Damodaran sovereign CDS panel, corroborated at

the instrument level by the matched conventional-bond spreads of the companion study [8].)

3.2 Realized recovery: direct evidence from the sukuk default record

The $\delta = 0.60$ baseline is the Moody’s senior-unsecured *sovereign* convention, and for the sovereign iCDS it is the only available anchor (no sovereign sukuk has defaulted). For the *corporate* iCDS, where recovery is first-order and default has occurred, we assemble the first instrument-level recovery evidence directly from the Cbonds default record. The defaulted population is small because sukuk default is rare: of the entire 800-bond USD-sukuk universe, exactly five issues have entered redemption default—Bank Asya, IIG (International Investment Group), NMC Health, and Serba Dinamik’s two issues—a 0.6% historical default rate, and this is the complete USD population, not a selected sample. Including the MYR and SGD record, Cbonds flags 17 redemption-defaults in all; realized recovery—measured where a workout figure is disclosed, otherwise by the post-default indicative (distressed) price, the market’s recovery estimate—is documented for six (Table 1). The remaining USD case, the 2012 IIG default (Kuwait, \$200 m), has no disclosed post-default price in our sources and is therefore recorded as a default without a recovery figure rather than assigned one.

Issuer (USD/SGD sukuk)	Default	Recovery	Basis
Serba Dinamik 6.30% '22	2022	$\approx 5\%$	distressed price (4.6–5.2c); wound up
Serba Dinamik 7.00% '25	2022	$\approx 5\%$	distressed price; wound up
NMC Healthcare 5.95% '23	2020	low	administration; debt-to-equity
Swiber Capital (2 sukuk)	2016	low	liquidation; unsecured
Bank Asya 7.5% '23	2015	$\approx$ par	<i>sovereign</i> backstop (Turkey)

Table 1: Realized recovery on defaulted sukuk (Cbonds default record; distressed-price proxy where no workout figure is disclosed). Asset-based corporate sukuk recover an order of magnitude below the 0.60 sovereign convention; the lone near-par recovery is a government guarantee, not asset recourse.

Three facts stand out. First, realized recovery on *asset-based* corporate sukuk is an order of magnitude below the 0.60 convention: the Serba Dinamik dollar sukuk traded at 4.6–5.2 cents at default and the issuer was wound up; NMC Healthcare entered administration with a debt-to-equity restructuring; Swiber was liquidated—across these cases recovery clusters near 10% ($\delta \approx 0.1$), not 60%. Second, the one near-par recovery, Bank Asya, came not from asset recourse but from a Turkish sovereign backstop of the bank’s foreign debt—the exception that proves the rule. Third, the structural reason is documented: van Wijnbergen & Zaheer [23] show that in default most sukuk (the *asset-based* majority—AAOIFI estimates

85% of GCC issuance) do not transfer the underlying assets to holders and are treated as *unsecured, subordinated* debt of the obligor, so holders “often had nothing to resort to”; only true-sale asset-*backed* structures (e.g. the US East Cameron sukuk, where the assets were assigned to holders) recover through the collateral.

The implication sharpens the contract design of Section 5. The spread $s^* = \kappa(1 - \delta)$ scales with loss-given-default $(1 - \delta)$, so a corporate iCDS priced at the sovereign $\delta = 0.60$ (LGD 0.40) against a realized asset-based LGD near 0.90 would *underprice* protection by roughly $0.90/0.40 \approx 2.3\times$. This is precisely why `iCDS.sol` exposes δ as a per-contract governable parameter (`recoveryRateWad`, Section 5.3) rather than hard-coding the sovereign convention: the recovery a buyer should price is near zero for asset-based corporate names and materially higher only for true-sale asset-backed or sovereign-backstopped issues. The clean sovereign record and the low asset-based corporate record are two faces of the same credit-intensity ordering the κ -framework predicts.

3.3 Independent CDS benchmark

For the non-circular validation of Section 4.2 we use sovereign CDS as an *independent* instrument against which to check the sukuk-implied $\hat{\kappa}$. Annual sovereign CDS spreads are sourced from Damodaran [17]; they are used only as a comparison benchmark (a different market), *not* as a sukuk-spread proxy that feeds the $\hat{\kappa}$ inputs.

3.4 SOFR (for the present-value premium only)

The present-value riba premium $\Pi_{\text{riba}}^{\text{PV}}(T)$ of Section 2.4 is evaluated at the discount rate $\iota = \text{SOFR}$; we use the 2025 annual mean SOFR of 4.24% (FRED [19]). SOFR enters *only* the present-value premium, never the par spread s^* , which is discount-independent.

4 Pricing Analysis

4.1 Descriptive Statistics: the real sukuk cross-section

Table 2 summarises the κ -spread on the *real* sovereign USD sukuk cross-section — 62 actively-quoted bonds across 12 sovereigns, pulled live from the Cbonds Bond Screener. The fair spread $s^* = \hat{\kappa}(1 - \delta)$ equals the observed sukuk G -spread by the identification $\hat{\kappa} = s/(1 - \delta)$, and the implied $\hat{\kappa}$ runs from 0.7% (Qatar, Malaysia) through 1.8–2.1% for the GCC/Indonesia investment-grade names to 5–9% for the high-yield sovereigns (Bahrain, Egypt, Pakistan)

Sovereign	n sukuk	median G (bp)	$\hat{\kappa}$ (%)	range G (bp)
Qatar	1	29	0.71	29–29
Hong Kong	1	32	0.79	32–32
Malaysia	4	37	0.92	13–50
Oman	2	71	1.77	64–78
Saudi Arabia	12	72	1.80	50–78
Philippines	1	72	1.81	72–72
Indonesia	21	83	2.06	26–294
Turkey	6	144	3.59	34–232
Bahrain	10	215	5.38	202–275
Egypt	2	245	6.13	195–295
Benin	1	256	6.41	256–256
Pakistan	1	365	9.12	365–365
All	62	78	1.94	13–365

Table 2: Descriptive statistics on the *real* sovereign USD sukuk cross-section (Cbonds Bond Screener, account 972307; 62 actively-quoted bonds, 12 sovereigns). $\hat{\kappa}$ from the median G -spread via $\hat{\kappa} = s/(1 - \delta)$, $\delta = 0.60$. The credit ordering (Qatar/Malaysia tight $\rightarrow$ Pakistan wide) is read directly from the sukuk market, with no synthetic panel and no CDS proxy.

— a clean, monotone credit ordering read directly from the sukuk market, with no synthetic panel and no CDS proxy.

4.2 Real-Data Pricing: the iCDS on Observed Sukuk

On the real sukuk cross-section described above, pricing the iCDS is direct: for a sovereign whose traded sukuk carries G -spread s over Treasury, the riba-free iCDS spread is, by the identification $\hat{\kappa} = s/(1 - \delta)$, exactly

$$s_{t=0}^* = \hat{\kappa}(1 - \delta) = s, \quad (12)$$

i.e. the iCDS charges the issuer’s *own observed sukuk spread* — the full expected loss — with no proxy and no benchmark rate.

Independent validation (non-circular). Because $s^* \equiv s$ by the identification, recovering the sukuk spread is an inversion identity, not a test. The genuine test asks whether $\hat{\kappa}$ read from the *sukuk market* agrees with the hazard implied by a *different*, independent instrument. We run it two ways on live Cbonds data (sovereign USD, outstanding; account 972307). First, against the CDS market: across the five sovereigns with both traded sukuk and a CDS quote, the median sukuk G -spread and the sovereign CDS co-move with correla-

tion 0.97 (Table 3), the sukuk trading at a stable basis of $\approx 0.70\times$ the CDS. Second — the cleaner test, entirely within the bond market — against the same sovereign’s conventional bonds: for the two issuers with deep sukuk coverage (Saudi Arabia, $n=12$; Indonesia, $n=21$), the sukuk and conventional G -spreads agree within 6–16 bp, the sukuk marginally tighter (Table 4). In both cases $\hat{\kappa}$ from sukuk recovers the hazard an independent instrument prices.

Sovereign	n sukuk	sukuk G -spread (bp)	$\hat{\kappa}$ (%)	CDS (bp)
Bahrain	10	215	5.38	290
Indonesia	21	83	2.06	120
Malaysia	4	37	0.92	80
Qatar	1	29	0.71	58
Saudi Arabia	12	72	1.80	65

Table 3: **Non-circular validation.** $\hat{\kappa}$ read from *real* Cbonds USD sovereign sukuk G -spreads (sukuk market) against sovereign CDS (CDS market) — two independent instruments. Across the 5 sovereigns with both, the median sukuk G -spread and the CDS co-move with correlation 0.97 (OLS slope 0.76, $R^2 = 0.94$), the sukuk trading at a stable basis of $\approx 0.70\times$ the CDS. Because it compares $\hat{\kappa}$ from one market (sukuk) to a benchmark from another (CDS) — rather than regressing a series on itself — the agreement is genuine evidence, not an inversion identity. Source: Cbonds Bond Screener (account 972307); $\delta = 0.60$.

Sovereign	n sukuk	sukuk G (bp)	n conv.	conv. G (bp)	sukuk–conv. (bp)
Saudi Arabia	12	72	30	78	-6
Indonesia	21	83	51	98	-16
Philippines [†]	1	72	40	94	-21
Qatar [†]	1	29	15	53	-24
Hong Kong [†]	1	32	3	6	+26

Table 4: **Non-circular validation (same market).** Median G -spread of *real* Cbonds USD sovereign sukuk vs the same sovereign’s *real* conventional USD bonds. For the two sovereigns with deep sukuk coverage (Saudi Arabia $n=12$, Indonesia $n=21$) the sukuk and conventional spreads agree within 16–6 bp, the sukuk trading marginally tighter — i.e. $\hat{\kappa}$ read from sukuk matches $\hat{\kappa}$ from conventional bonds, two independent instruments in one market. [†]Single-issue sukuk: medians are a single bond and reflect maturity mismatch, not a basis. Source: Cbonds Bond Screener (account 972307), live pull; $\delta = 0.60$.

Present-value riba on the real universe. The riba content lives in the present value (Section 2.4): Table 5 reports both the par κ -spread and the maturity-matched present-value riba premium $\Pi_{\text{riba}}^{\text{PV}}(T)$ across the rating spectrum of the real sukuk universe.




Figure 1: Non-circular validation: $\hat{\kappa}$ from *real* Cbonds sovereign sukuk *G*-spreads (sukuk market) against independent sovereign CDS (CDS market). $n=5$ sovereigns; Cbonds account 972307; $\delta = 0.60$.

Rating	$s^* = \hat{\kappa}(1 - \delta)$ (bp)	$\hat{\kappa}$ (%)	$\Pi_{\text{riba}}^{\text{PV}}(5\text{y})$ (% notional)	$\Pi_{\text{riba}}^{\text{PV}}(10\text{y})$ (% notional)	$\Pi_{\text{riba}}^{\text{PV}}(\infty)$ (% notional)
A (IG)	36	0.90	0.17	0.63	33.0
BBB (IG)	76	1.90	0.35	1.24	27.6
BB (HY)	196	4.90	0.83	2.65	18.6
B (HY)	216	5.40	0.90	2.83	17.6
CCC (HY)	296	7.40	1.15	3.43	14.6

Table 5: The par spread vs. the present-value riba premium on real sovereign sukuk, by rating bucket (median *G*-spread from the 62-bond Cbonds cross-section; $\delta = 0.60$; $\iota = \text{SOFR}_{2025} = 4.24\%$). The fair κ -spread $s^* = \hat{\kappa}(1 - \delta)$ is riba-neutral and identical for the iCDS and a correctly-priced conventional CDS (Eq. 8). The riba is confined to the present value: $\Pi_{\text{riba}}^{\text{PV}}(T)$ (Eq. 10) is the haircut interest discounting applies to the future loss, as a fraction of notional. At realistic 5–10yr sukuk maturities it is small (well under 4%); the perpetual figure $\Pi_{\text{riba}}^{\text{PV}}(\infty) = (1 - \delta)\iota/(\kappa + \iota)$ is an *upper bound*, inflated because $1/\kappa$ pushes the certain loss decades into the discount.

Three points are worth emphasising. First, the iCDS charges the same actuarially fair $\kappa(1 - \delta)$ as a correctly-priced conventional CDS, with no riba premium in the par spread at any maturity. Second, the riba is nonetheless real — it is the present-value haircut $\Pi_{\text{riba}}^{\text{PV}}(T)$ that discounting applies to the future loss (Eq. 10); at realistic 5–10 yr sukuk maturities it is small (under 4% of notional even at CCC), and the perpetual figure in the last column is an upper bound, inflated because $1/\kappa$ pushes the certain loss decades into the discount. Figure 2 traces $\Pi_{\text{riba}}^{\text{PV}}(T)$ across maturities. What additionally makes the iCDS riba-free at the contract level is the structure of Section 2. Third, the long-run iCDS price built from the *real* constant-maturity $\hat{\kappa}$ — Malaysia 36 bp, Saudi Arabia 92 bp, Indonesia 112 bp, Bahrain 209 bp — recovers the same credit ordering from an independent dynamic estimation, and the real mean-reversion half-life of a few months for the investment-grade issuers means the dynamic premium of Section 5.3 re-prices to new credit information within a quarter, not a year. The iCDS mechanism is therefore not only theoretically riba-free but quantitatively calibratable from observed sukuk markets today.




Figure 2: The present-value riba premium $\Pi_{\text{riba}}^{\text{PV}}(T)$ (Eq. 10) as a function of maturity, by rating bucket, at $\iota = \text{SOFR}_{2025} = 4.24\%$ and $\delta = 0.60$. Dotted lines are the perpetual upper bounds $(1 - \delta)\iota/(\kappa + \iota)$; the shaded band is the typical 5–10 yr sukuk tenor, where the premium is a small fraction of notional. Within the plotted range the premium is larger for riskier credits; the ordering reverses only at maturities far beyond any relevant tenor ($\gtrsim 85$ yr), where small κ defers the certain loss deep into the discount — the $T \rightarrow \infty$ limit reflected in the perpetual bounds (dotted), which are larger for safer credits.

4.3 Welfare Analysis

We compute the certainty-equivalent wealth (CEW) for a risk-averse investor holding \$10 M of Saudi Arabia sovereign sukuk with a 5-year horizon under two strategies: (1) unhedged and (2) hedged at the fair spread $s^* = \hat{\kappa}(1 - \delta)$ (the iCDS, which a correctly-priced conventional CDS would match). We use the *real* Saudi Arabia sukuk $\hat{\kappa} = 1.80\%$ (median Cbonds G -spread 72 bp).

Under CRRA utility $u(W) = W^{1-\gamma}/(1-\gamma)$, the 5-year default probability is $1 - e^{-\hat{\kappa}_{SA} \cdot 5} = 8.6\%$. The fair hedging premium over 5 years is $\hat{\kappa}(1 - \delta) \cdot W_0 \cdot T = 1.80\% \times 0.40 \times \$10M \times 5 = \$360,000$.

Table 6 reports CEW for $\gamma \in \{2, 5, 10\}$. For all risk aversion levels, hedging at the fair spread increases CEW relative to the unhedged position: the 8.6% default probability and 40% LGD generate a welfare cost that hedging eliminates. The iCDS delivers this protection with no interest accruing on the posted collateral; its advantage over a conventional CDS is structural (riba-free collateral, full backing, transparent settlement) at the *same* fair price, not a difference in spread.

Table 6: Welfare Analysis: Certainty-Equivalent Wealth for Saudi Arabia 5yr Sukuk.

γ	Fair premium	CEW: Unhedged	CEW: Hedged
2	\$360,000	\$9,457,345	\$9,640,000
5	\$360,000	\$8,922,167	\$9,640,000
10	\$360,000	\$7,791,049	\$9,640,000

Notes: CRRA utility $u(W) = W^{1-\gamma}/(1-\gamma)$. Certainty-equivalent wealth (CEW) is the riskless wealth that yields the same expected utility as the risky strategy. Reference entity: Saudi Arabia sovereign; notional $W_0 = \$10,000,000$; $T = 5$ years; $\kappa = \hat{\kappa}_{SA} = 1.80\%$ (real Cbonds sukuk median G -spread, 72 bp); $\delta = 0.60$. The fair hedging premium $= \hat{\kappa}(1-\delta) \cdot W_0 \cdot T$ is the discount-independent κ -spread; a correctly-priced conventional CDS charges the same, so “Hedged” applies to both the iCDS and a fairly-priced CDS. The iCDS’s advantage is structural (riba-free collateral), not a lower price.

5 Smart Contract Implementation

5.1 Architecture

`iCDS.sol` is a Solidity 0.8.24 smart contract deployed on Arbitrum Sepolia (chain ID 421614) at address `0x77a98785...521b`. It inherits from OpenZeppelin’s `Ownable2Step`, `Pausable`, and `ReentrancyGuard`, providing access control, emergency halt, and reentrancy protection. Two immutable dependencies are set at deployment:




Figure 3: Certainty-equivalent wealth by hedging strategy for CRRA investors ($W_0 = \$10M$, Saudi Arabia 5yr sukuk, real $\hat{\kappa} = 1.80\%$, $\delta = 0.60$, $T = 5$ yr). Hedging at the fair spread dominates the unhedged position for all γ .

- `evOption (0x8B45...dEC8)`: the `EverlastingOption` contract implementing AHJ [7], Proposition 6 put pricing at $\iota = 0$.
- `oracle (0x4f6C...2999)`: the `OracleAdapter` providing real-time and historical asset prices via Chainlink feeds.

Each protection is stored as a `Protection` struct with 11 fields:

```

1 struct Protection {
2     address seller;           // protection seller (posted notional)
3     address buyer;           // protection buyer (pays premiums)
4     address refAsset;         // reference market asset (oracle key)
5     address token;           // collateral / premium token (USDC)
6     uint256 notional;         // max payout (raw token units)
7     uint256 recoveryRateWad; // delta in WAD (e.g. 60e16 = 60%)
8     uint256 recoveryFloorWad; // put strike = recoveryRate *
                               // spot_at_open
9     uint256 tenorEnd;         // protection expiry (Unix timestamp)
10    uint256 lastPremiumAt;    // timestamp of last premium payment
11    uint256 premiumsCollected; // cumulative premiums paid to seller
12    Status status;           // Open | Active | Triggered | Settled |
                               // Expired
13 }
```

Listing 1: Protection struct (iCDS.sol)

5.2 Full Lifecycle

Figure 4 illustrates the complete contract lifecycle.

Step 1 – openProtection. The seller calls `openProtection(refAsset, token, notional, recoveryRateWad, tenorDays)`. The contract deposits `notional` tokens from the seller into escrow (anti-naked-CDS rule). The *recovery floor* is computed once:

$$\text{recoveryFloor} = \text{recoveryRateWad} \times \text{oracle.getIndexPrice(refAsset)} / \text{WAD}, \quad (13)$$

locking the CDS strike at the spot price prevailing when protection is opened.

Step 2 – acceptProtection. The buyer calls `acceptProtection(id)`. The *first premium* is immediately computed and transferred to the seller. The fair per-period premium implied by the everlasting-put reserve is the flow relation of the companion credit-equivalence paper (Ahmed & Bhuyan [9], the takaful-flow theorem $\pi = \kappa \Pi$):

$$\text{premium}_0 = \hat{\kappa} \cdot \frac{\Pi_{\text{put}}(\text{spot}, \text{recoveryFloor}) \times \text{notional}}{\text{WAD}} \cdot \Delta t, \quad \Delta t = \frac{90}{365}, \quad (14)$$

where Π_{put} is the everlasting put price from `EverlastingOption.quotePut()`, implementing the closed-form two-region formula of Akerer, Hugonnier & Jermann [7], Proposition 6 at $\iota = 0$ (Appendix A).¹

Step 3 – payPremium (quarterly, dynamic). Every 90 days, the buyer pays:

$$\text{premium}_t = \hat{\kappa} \cdot \frac{\Pi_{\text{put}}(\text{spot}_t, \text{recoveryFloor}) \times \text{notional}}{\text{WAD}} \cdot \Delta t. \quad (15)$$

Because Π_{put} rises as spot_t falls toward `recoveryFloor`, the premium automatically rises as credit quality deteriorates – a built-in early-warning feedback mechanism absent from fixed-spread conventional CDS.

Step 4 – triggerCreditEvent. An authorised keeper calls `triggerCreditEvent(id)` once the oracle confirms $\text{spot}_t \leq \text{recoveryFloor}$. This replaces the opaque ISDA credit-events committee with an objective, on-chain, oracle-verified condition, directly addressing the gharar objection.

¹*Deployment note.* An earlier prototype (0x7A02...b15F) charged the *unscaled* put price $\Pi_{\text{put}} \times \text{notional} / \text{WAD}$ per period — a conservative overcharge of order $1/(\hat{\kappa} \Delta t)$ relative to the fair flow above. The deployed immutable contract above implements the fair flow, and the gas and lifecycle figures reported here are measured on it (0x77a98785...521b).

Step 5 – settle. The buyer calls `settle(id)`. Loss-given-default is distributed:

$$\text{LGD payout} = \text{notional} \times (1 - \delta); \quad \text{sellerReturn} = \text{notional} \times \delta. \quad (16)$$

The seller is not penalised beyond the actual LGD – consistent with the mutual takaful principle that financial burden is shared according to the pre-agreed formula, not punitive.

Step 6 – expire. If no credit event occurs by `tenorEnd`, the seller reclaims the full notional collateral via `expire(id)`.

Open	→	Active	→	Triggered	→ Settled
Seller posts notional		Buyer accepts; pays 1st prem.		Keeper confirms $\text{spot} \leq \text{floor}$	Buyer receives LGD payout
			↓ no default at maturity		
			Expired: seller reclaims notional		

Figure 4: iCDS lifecycle state machine. Transitions: Open → Active (`acceptProtection`); Active → Triggered (`triggerCreditEvent`); Triggered → Settled (`settle`); Active → Expired (`expire` after `tenorEnd`).

5.3 Dynamic Premium Mechanics and Shariah Compliance

The dynamic put-pricing premium eliminates *riba* structurally, not by relabelling a fixed spread.² In a conventional CDS, the fixed spread is set at inception and paid regardless of subsequent credit-quality changes; the protection seller earns interest on the posted margin in the interim, which is the *riba* element (Section 2).

The iCDS premium re-prices at each payment date based on current oracle prices, and the seller’s collateral sits in a non-interest-bearing *tabarru* pool. Because $\iota = 0$ in the pricing formula and no interest accrues on the float, there is no *riba* component. The premium paid at time t is the pure risk-neutral expected loss over the next quarter – the actuarially fair takaful contribution (*tabarru*) for that specific risk period, priced as if default could occur at any instant with no preference given to near-term versus far-term losses. This is consistent with the AAOIFI takaful model where each participant contributes according to their contemporaneous risk exposure.

²The theoretical formula $s^* = \kappa(1 - \delta)$ is derived for a perpetual CDS in the AHJ framework (Theorem 2.1, Section 2). The contract implementation uses a finite tenor (`tenorDays`) with quarterly premium resets, which is consistent: each reset re-prices the quarterly premium from the current everlasting put price via the flow relation $\pi = \kappa \Pi_{\text{put}}$ of the companion paper [9]. The put price itself is a present value — the reserve covering all future protection — not a rate; multiplying by κ converts that reserve into the fair per-unit-time flow. This dynamic structure approximates perpetual pricing at each payment date without requiring the contract itself to be perpetual.



Figure 5: Dynamic iCDS pricing on a real sukuk credit history. Black (left axis): the iCDS fair spread $s^*(t) = \kappa(t)(1 - \delta)$ on Indonesia’s real 7y constant-maturity κ (Cbonds, 2018–2026), spanning the 2020 COVID blow-out and the 2022 rate shock. Red (right axis): the present-value riba premium $\Pi^{\text{PV}}(t, 7y)$ at $\iota = \text{SOFR}$. Both re-price each period to observed credit; $\delta = 0.60$.

Dynamic re-pricing on real credit history. Figure 5 makes the re-pricing concrete on real data. Feeding the contract Indonesia’s actual 7-year constant-maturity κ history (Cbonds daily G -spreads, 2018–2026), the iCDS fair spread $s^*(t) = \kappa(t)(1 - \delta)$ tracks the issuer’s credit through the cycle — blowing out to 368 bp at the 2020 COVID peak ($\kappa \approx 9.2\%$) and compressing back toward 60–90 bp as spreads normalised — while the present-value riba premium $\Pi^{\text{PV}}(t, 7y)$ moves with it between 0.6% and 2.3% of notional. The contract is therefore not a static-spread instrument: it marks the protection premium to current credit each period, exactly as a takaful pool re-assesses contributions, and does so with no interest term anywhere in the path.

5.4 Credit Event Oracle: Design Options

The credit event oracle is the central engineering challenge for an on-chain credit derivative. We identify four options, ordered by decentralisation:

- (A) **Chainlink decentralised vote.** Node operators vote on whether a sovereign default event has occurred. Fully on-chain and decentralised, but can lag by hours. Appropriate for testnet and low-notional deployment.

- (B) **ISDA credit events committee → Chainlink oracle update.** The ISDA determinations committee confirms a credit event (authoritative, immediate); a designated oracle operator updates the Chainlink feed. Introduces a centralised step but aligns with conventional market practice.
- (C) **Hybrid (Recommended).** The ISDA committee triggers a multi-sig vote among 5 authorised keepers; 3-of-5 consensus required to call `triggerCreditEvent`. Balances decentralisation with the authority of established credit conventions.
- (D) **zkProof of SWIFT MT103 payment failure (future).** A zero-knowledge proof of a missed coupon payment (SWIFT message failure) could trigger the credit event automatically and trustlessly. Technically feasible but requires SWIFT-to-blockchain data bridges not yet in production.

For the current Arbitrum Sepolia deployment, Option A is used. For commercial mainnet deployment, Option C is recommended.

5.5 Gas Cost Benchmarks

Table 7 reports the iCDS lifecycle gas *measured* by `forge test --gas-report`, priced at a live gas snapshot. The full lifecycle (one year, three quarterly premiums, credit-event settlement) consumes $\approx 897,000$ gas units, with the storage-heavy `openProtection` and `acceptProtection` dominating. At the snapshot this is $\approx \$0.036$ on Arbitrum versus $\approx \$0.28$ on Ethereum mainnet; the saving equals the gas-price ratio (uniform across functions, time-varying with gas prices). This keeps micro-notional iCDS contracts (e.g. \$1,000 notional for smallholder sukuk) economically viable: the gas cost is a small fraction of notional even at the smallest scale.

Table 7: Gas Cost Benchmarks: iCDS Lifecycle on Arbitrum One vs. Ethereum Mainnet.

Function	Gas units	Arb. (\$)	ETH main. (\$)	Saving
<code>openProtection</code>	263,800	\$0.0107	\$0.08	87%
<code>acceptProtection</code>	182,929	\$0.0074	\$0.06	87%
<code>payPremium</code>	104,823	\$0.0043	\$0.03	87%
<code>triggerCreditEvent</code>	68,487	\$0.0028	\$0.02	87%
<code>settle</code>	67,416	\$0.0027	\$0.02	87%
<code>expire</code>	54,031	\$0.0022	\$0.02	87%
Full lifecycle [†]	897,101	\$0.0364	\$0.28	87%

Gas units are Foundry `forge --gas-report` medians from `test/unit/iCDS.t.sol`. Costs at the 2026-05-30 snapshot: Arbitrum One 0.020 gwei, Ethereum mainnet 0.154 gwei, ETH/USD = \$2,030. [†]Full lifecycle = open + accept + 3×payPremium + triggerCreditEvent + settle.




Figure 6: Per-function gas: real measured values (`forge --gas-report`), Arbitrum One versus Ethereum mainnet. The L2 dollar saving over mainnet equals the live gas-price ratio (uniform across functions, time-varying): $\approx 87\%$ at the 2026-05-30 snapshot (Arbitrum 0.020 / mainnet 0.154 gwei), exceeding 99% at a typical 20 gwei mainnet, keeping the frequent premium payments at sub-cent cost.

6 Legal Analysis

6.1 Assessment against AAOIFI Shariah Standards

The Accounting and Auditing Organisation for Islamic Financial Institutions has issued no standard that explicitly addresses credit default swaps. The most closely applicable of its Shariah Standards [12] are Standard No. 20 (*Sale of Commodities in Organized Markets* — the standard governing futures- and options-like structures and the canonical source of the derivatives objections), Standard No. 5 (*Guarantees*, whose *kafala* analysis is the nearest analogue of the protection leg), Standard No. 49 (*Promises and Bilateral Promises*), and Standard No. 26 (*Islamic Insurance / takaful* [11]). We assess the iCDS against the relevant criteria, deferring the final reading of each standard to a qualified Shariah supervisory board.

(a) Wa’d structure. The iCDS uses a wa’d-based architecture: the seller makes a unilateral promise (*wa’d*) to pay the LGD upon a credit event; the buyer makes a separate unilateral promise to pay periodic premiums. Under Standard No. 49, a unilateral promise can be binding, but *muwa’adah* (mutually binding bilateral promises) is restricted where the paired promises amount to the underlying contract itself [12]. The iCDS open/accept architecture is structured with that constraint in view: no bilateral obligation exists at `openProtection` (the open offer is revocable by the seller until acceptance), and the binding

relationship arises only at `acceptProtection`, the operative moment of contract formation — whether this structure satisfies Standard No. 49 in a given school is a determination for the supervisory board.

(b) Premium as halal consideration. The premium is computed as the put-option-priced expected loss per period, not as a pre-specified annuity. This distinction is crucial: a fixed annuity on notional resembles a debt-service obligation with *riba* characteristics, whereas a dynamically re-priced risk premium is the actuarially fair *tabarru* (charitable contribution for mutual insurance purposes). The AAOIFI *takaful* model [11] (Standard No. 26) explicitly permits dynamic *tabarru* contributions adjusted to current risk.

(c) Real risk, not simulated gambling. The reference entity’s default (sovereign debt restructuring) is a real, verifiable economic event. The seller must post full notional as collateral, so no uncovered protection can be written; the buyer-side insurable-interest requirement (condition (2) below) excludes naked protection buying. Together these structurally address the *maysir* objection — neither party is positioned for a purely speculative wager — with the final determination resting with the supervisory board.

(d) Gharar minimisation. The credit event is defined as *oracle price* $\leq$ *recovery floor*: an objective, continuously observable, publicly verifiable condition. This replaces the ambiguous ISDA credit events committee with a mathematical threshold, directly addressing the *gharar* objection identified in Jobst & Solé [21].

Conclusion. The iCDS structure addresses each criterion examined above, subject to two conditions: (1) the seller posts full notional as collateral (anti-naked-CDS rule, enforced by the contract); and (2) the buyer holds the reference *sukuk* off-chain and provides a Shariah attestation to this effect at the time of `acceptProtection`. Formal determination of conformity with AAOIFI standards rests with a qualified Shariah supervisory board; no fatwa is claimed here.

These two conditions, together with the objective oracle settlement and fixed tenor, instantiate the four product-layer rules of the companion Shariah-boundaries study [10]: asset-backing (a real reference obligation), possession-gating (the buyer’s attested holding), insurable interest / no naked positions (both conditions above), and expiry with objective settlement (or *tabarru* risk-sharing). In the launch taxonomy of that study, the iCDS is risk transfer on a real, owned *sukuk* exposure — a hedging utility in the *takaful* spirit, not a cash-settled wager — and the same $\iota = 0$ pricing applies unchanged to protection on tokenised

sukuk holdings, the natural first deployment alongside sukuk-linked instruments.

6.2 ISDA/IIFM Tahawwut Documentation

The ISDA/IIFM Tahawwut Master Agreement [20] was designed for OTC Islamic hedging transactions and can be adapted to cover iCDS. Three modifications distinguish the iCDS from a standard CDS confirmation under the Tahawwut framework:

1. **Premium formula.** Replace the market-quoted spread with $s^* = \hat{\kappa}(1 - \delta)$ computed from the oracle price and the AR(1) OLS estimated convergence intensity.
2. **Credit event definition.** Replace the ISDA credit events list (Failure to Pay, Restructuring, Repudiation) with the on-chain oracle condition: `spot ≤ recoveryFloor`. A supplemental schedule should specify the oracle contract address and the reference asset.
3. **Settlement.** Replace the ISDA auction settlement with the LGD formula: payout = notional $\times (1 - \delta)$, settled immediately on-chain upon keeper confirmation.

Appendix C provides a draft term sheet in Tahawwut format.

6.3 Jurisdictional Analysis

Table 8 summarises the regulatory status across six jurisdictions.

Table 8: Jurisdictional Analysis: On-Chain iCDS Regulatory Status.

Jurisdiction	Primary Framework	Status	Key Requirement
UAE / VARA	VARA Virtual Asset Reg. 2023	Permissible	Virtual asset derivative; seller must be licensed
Malaysia / BNM	IFSA 2013 + SC Guidelines	Permissible	AAOIFI-compliant structure acceptable; EMI
Singapore / MAS	CMS Act + SF Act	<i>Conditional</i>	OTC derivative: requires CMS licence; EMIR
UK / FCA	UK MAR + CMAR	<i>Conditional</i>	Treated as OTC swap; EMIR UK reporting
US / CFTC	CEA + Dodd-Frank	Uncertain	Likely classified as swap; SEF registration
Bahrain / CBB	CBB Rulebook Vol. 2	Permissible	Gulf-friendly pilot jurisdiction; CBB Module

Sources: VARA Virtual Asset and Related Activities Regulation (2023); BNM Islamic Financial Services Act 2013; MAS Capital Markets Services Act; FCA UK CMAR; CFTC Commodity Exchange Act §2(h); CBB Rulebook Vol. 2 (Capital Markets, Module CM). Status reflects the authors’ legal assessment; independent regulatory counsel advised before commercial launch.

UAE / VARA. The Virtual Assets and Related Activities Regulation [22] (effective June 2023) establishes a comprehensive licensing regime for virtual asset service providers (VASPs) in Dubai. VARA’s virtual-asset derivatives regime — formalised in its Exchange Services Rulebook — permits exchange-traded derivatives (contracts for difference, futures, and options referenced to virtual assets) for appropriately licensed virtual asset service providers (VASPs), with the relevant authorisation falling under VARA’s activity-based licences (notably Exchange Services and/or Broker-Dealer Services) and subject to retail leverage caps and suitability requirements. A smart-contract-based credit derivative settled in USDC, such as the iCDS, falls within the scope of that derivatives framework. VARA’s KYC/AML and market-conduct obligations apply. The DIFC (Dubai International Financial Centre) *may* additionally require authorisation under the DFSA Rulebook if the contract is marketed to DIFC-based clients, but the Arbitrum blockchain deployment is outside DIFC territorial jurisdiction for domestic users.

Malaysia / BNM. The Islamic Financial Services Act 2013 (IFSA) and Securities Commission Guidelines on Digital Assets (2020) provide the applicable framework. Section 10 of IFSA permits Islamic derivative transactions that are AAOIFI-compliant. BNM would require a Shariah advisory board endorsement of the iCDS structure prior to commercial deployment. The SC’s digital asset guidelines classify smart-contract-settled instruments as capital-markets products requiring a Capital Markets Services licence.

Singapore / MAS. The Securities and Futures Act 2001 (SFA) governs OTC derivatives. A credit default swap – even if Islamic – is classified as an OTC derivative, and dealing in it requires a Capital Markets Services (CMS) licence issued under the SFA. OTC derivative trade reporting to a licensed trade repository (DTCC Data Repository (Singapore)) applies. MAS has not issued specific Islamic finance guidance for DeFi instruments; the iCDS would be assessed on its economic substance as a credit derivative.

UK / FCA. The Financial Services and Markets Act 2000 (FSMA) and the UK’s post-Brexit retained derivatives framework (UK EMIR, UK MiFIR) regulate OTC credit derivatives. The iCDS would require FCA authorisation under FSMA Section 19 for any dealing activity and trade reporting to an FCA-approved Trade Repository under UK EMIR Article 9. The FCA has actively supported Islamic finance development: its Discussion Paper DP22/2 (2022) explored regulatory accommodations for Shariah-compliant instruments, and HM Treasury’s Islamic Finance Task Force positioned London as a global Islamic finance hub. An iCDS would be assessed on its economic substance as a credit derivative, with no

current Islamic finance carve-out from standard OTC derivative rules. London’s concentration of AAOIFI-accredited advisers and sukuk-experienced legal counsel makes the UK a viable secondary pilot jurisdiction once commercial mainnet deployment proceeds.

US / CFTC. The Commodity Exchange Act Section 2(h) and Dodd-Frank Title VII classify most single-name credit default swaps as swaps subject to CFTC jurisdiction. A Shariah-compliant structure does not alter this classification. To operate commercially, the iCDS would require either registration on a Swap Execution Facility (SEF) or bilateral documentation under an ISDA/IIFM Master Agreement with eligible contract participants (ECPs). The CFTC has shown interest in blockchain-based derivatives (RFI on DeFi, June 2023) but has not provided specific Islamic finance guidance.

Bahrain / CBB. The Central Bank of Bahrain Rulebook Volume 2 (Capital Markets, Module CM) permits credit derivatives for licensed entities. Bahrain’s long-standing tradition as an Islamic finance centre and the CBB’s Shariah governance framework (Module CM-4) create a permissive environment. We identify Bahrain as a viable alternative pilot jurisdiction to UAE if VARA licensing timelines extend beyond the project’s commercial roadmap.

6.4 Recommended Pilot: UAE / VARA

We recommend UAE/VARA as the optimal first commercial jurisdiction on the basis of four factors. *Speed:* VARA operates a published, formal licensing process for virtual-asset derivative activity, providing a predictable authorisation pathway. *Geographic alignment:* the primary reference entities (UAE, Saudi Arabia, Qatar sovereigns) are GCC-domiciled, minimising cross-border regulatory friction. *Legal clarity:* VARA explicitly regulates virtual asset derivatives, removing uncertainty about whether the smart-contract settlement mechanism constitutes a licensing trigger. *Shariah ecosystem:* the UAE has an established network of AAOIFI-accredited Shariah scholars available for the mandatory advisory board function.

The regulatory pathway for commercial launch is therefore:

1. Obtain the applicable VARA VASP licence authorising virtual-asset derivative activity.
2. Engage AAOIFI-accredited Shariah board; obtain written fatwa endorsing the iCDS structure as described herein.
3. Upgrade the credit event oracle from Option A to Option C (3-of-5 keeper multi-sig referencing ISDA committee determinations).
4. Deploy on Arbitrum One mainnet (not Sepolia testnet).

5. Launch with initial reference entities: UAE sovereign, Saudi Arabia sovereign, Qatar sovereign.

7 Liquidity Bootstrapping

7.1 Market Structure

Two market structures are viable for the iCDS market. For small notionals (\$1k–\$500k), a single-sided automated market maker (AMM) is optimal: a protection seller deposits notional into an AMM pool; buyers can atomically open protection at the oracle-quoted put price. For large notionals (\$500k+), a central-limit order book (CLOB) allows price discovery across heterogeneous recovery rate assumptions. The current deployment supports peer-to-peer bilateral matching; an AMM wrapper is planned for a future version.

7.2 Governance Token Incentive Mechanism

A fixed-supply protocol governance token (GT) provides liquidity incentives aligned with the long-run sustainability of the iCDS market.³

- **Protection sellers (risk-takers):** earn GT rewards at rate λ_{GT} proportional to notional deposited and time open: $R_{\text{seller}} = \lambda_{GT} \cdot N \cdot \Delta t$.
- **Protection buyers (hedgers):** receive a GT fee discount on premiums paid, proportional to GT holdings.
- **Keepers:** earn GT per successful `triggerCreditEvent` call with valid oracle proof.

The GT emission schedule tapers over four years (60% / 25% / 10% / 5% of the liquidity allocation), after which fee revenue ($s^* \times$ total notional outstanding) sustains the market without token subsidy.

7.3 Participation Condition and Market Viability

Proposition 7.1 (Seller Participation Condition). *Let λ_{GT} denote the governance-token reward rate to sellers (GT per dollar-year of notional posted) and p_{GT} the prevailing GT*

³In the Baraka Protocol reference implementation, GT = BRKX (100M fixed supply; ERC-20+Votes, deployed on Arbitrum Sepolia). Specific parameters: 10M GT allocated to liquidity incentives over four years; fee tiers: 20% discount for $\geq 1,000$ GT held, 30% for $\geq 50,000$ GT. Source: <https://github.com/Arcus-Quant-Fund/BarakaDapp>.

market price. A capital allocator with outside return r_f weakly prefers participating as an iCDS protection seller if and only if:

$$\lambda_{GT} \cdot p_{GT} \geq r_f. \quad (17)$$

When equation (17) holds, participation is a best response for every seller regardless of others' choices; hence a symmetric equilibrium with positive aggregate liquidity exists.

Proof sketch. The seller's expected profit per unit notional per period is the premium received, minus the expected default loss, plus the governance-token reward:

$$\begin{aligned} \pi_s &= s^* \cdot \Pr[\text{survive}] - \text{LGD} \cdot \Pr[\text{default}] + \lambda_{GT} p_{GT} \\ &= \kappa(1 - \delta)(1 - p_D) - (1 - \delta)p_D + \lambda_{GT} p_{GT}, \end{aligned}$$

with $p_D \approx \kappa \Delta t$ for a period of length Δt . At the fair spread $s^* = \kappa(1 - \delta)$ the insurance leg is zero-NPV: per unit time the premium rate $\kappa(1 - \delta)$ exactly offsets the expected-loss rate (intensity κ times LGD $1 - \delta$), so $\pi_s = \lambda_{GT} p_{GT} + O(\kappa^2)$. The posted collateral earns no interest, so the seller's only excess return is the governance-token yield. Setting $\pi_s \geq r_f$ therefore yields the participation condition (17): the token reward alone must cover the outside return, since the fair insurance spread contributes nothing on a risk-adjusted basis.

□

□

As a calibration check, at 2025 SOFR ($r_f = 4.24\%$), equation (17) requires the governance-token yield to sellers to cover the full outside return, $\lambda_{GT} p_{GT} \geq 4.24\%$ per annum (the fair insurance spread $\kappa(1 - \delta)$ is zero-NPV and does not contribute). With a Year-1 GT emission of 6% of total supply directed to sellers and a target TVL of \$1M, the required GT-price threshold is approximately $p^* \approx \$0.007$ – a low bar relative to comparable DeFi governance token valuations, implying that even modest market adoption sustains the Nash equilibrium condition.

7.4 Liquidity Milestones

1. **Minimum viable (\$1M notional):** price discovery becomes possible; bid-offer spreads on iCDS terms compressible below 10 bps.
2. **Useful scale (\$50M notional):** comparable to a small conventional CDS market; GT rewards self-sustaining from fee revenue.

3. **Target (\$500M by Year 3):** comparable to a medium CDS market on individual reference entities; attractive to institutional Islamic finance participants seeking hedging against GCC sovereign sukuk exposure.

8 Systemic Risk

8.1 Lessons from 2008

Beyond bootstrapping liquidity, an Islamic CDS market must avoid the systemic failures that discredited the conventional one. The 2008 financial crisis demonstrated three systemic failure modes of conventional CDS markets. First, *uncollateralised selling*: AIG and others wrote billions in CDS notional without posting collateral against the contingent liability, creating uncovered exposures that failed exactly when triggered [15]. Second, *underestimated correlation*: Gaussian copula models underestimated joint default probabilities, leading to catastrophic concentration of credit risk in protection sellers. Third, *opacity*: the OTC bilateral market had no central counterparty; regulators could not monitor aggregate exposure.

8.2 iCDS Safeguards

The iCDS smart contract addresses each failure mode by design.

Full collateralisation. Every seller must deposit the full notional in the contract before a protection can be opened (line 209 of `iCDS.sol`, which calls `safeTransferFrom` to escrow the full notional). Naked positions are mechanically impossible.

Transparent exposure. All open protections are visible on-chain: any observer can query `protections[id]` and `nextId` to enumerate the full aggregate exposure. Regulators can monitor positions in real time without relying on mandatory reporting regimes.

No synthetic leverage chains. Because each seller funds their own position individually, there are no synthetic CDS-on-CDS structures. The aggregate exposure of the iCDS market is bounded by the sum of collateral actually deposited, not by a notional multiple of underlying assets.

8.3 Circuit Breakers and Position Limits

The `Pausable` mechanism allows the contract owner to halt all new interactions in the event of an oracle manipulation attempt or black-swan default event. Future versions will implement:

- Per-seller notional caps (e.g., max 20% of TVL from a single seller).
- Per-reference-entity notional caps (diversification across sovereigns).
- Automatic pause if oracle price moves $> 30\%$ in 24 hours (flash-crash protection).

These circuit breakers directly address the correlation and concentration risks that amplified the 2008 crisis.

9 Conclusion

We have presented the complete scientific, engineering, and regulatory case for the first Islamic Credit Default Swap on public blockchain infrastructure.

Theoretically, the iCDS spread $s^* = \kappa(1 - \delta)$ is the actuarially fair credit protection premium, and — because both legs terminate at default and the discount cancels — the par spread is independent of the rate ι at every maturity, so a correctly-priced conventional CDS quotes the same fair spread. The par spread is therefore automatically *riba*-neutral; the *riba* is confined to the present value, where we define $\Pi_{\text{riba}}^{\text{PV}}(T)$, the discounting haircut on the future loss — increasing in ι , vanishing at $\iota = 0$, and a small fraction of notional at realistic maturities. The iCDS removes it both by pricing at $\iota = 0$ and structurally: collateral held in a non-interest-bearing *tabarru* pool, full collateralisation, and an objective on-chain credit event. Empirically, the analysis runs on a cross-section of 62 real, actively-quoted sovereign USD sukuk (Cbonds, 12 issuers), with no synthetic panel: κ read from those sukuk G -spreads recovers the hazard independent instruments imply — agreeing within 6–16 bp with same-issuer conventional bonds and co-moving with sovereign CDS at correlation 0.97. The discount-independence of the par spread also explains why sovereign CDS stayed at 48–100 bps in the near-zero-rate era 2012–2018 while policy rates sat near zero.

Technically, `iCDS.sol` implements the full lifecycle (open, accept, quarterly dynamic premiums, credit event trigger, LGD settlement) on Arbitrum, removing *riba* and structurally addressing *maysir* and *gharar* through smart-contract mechanics rather than legal workarounds. The full lifecycle is $\approx 897,000$ gas (Foundry-measured) — a few US cents on Arbitrum, with an L2-vs-L1 saving equal to the live gas-price ratio (order 90–99%), enabling economically viable micro-notional Islamic credit hedging.

Legally, UAE/VARA is identified as the most permissive immediate jurisdiction. The five-step commercial pathway (VARA licence, Shariah board fatwa, Oracle upgrade, mainnet deployment, GCC sovereign reference entities) is executable within 12–18 months of this paper’s publication.

Taken together, these contributions provide the academic foundation, deployed engineering artefact, and regulatory roadmap for an Islamic CDS market whose *riba* content is settled at the pricing layer, whose product-layer safeguards are designed for Shariah-board review, and which is technically operational (testnet) and legally mappable. The iCDS described in this paper is not purely a theoretical construct: it is running as a pre-audit prototype on Arbitrum Sepolia today; a commercial launch additionally requires an independent security audit (including the premium-flow correction of Section 5), the oracle upgrade to Option C, Shariah-board review, regulatory authorisation, and the liquidity bootstrapping mechanisms described here.

References

- [1] al-Dharir, S. M. al-Amin (1997). *Al-Gharar in Contracts and Its Effect on Contemporary Transactions*. Islamic Research and Training Institute (IRTI), Eminent Scholars’ Lecture Series 16.
- [2] Bacha, O. I. (1999). Derivative instruments and Islamic finance: Some thoughts for a reconsideration. *International Journal of Islamic Financial Services* 1(1).
- [3] Dusuki, A. W., & Abozaid, A. (2007). A critical appraisal on the challenges of realizing Maqasid al-Shariah in Islamic banking and finance. *IIUM Journal of Economics and Management* 15(2), 143–165.
- [4] El-Gamal, M. A. (2006). *Islamic Finance: Law, Economics, and Practice*. Cambridge University Press.
- [5] Obaidullah, M. (2001). Financial engineering with Islamic options. *Islamic Economic Studies* 8(1), 73–103.
- [6] Usmani, M. T. (2002). *An Introduction to Islamic Finance*. Kluwer Law International.
- [7] Akerer, D., Hugonnier, J., & Jermann, U. (2025). Perpetual futures pricing. *Mathematical Finance*, 1–17. DOI:10.1111/mafi.70018. (NBER Working Paper 32936, 2024; SSRN 4603820, 2023.)

- [8] Ahmed, S., & Bhuyan, R. (2026a). The κ -yield curve: Construction, estimation methodology, and real-data evidence from sovereign sukuk. *SSRN Working Paper 6322938*. https://papers.ssrn.com/sol3/papers.cfm?abstract_id=6322938
- [9] Ahmed, S., & Bhuyan, R. (2026b). Random stopping time as credit event: A riba-free credit pricing framework. *SSRN Working Paper 6322858*. https://papers.ssrn.com/sol3/papers.cfm?abstract_id=6322858
- [10] Ahmed, S., & Bhuyan, R. (2026c). The jurisprudential boundaries of an interest-free substrate: Riba, gharar, maysir, and qabd in cash-settled synthetic instruments. Working paper (companion research note).
- [11] AAOIFI. (2010). *Shariah Standard No. 26: Islamic Insurance (Takaful)*. Accounting and Auditing Organisation for Islamic Financial Institutions, Bahrain.
- [12] AAOIFI. (2017). *Shari'ah Standards (Full Text)* — including Standard No. 5 (Guarantees), No. 20 (Sale of Commodities in Organized Markets), and No. 49 (Promises and Bilateral Promises). Accounting and Auditing Organisation for Islamic Financial Institutions, Bahrain.
- [13] Bank for International Settlements. (2024). *OTC Derivatives Statistics at End-June 2024*. BIS Statistical Release, November 2024. URL: https://www.bis.org/publ/otc_hy2411.htm
- [14] Cakir, S., & Raei, F. (2007). Sukuk vs. Eurobonds: Is there a difference in value-at-risk? *IMF Working Paper WP/07/237*.
- [15] Cont, R. (2010). Credit default swaps and financial stability. *Banque de France Financial Stability Review*, 14, 35–43.
- [16] Cox, J. C., Ingersoll, J. E., & Ross, S. A. (1985). A theory of the term structure of interest rates. *Econometrica*, 53(2), 385–408.
- [17] Damodaran, A. (2026). Country risk: Determinants, measures and implications – sovereign CDS spreads dataset. *Stern School of Business, New York University*. URL: <https://pages.stern.nyu.edu/~adamodar/> (accessed 1 March 2026).
- [18] Duffie, D., & Singleton, K. J. (1999). Modeling term structures of defaultable bonds. *Review of Financial Studies*, 12(4), 687–720.
- [19] Federal Reserve Bank of St. Louis. (2026). *Secured Overnight Financing Rate (SOFR) and Federal Funds Rate*. FRED Economic Data. URL: <https://fred.stlouisfed.org>

- [20] ISDA/IIFM. (2010). *Tahawwut (Hedging) Master Agreement*. International Swaps and Derivatives Association / International Islamic Financial Market.
- [21] Jobst, A., & Solé, J. (2012). Operative principles of Islamic derivatives: Towards a coherent theory. *IMF Working Paper WP/12/63*.
- [22] Virtual Assets Regulatory Authority. (2023). *Virtual Assets and Related Activities Regulation 2023*. Government of Dubai, UAE. URL: <https://www.vara.ae>
- [23] van Wijnbergen, S., & Zaheer, S. (2013). Sukuk defaults: On distress resolution in Islamic finance. *Tinbergen Institute / Duisenberg School of Finance Discussion Paper TI 13-087 / DSF 57*.

A iCDS.sol Key Functions (Solidity)

The following listings reproduce the core functions of `iCDS.sol` (deployed at `0x77a98785...521b` on Arbitrum Sepolia). *Full source*: <https://github.com/Arcus-Quant-Fund/BarakaDapp>.

Everlasting put formula at $\iota = 0$. The `EverlastingOption.quotePut()` function implements Proposition 6 of Ackerer, Hugonnier & Jermann [7] at $\iota = 0$, in the two-region closed form derived in the companion paper [9]. For an everlasting put with current spot S and strike K (recovery floor),

$$\Pi_{\text{put}}(S, K; \kappa) = \begin{cases} C_p S^{\beta_p}, & S \geq K, \\ (K - S) + \frac{K^{1-\beta_c}}{\beta_c - \beta_p} S^{\beta_c}, & S < K, \end{cases} \quad C_p = \frac{K^{1-\beta_p}}{\beta_c - \beta_p}, \quad (18)$$

where

$$\beta_p = \frac{1}{2} - \sqrt{\frac{1}{4} + \frac{2\kappa}{\sigma^2}} < 0, \quad \beta_c = \frac{1}{2} + \sqrt{\frac{1}{4} + \frac{2\kappa}{\sigma^2}} > 1, \quad (19)$$

are the negative and positive roots of the characteristic equation $\frac{\sigma^2}{2}\beta(\beta-1) = \kappa$ at $\iota = 0$, and σ^2 is the variance parameter of the reference asset price process. The two-region form (value-matching and smooth-pasting at $S = K$, with the intrinsic value $K - S$ as the particular solution in the money) is essential: a single power law extended through $S < K$ would omit the intrinsic value and misprice the floor [9]. The deployed `EverlastingOption` contract computes exactly this two-region formula (`_exponents()` returns β_p and β_c ; `_quotePut()`

switches regions at the strike). The contract uses $S = \text{spotWad}$ and $K = \text{recoveryFloorWad}$ from the oracle, with σ^2 estimated from the historical volatility of the reference asset price.

```

1 function openProtection(
2     address refAsset, address token,
3     uint256 notional, uint256 recoveryRateWad, uint256 tenorDays
4 ) external nonReentrant whenNotPaused returns (uint256 id) {
5     uint256 spotWad = oracle.getIndexPrice(refAsset);
6     uint256 recoveryFloorWad = (spotWad * recoveryRateWad) / WAD;
7     id = _nextId++;
8     protections[id] = Protection({
9         seller: msg.sender, buyer: address(0),
10        refAsset: refAsset, token: token,
11        notional: notional, recoveryRateWad: recoveryRateWad,
12        recoveryFloorWad: recoveryFloorWad,
13        tenorEnd: block.timestamp + tenorDays * 1 days,
14        lastPremiumAt: 0, premiumsCollected: 0, status: Status.Open
15    });
16    IERC20(token).safeTransferFrom(msg.sender, address(this), notional);
17 }

```

Listing 2: openProtection: seller deposits notional as collateral

```

1 function _computePremium(Protection storage prot)
2     internal view returns (uint256)
3 {
4     uint256 spotWad = oracle.getIndexPrice(prot.refAsset);
5     uint256 putRateWad = evOption.quotePut(
6         prot.refAsset, spotWad, prot.recoveryFloorWad
7     );
8     return (putRateWad * prot.notional) / WAD;
9 }

```

Listing 3: _computePremium: put-option-priced premium at $\text{iota}=0$

```

1 function triggerCreditEvent(uint256 id) external whenNotPaused {
2     require(authorisedKeepers[msg.sender], "iCDS: not keeper");
3     Protection storage prot = protections[id];
4     uint256 spotWad = oracle.getIndexPrice(prot.refAsset);
5     require(spotWad <= prot.recoveryFloorWad, "iCDS: no default");
6     prot.status = Status.Triggered;
7 }

```

Listing 4: triggerCreditEvent: on-chain oracle credit event

B Real Constant-Maturity $\hat{\kappa}$ Dynamics

Table 9: Real Constant-Maturity $\hat{\kappa}$ Dynamics (Cbonds daily G -spread histories).

Issuer	Tenor	$\hat{\kappa}$ (%)	$\hat{\alpha}$ (p.a.)	Half-life (mo., bc)	Feller
Malaysia	7y	0.89	5.83	1.4	Yes
Saudi Arabia	7y	2.31	2.53	3.3	Yes
Indonesia	7y	2.80	1.95	4.3	Yes
Bahrain	7y	5.22	3.11	2.7	Yes
Turkey	5y	7.44	0.79	10.5	Yes

Source: Cbonds daily constant-maturity G -spread histories (roll-down removed), 2018–2026, via the companion κ -yield-curve study [8]. $\hat{\alpha}$ and half-life are Kendall bias-corrected (bc); Feller condition $2\hat{\alpha}\hat{\kappa} \geq \hat{\nu}^2$ satisfied in every series. All values are estimated from real traded sukuk, not a synthetic panel.

C Draft iCDS Term Sheet (ISDA/IIFM Tahawwut Format)

Field	Value / Description
Transaction Type	Islamic Credit Default Swap (iCDS) under ISDA/IIFM Tahawwut Master Agreement [20]
Reference Entity	[e.g., Kingdom of Saudi Arabia]
Reference Obligation	[e.g., Saudi Arabia 5yr USD Sukuk, ISIN: XS0000000000]
Protection Seller	[VARA-licensed entity or bilateral ECP]
Protection Buyer	[Must hold Reference Obligation off-chain; Shariah attestation required]
Notional Amount	[e.g., USD 1,000,000]
Premium Formula	$s^* \times \text{Notional} \times \Delta t$ per payment date, with $s^* = \hat{\kappa}(1 - \delta)$ and $\Delta t = 0.25$ (quarterly), $\hat{\kappa}$ computed from the oracle at each payment date (dynamic, not fixed)
Premium Payment Dates	Quarterly (every 90 days from Effective Date)
Convergence Intensity ($\hat{\kappa}$)	AR(1) OLS estimate from Table 9; or live oracle feed (future)
Recovery Rate (δ)	0.60 (60%), per Moody’s sovereign tables

Field	Value / Description
Credit Event Definition	On-chain: $\text{oracle spot price} \leq \text{Recovery Floor} = \delta \times \text{Spot at Effective Date}$
Credit Event Oracle	<code>OracleAdapter</code> at <code>0x4f6C...2999</code> ; confirmed by 3-of-5 keeper multi-sig (recommended Option C)
Settlement	Cash settlement (physical-equivalent economics): Protection Buyer receives $\text{Notional} \times (1 - \delta)$ in USDC on-chain within 1 block of Buyer calling <code>settle(id)</code>
Tenor	[e.g., 5 years from Effective Date]
Governing Law	[e.g., DIFC law / English law]
Shariah Adviser	[AAOIFI-accredited scholar; fatwa reference number]
Smart Contract	<code>iCDS.sol</code> at <code>0x77a98785...521b</code> (Arbitrum Sepolia test-net; mainnet address TBD)
VARA Licence	[VARA VASP derivatives licence no.; TBD]

CRedit Author Contribution Statement. S.A.: conceptualisation, software, methodology, writing (original draft); R.B.: supervision, writing (review), resources.

Conflict of Interest Statement. S.A. is the principal architect of the smart contracts described in this paper and holds governance tokens in the project. R.B. declares no conflicts.

Data Availability. All data, code, and pipeline scripts are available at <https://github.com/Arcus-Quant-Fund/BarakaDapp>.

EDITORIAL NOTE TO THIS PRINTING

Paper 3 · The Simulation Framework

Working paper, revised July 2026 SSRN 6323618

Role in the programme. The adversarial test-bench: what happens when the key number is a measurement and someone wants to bend it. **What it establishes.** An integrated cadCAD, reinforcement-learning and game-theoretic framework; a closed-form manipulation frontier, with the cost-to-move multiplier running from tens to thousands depending on book depth. **Corrected or extended later.** The abstract's net-transfer claim was restated with confidence-interval honesty in July 2026 (a small positive mean traced to discretisation drift, not a free lunch). **Not established here.** Live-market validation of the simulated frontier; the chain's defences implement what the simulations recommend.

Replication for every number in this paper: DOI 10.5281/zenodo.21290557 · github.com/shehzadahmed-xx/kappa-replication

An Integrated Simulation Framework for Decentralised Finance Protocols: cadCAD, Reinforcement Learning, Game Theory, and Mechanism Design

Shehzad Ahmed

Arcus Quant Fund
<https://arcusquantfund.com>

March 2026

Abstract

We present the Integrated Economic System (IES), the first four-layer simulation framework for decentralised finance (DeFi) protocols, combining cadCAD state-machine modelling, reinforcement learning (RL) traders, game-theoretic Nash equilibrium analysis, and mechanism design optimisation into a single closed-loop system. Each layer captures a distinct economic phenomenon that no single method can address alone: cadCAD provides the market-level ground truth; RL agents model rational individual traders with optimal stopping intuition; game theory characterises strategic leverage selection; and mechanism design searches the protocol parameter space for the Pareto-optimal regime that simultaneously minimises insolvency risk, liquidation frequency, and funding-rate deviation from zero.

The framework applies to any DeFi protocol whose smart contract logic can be expressed as a cadCAD partial-state-update block. As a case study, we apply the IES to a riba-free-by-design ($\iota = 0$) perpetual futures exchange (the Baraka Protocol, deployed on the Arbitrum Sepolia testnet, pre-audit), producing four primary simulation results over 5 episodes $\times$ 720 steps (150 simulated trading days): (i) zero insolvency events across all Monte Carlo paths; (ii) *mean* Nash-equilibrium leverage of $2.72\times$ (long) and $3.28\times$ (short) across intervals (headline seed; $3.16\times/2.84\times$ over $R = 30$ seeds, Robustness 1), well below the $4.97\times$ longs are driven to in the $\iota > 0$ counterfactual; (iii) a small positive mean net transfer (\$865 per \$100,000 notional, i.e. 0.87%; 95% CI \$245–\$1,485, over $R = 30$ seeds — statistically distinguishable from zero at finite horizon, attributable to the discretisation drift analysed in §8, and an order of magnitude below the interest-driven transfer; the individual intervals are large and opposite-signed) versus a systematic \$8,420 paid by longs under $\iota > 0$; and (iv) endogenous mechanism design convergence to a tighter circuit breaker ($75 \rightarrow 41$ bps), higher maintenance margin ($2.0 \rightarrow 3.1\%$), and greater insurance split ($50 \rightarrow 61\%$), consistent with the κ -rate parameterisation implied by Ahmed and Bhuyan [5]. We extend these with three robustness analyses — Monte Carlo confidence intervals over $R = 30$ seeds, a Sobol sensitivity decomposition of the mechanism-design objective, and a heterogeneous five-type trader population — and report an honest negative result on reinforcement learning (a trained PPO agent underperforms the rule-based baseline). We release the full simulation codebase as open-source.

Keywords: cadCAD, reinforcement learning, game theory, mechanism design, Islamic finance, DeFi, agent-based simulation, perpetual futures, Shariah compliance, $\iota = 0$, computational economics

1 Introduction

The design and analysis of decentralised finance (DeFi) protocols sits at the intersection of mechanism design, game theory, and stochastic control, yet is rarely treated as such. Most DeFi research either (a) studies individual agent behaviour in isolation [9, 20], (b) applies equilibrium theory to single-parameter models [16], or (c) runs agent-based Monte Carlo simulations without equilibrium analysis [6]. No existing work couples all four methods into a single feedback system whose state-update blocks

mirror the logic of a deployed (testnet) protocol implementation.

The problem is particularly acute in Islamic DeFi, where protocol correctness carries a dual obligation: economic soundness and jurisprudential compliance. A protocol that is economically stable but fails the $\iota = 0$ riba test is invalid; a protocol that is Shariah-compliant but mechanistically insolvent is undeployable. Verifying both simultaneously requires a simulation infrastructure that is richer than any single method provides.

This paper presents such an infrastructure: the **Inte-**

grated Economic System (IES) simulation framework. The IES couples:

1. A **cadCAD state machine** that implements the exact on-chain protocol logic, serving as the simulation ground truth.
2. A **reinforcement learning (RL) agent** embedded inside the cadCAD market, whose actions are drawn from an observation of live protocol state and fed back as on-chain positions.
3. A **game theory layer** that, every N steps, extracts the recent price window, builds a 5×5 payoff matrix over the leverage space, and computes the Nash equilibrium leverage distribution — which is then returned to the cadCAD simulation as the background-trader respawn policy.
4. A **mechanism design optimiser** that, at each episode boundary, receives aggregate cadCAD metrics, minimises a multi-objective loss over the protocol parameter space, and blends the result into the next episode’s parameters.

The key conceptual contribution is the *coupling architecture*: a directed graph of information flows connecting the four layers so that each layer both consumes outputs from the others and produces inputs for them. This bidirectional coupling distinguishes the IES from sequential multi-method analysis: the layers are not independent. The Nash equilibrium changes the cadCAD trader policy, which changes the price path, which changes the payoff matrices that determine the next Nash equilibrium.

Application. We apply the IES to a *riba-free-by-design* ($\iota = 0$) perpetual futures exchange [4] deployed on the Arbitrum Sepolia testnet in February 2026 (pre-audit), built around the $\iota = 0$ no-arbitrage condition of Akerer, Hugonnier & Jermann [1] — simultaneously the source of its *riba*-freedom at the pricing layer and the theoretical object we seek to validate. (The $\iota = 0$ design addresses the *riba* prohibition at the pricing layer; *gharar*, *maysir*, and *qabd* are product-layer questions reserved for Shariah scholars.)

Structure. Section 2 reviews each simulation method and its limitations when used alone. Section 3 presents the IES coupling architecture in full. Sections 4–7 detail each layer. Section 8 presents the four primary experimental results, three robustness analyses, and an honest negative reinforcement-learning result. Section 10 discusses implications for DeFi mechanism design. Section 12 concludes.

2 Background and Related Work

2.1 cadCAD: Complex Adaptive System Modelling

cadCAD (complex adaptive dynamics/computer-aided design) is a Python framework for specifying economic state machines as a set of *partial-state-update blocks* [26, 8]. Each block consists of *policy functions* (which compute signals from current state) and *state update functions* (which apply signals to produce the next state). The framework natively supports Monte Carlo over independent random number streams and parameter sweeps.

cadCAD is well-suited to protocol analysis because it forces the modeller to specify every state variable and every transition function explicitly, producing a simulation that mirrors smart contract logic. However, cadCAD lacks:

- **Rational agent learning.** Background trader policies in cadCAD are typically hard-coded rules, not learned strategies.
- **Strategic interaction.** cadCAD does not model multi-agent games or equilibrium selection.
- **Endogenous parameter adaptation.** Protocol parameters are fixed within a simulation run.

2.2 Reinforcement Learning for DeFi

Reinforcement learning (RL) has been applied to DeFi trading [25], extractable-value and liquidation analysis [21], and AMM liquidity provision [13]. The key advantage is that RL agents *learn* optimal policies from experience rather than requiring closed-form strategy specifications.

Stable-Baselines3 [22] provides the Proximal Policy Optimisation (PPO) implementation used in this work. The RL environment (described in Section 5) is a gymnasium-compatible wrapper around the cadCAD market state.

The limitation of RL alone is that it models one rational agent against a fixed background. It does not characterise the multi-agent equilibrium that emerges when all traders adapt simultaneously.

2.3 Game Theory in DeFi

Nash equilibrium analysis has been applied to AMM fee structures [16], MEV extraction [9], and liquidation dynamics [19]. For perpetual futures, the relevant game is between long and short players choosing leverage, with the funding rate as the strategic variable.

The pricing foundation is the Akerer et al. [1] result that $\iota = 0$ achieves unique no-arbitrage perpetual-futures pricing (their Theorem 3), under which the funding rate carries no interest component. AHJ do not, however, frame $\iota = 0$ as a Nash equilibrium — their Theorem 3 is a no-arbitrage pricing result, not a game-theoretic one — so the equilibrium analysis here is our own. Our game theory layer studies the *leverage* game among traders op-

erating under the $\iota = 0$ funding rule, computing the Nash equilibrium numerically over live cadCAD price paths using the nashpy vertex enumeration algorithm [15].

2.4 Mechanism Design for DeFi Protocols

Mechanism design [14, 18] studies how to construct the rules of a game so that self-interested agents produce socially desirable outcomes. Applied to DeFi, the “mechanism” is the protocol parameter vector $\theta = (F_{\max}, m, p, s)$ (circuit breaker, maintenance margin, liquidation penalty, insurance split), and the “social objective” is the multi-objective loss $\mathcal{L}(\theta)$ of equation (14) in Section 7.

Prior work in DeFi mechanism design includes AMM curve optimisation [7], liquidation parameter tuning [12], and governance token design [11]. We extend this to the Islamic DeFi context, where the mechanism must satisfy a Shariah constraint (funding-rate symmetry, $\iota = 0$) in addition to economic objectives.

3 The Integrated Economic System Architecture

3.1 Design Principles

The IES is built on three design principles:

1. **One ground truth.** The cadCAD state machine is the single source of market prices, positions, funding rates, and insurance-fund balances. All other layers read from and write to this state. There is no separate “market” in the RL environment or game theory module; both operate on the live cadCAD state.
2. **Asynchronous coupling.** Layers run at different time scales. The RL agent acts at every cadCAD step (hourly). The game theory layer computes a new Nash equilibrium every $\Delta_{\text{GT}} = 50$ steps. The mechanism design optimiser runs once per episode (720 steps). This mirrors the natural time scales of individual traders, strategic leverage rebalancing, and protocol governance updates.
3. **Closed-loop feedback.** Every layer both consumes outputs from and produces inputs for the cadCAD engine. The coupling graph is:

$$\begin{aligned} \text{cadCAD} &\xrightarrow{s_t} \text{RL} \xrightarrow{a_t} \text{cadCAD} \\ \text{cadCAD} &\xrightarrow{\{p_s\}_{t-W}^t} \text{GT} \xrightarrow{(\ell_L^*, \ell_S^*)} \text{cadCAD} \\ \text{cadCAD} &\xrightarrow{\mu_T} \text{MD} \xrightarrow{\theta^*} \text{cadCAD} \end{aligned}$$

where s_t is state observation, a_t is action, $\{p_s\}$ is the price window of length W , (ℓ_L^*, ℓ_S^*) are Nash equilibrium leverages, μ_T are episode aggregate metrics, and θ^* is the updated parameter vector.

3.2 Episode Structure

Each simulation run consists of E episodes, each of T steps. Within an episode:

1. cadCAD initialises market state (price, positions, insurance fund) with current parameter vector θ_e .
2. At each step t : market step, RL action, GT recomputation (if $t \bmod \Delta_{\text{GT}} = 0$).
3. At episode end: MD optimiser produces θ_{e+1} .

The episode boundary is the Islamic governance analogue: a Shariah board review that can adjust protocol parameters based on observed performance, subject to the $\iota = 0$ and $\ell_{\max} = 5$ hard constraints.

4 Layer 1: cadCAD Market Engine

Implementation note. The Layer-1 market is specified as five cadCAD partial-state-update blocks (Section 4, Partial State Update Blocks). We run these blocks in two ways. The standalone Layer-1 Monte Carlo is executed through the actual cadCAD library (cadCAD 0.5.3) using its `Executor` with sequential substep semantics; a 20-run $\times$ 720-step ensemble yields zero insolvency in 20/20 runs and a funding rate confined to the symmetric ± 75 bps band ($\iota = 0$), consistent with the integrated results reported below. The *integrated* four-layer runs (which produce Table 1) instead use a pure-Python implementation of the *same* partial-state-update blocks, executed in a single process so that the reinforcement-learning, game-theory, and mechanism-design layers can be coupled tightly within each step. Both share the identical cadCAD state space and block structure.

4.1 State Space

The cadCAD state $\mathcal{S}$ consists of:

$$\mathcal{S} = \{p_t^m, p_t^i, F_t, \text{CF}_t, \mathcal{P}_t, I_t, L_t, v_t\} \quad (1)$$

where p_t^m and p_t^i are mark and index prices, F_t is the instantaneous funding rate, CF_t is the cumulative funding index, $\mathcal{P}_t$ is the pool of open positions, I_t is the insurance fund balance, L_t is the liquidation count, and $v_t \in \{0, 1\}$ is the solvency flag.

4.2 Partial State Update Blocks

The state transition is implemented as five cadCAD partial-state-update blocks:

Block 1 (Oracle). The index price follows geometric Brownian motion:

$$p_{t+1}^i = p_t^i \exp(\sigma_i \epsilon_t^i), \quad \epsilon_t^i \sim \mathcal{N}(0, 1) \quad (2)$$

Note this discretisation omits the Itô compensator and is therefore not a martingale: $\mathbb{E}[p_{t+1}^i | p_t^i] = p_t^i e^{\sigma_i^2/2}$, an upward index drift of ≈ 2 bps per hourly step at $\sigma_i = 0.02$.

This drift contributes, together with the mark-process premium drift, to the small positive finite-horizon transfer in Robustness 1, and inflates the absolute PnL of long-biased policies (including the random baseline of Section 11); the $\iota = 0$ versus $\iota > 0$ *contrast*, which is the object of interest, is unaffected since both arms share the same oracle process. The mark price mean-reverts to index:

$$p_{t+1}^m = 0.9 p_t^m \exp(\sigma_m \epsilon_t^m) + 0.1 p_{t+1}^i \quad (3)$$

with $\sigma_m = 0.7 \sigma_i$ (reduced volatility, market maker dampening).

Block 2 (Funding). The funding rate at $\iota = 0$:

$$F_t = \text{clip}\left(\frac{p_t^m - p_t^i}{p_t^i}, -F_{\max}, +F_{\max}\right) \quad (4)$$

$F_{\max}$ is the protocol circuit breaker (initially 75 bps, adaptive via MD layer). The $\text{clip}(\cdot)$ bound is exactly the funding-rate *clamp* whose no-arbitrage perpetual pricing under stochastic rates is formalised by Hugonnier & Jermann [2]; our $\iota = 0$ funding rule is the $I = 0$ case of that clamped specification.

Block 3 (Traders). Background positions are accrued: $c_{i,t+1} = c_{i,t} - \text{sign}(\ell_i) \Delta \text{CF}_t n_i$ where c_i is collateral, n_i is position size, and $\Delta \text{CF}_t = F_t$. The RL agent’s position is managed by Layer 2.

Block 4 (Liquidations). Position i is liquidated if $c_i < n_i \cdot m$ (maintenance margin fraction m):

$$\Delta I_t = \sum_{i \text{ liq.}} n_i p s \quad (5)$$

$$c_i \leftarrow 0, \quad n_i \leftarrow n_{\text{respawn}} \quad (6)$$

where p is the liquidation penalty and s is the insurance split.

Block 5 (Bookkeeping). Solvency flag: $v_t = \mathbf{1}[I_t < 0]$.

4.3 Respawn Policy with Nash Leverage

When a position is liquidated (Block 4), it is immediately respawned with a new collateral of \$5,000 and leverage drawn from the current Nash equilibrium distribution $\mathcal{D}_{\text{Nash}}(\ell_L^*, \ell_S^*)$ supplied by the game theory layer. This coupling ensures the background trader population adapts its strategic behaviour in response to observed market conditions.

5 Layer 2: Reinforcement Learning Trader

5.1 RL Environment

The RL environment (`IESTraderEnv`) is a `gymnasium.Env` wrapping the `cadCAD` state. The environment is injected into the `cadCAD` simulation loop so that the agent’s actions directly modify the `cadCAD` position pool at each step.

Observation space. $\mathcal{O} \subset \mathbb{R}^8$:

$$s_t = (p_t^m/p_0, p_t^i/p_0, F_t, \text{CF}_t, q_t/p_0, c_t/p_0, \mathbf{1}[\text{long}], w_t/W_0) \quad (7)$$

where q_t and c_t are position size and collateral, w_t is free collateral, $p_0 = \$60,000$, $W_0 = \$100,000$.

Action space. $\mathcal{A} = \{0, 1, \dots, 11\}$:

- $a = 0$: hold
- $a \in [1, 5]$: open long at leverage $a \times$
- $a \in [6, 10]$: open short at leverage $(a - 5) \times$
- $a = 11$: close current position

The leverage range $[1, 5]$ is constrained to the Shariah hard cap $\ell_{\max} = 5$ (ShariahGuard on-chain).

Reward function. The reward at each step is the sum of funding PnL (shaped reward for survival) and realised price PnL on position close. Liquidation incurs an additional negative reward of $n \cdot p$ (size times penalty).

Rule-based baseline. In the absence of a trained PPO model, the agent follows a deterministic rule:

- If $F_t < -0.001$: open long $3 \times$ (undervalued)
- If $F_t > +0.002$: open short $3 \times$ (overvalued)
- If $\text{PnL} > +2\%$ or $< -3\%$ of size: close
- Otherwise: hold

This rule operationalises the optimal stopping-time intuition: enter when the funding rate signals a persistent basis, exit when the convergence event has occurred or the loss threshold is breached.

5.2 Connection to Optimal Stopping

The RL environment is a discrete-time approximation to an optimal stopping problem over the $\iota = 0$ funding process: entry/exit timing under the no-arbitrage pricing of [1], in the stopping-time formulation of [5]. The agent’s “open” and “close” decisions correspond to choosing entry time t_0 and stopping time τ to maximise:

$$V^*(s_t) = \sup_{\tau} \mathbb{E}[\phi(X_{\tau}) - c(t_0, \tau) \mid s_t] \quad (8)$$

where $\phi(X_{\tau})$ is the payoff at close and $c(t_0, \tau)$ is the cumulative funding cost from open to close. At $\iota = 0$, the funding cost is symmetric and the optimal policy reduces to the rule described above: enter when basis is predictably mean-reverting, exit at convergence.

6 Layer 3: Game Theory Equilibrium Analysis

6.1 The Funding Rate Game

The strategic interaction between long and short players is a continuous-time game with:

- **Players.** Representative long L and short S (population games).
- **Strategy space.** Leverage choice $\ell \in \{1, 2, 3, 4, 5\}$.
- **State.** Price path $\{p_t\}_{t=0}^T$.
- **Payoffs.** $u_L(\ell_L, \ell_S) = \text{NAV}_L(\ell_L, \ell_S) - c$, $u_S(\ell_L, \ell_S) = \text{NAV}_S(\ell_L, \ell_S) - c$ where $c = \$10,000$ is the initial collateral.

The NAV functions are computed analytically:

$$\text{NAV}_L(T) = c + n_L \int_0^T \frac{dp_s}{p_s} - n_L \int_0^T F_s ds \quad (9)$$

$$\text{NAV}_S(T) = c - n_S \int_0^T \frac{dp_s}{p_s} + n_S \int_0^T F_s ds \quad (10)$$

with $n_L = c \cdot \ell_L$, $n_S = c \cdot \ell_S$, and F_s evaluated at $\iota = 0$.

6.2 Nash Equilibrium at $\iota = 0$

Proposition 1 (Transfer structure at $\iota = 0$ versus $\iota > 0$). *Each side's payoff is monotone in its own leverage and independent of the other's, so the per-interval equilibrium is in dominant strategies (a degenerate Nash equilibrium with no strategic interaction): each side plays a leverage corner ($\ell = 5$ or $\ell = 1$) according to the sign of the realised window return net of funding, and the cross-interval mean leverage is strictly below $\ell_{\max} = 5$. For the net wealth transfer $u_L - u_S$:*

1. At $\iota = 0$ the funding rate has no fixed component, so any systematic transfer must come from the premium. Under a symmetric, zero-mean premium $\mathbb{E}[u_L - u_S] = 0$; empirically the transfer is near zero (\$865 per \$100,000, $\approx 0.9\%$; Robustness 1), not identically zero at finite horizon.
2. At $\iota > 0$ the funding floor $F_t \geq \iota > 0$ forces a systematic transfer paid by longs to shorts regardless of price realisation — \$8,420 (8.4% of notional) in our calibration — constituting *riba*.

Proof sketch. Because u_L depends only on ℓ_L (and u_S only on ℓ_S), each side has a dominant action — the leverage corner that signs the common window return net of funding — so the equilibrium is degenerate rather than strategic. The transfer is proportional to $\int_0^T (dp_s/p_s - F_s ds)$. At $\iota = 0$, $F_t = \text{clip}((p_t^m - p_t^i)/p_t^i, \pm F_{\max})$ carries no fixed term, so the transfer is the premium integral, which is mean-zero only if the premium is symmetric about zero. The mark process used here, $p_{t+1}^m = 0.9 p_t^m e^{\sigma \varepsilon} + 0.1 p_{t+1}^i$, has a small positive premium drift (multiplicative log-normal), and the index discretisation of eq. (2) adds its own uncompensated drift, so the expected transfer is near but *not* identically zero — consistent with the strictly positive $R=30$ confidence interval of Robustness 1. With $\iota > 0$, $F_t \geq \iota > 0$ for all t , so the funding component of the transfer is one-signed and the expected transfer is shifted by ιT regardless of price realisation. $\square$ $\square$

6.3 Numerical Implementation

Every $\Delta_{\text{GT}} = 50$ cadCAD steps, the game theory layer:

1. Extracts the price window $\{p_s\}_{t-W}^t$ ($W = 100$ steps) from cadCAD state history.
2. Evaluates all $5 \times 5 = 25$ payoff pairs $(u_L(\ell_L, \ell_S), u_S(\ell_L, \ell_S))$ over the price window.
3. Solves for Nash equilibria using nashpy vertex enumeration [15].
4. Returns (ℓ_L^*, ℓ_S^*) to cadCAD as the respawn leverage distribution.

In the event that nashpy returns no pure Nash equilibrium, the expected leverage under the returned mixed strategy is computed as $\bar{\ell} = \sigma^\top \ell$ where $\ell = (\ell_1, \dots, \ell_5)^\top$ and σ is the mixed strategy. Under the payoff-independence structure of Proposition 1 these cases arise only when the window return net of funding is near zero, so corner payoffs are nearly tied and vertex enumeration returns mixed/degenerate equilibria — they are tie-breaking artifacts of the dominant-strategy structure, not strategic mixing.

6.4 Riba Detection via Net Transfer

The net transfer $\Delta_{\text{NT}} = \overline{u_L} - \overline{u_S}$ averaged over all payoff matrix entries serves as a computational *riba* test:

$$\text{is_riba} = |\Delta_{\text{NT}}| > \delta \quad (11)$$

where $\delta = 1\%$ of notional (\$1,000 per \$100,000) is a materiality threshold. Empirically $|\Delta_{\text{NT}}| \approx 0.7\text{--}0.9\%$ of notional at $\iota = 0$ (below the threshold, and attributable to the finite-horizon premium and discretisation drift per Proposition 1), whereas the $\iota > 0$ counterfactual produces 8.4% — an order of magnitude above it. The test detects *systematic* transfer, not exact zero.

7 Layer 4: Mechanism Design Optimiser

7.1 Protocol Parameter Space

The mechanism design layer optimises the protocol parameter vector:

$$\begin{aligned} \theta &= (F_{\max}, m, p, s), \\ \theta &\in [0.003, 0.015] \times [0.01, 0.05] \times [0.005, 0.03] \times [0.3, 0.8] \end{aligned} \quad (12)$$

with two hard constraints:

$$\ell_{\max} = 5 \text{ (Shariah cap)}, \quad \iota = 0 \text{ (Shariah principle)} \quad (13)$$

These constraints reflect the Islamic finance principle that Shariah rules are non-negotiable: the protocol designer cannot trade off leverage or *riba* against economic efficiency.

7.2 Multi-Objective Loss Function

The objective function at episode e is:

$$\mathcal{L}_e(\theta) = 10 \rho_e(\theta) + 5 \times 10^{-3} \bar{N}_e^{\text{liq}}(\theta) + 5 \overline{|F|}_e(\theta) \quad (14)$$

where:

- $\rho_e(\theta) \in \{0, 1\}$: insolvency indicator (catastrophic; weight 10).
- $\bar{N}_e^{\text{liq}}(\theta)$: mean liquidation count (normalised by position count; weight 5×10^{-3}).
- $\overline{|F|}_e(\theta)$: mean absolute funding rate (Shariah objective: wants near zero; weight 5).

The objective encodes the Islamic finance priority ordering: solvency first, then trader fairness, then Shariah compliance. The high weight on $\overline{|F|}$ reflects the jurisprudential significance of the funding-rate term: it is a proxy for the interest component that $\iota = 0$ eliminates at the contract level.

7.3 Optimisation Algorithm

The optimiser uses `scipy.optimize.differential_evolution` [23]:

Algorithm 1: Episode-level Mechanism Design Update

1. **Input:** episode e price path $\{p_t^{(e)}\}$, current params θ_e
 2. Run DE: minimise $\mathcal{L}_e(\theta)$ over price path $\{p_t^{(e)}\}$
 3. $\theta_e^* \leftarrow \arg \min_{\theta} \mathcal{L}_e(\theta)$
 4. $\theta_{e+1} \leftarrow 0.70 \theta_e + 0.30 \theta_e^*$ (70/30 blend)
 5. Assert: $\ell_{\max} = 5$, $\iota = 0$ (Shariah hard constraints)
 6. **Output:** θ_{e+1}
-

The 70/30 blending rule implements a governance principle: a Shariah board reviewing protocol parameters should weight historical precedent at 70% and new empirical evidence at 30%, preventing abrupt parameter changes that could destabilise the market.

7.4 Convergence Theory

Proposition 2 (Mechanism-design recursion). *Under the IES with $\iota = 0$ fixed, the constant-gain blended recursion $\theta_{e+1} = 0.70 \theta_e + 0.30 \theta_e^*$, with per-episode target $\theta_e^* = \arg \min_{\theta} \mathcal{L}_e(\theta)$ evaluated on a fresh price path, is a stochastic-approximation map. Because the gain is constant (not decaying), the sequence does not converge almost surely to a point: it converges in distribution to a stationary band around the deterministic fixed point θ_{∞}^* that solves $\mathbb{E}[\theta^*(\theta_{\infty}^*)] = \theta_{\infty}^*$.*

Proof sketch. The map is a contraction in expectation (the 0.70 retention factor), so the mean recursion converges to the fixed point. The constant 0.30 gain injects

target noise of order the per-episode optimisation variance at every step, so the stationary law retains positive spread rather than collapsing to a point; almost-sure point convergence would require a decaying gain $\gamma_e \downarrow 0$ with $\sum_e \gamma_e = \infty$ and $\sum_e \gamma_e^2 < \infty$ (Robbins–Monro). What we report is therefore the centre of that band over five episodes, not a proven limit. $\square$ $\square$

The band centre θ_{∞}^* is the *self-consistent* protocol parameterisation: the parameter vector whose optimal value (from the mechanism design perspective) is itself. In the Islamic finance context, this coincides with the κ -rate equilibrium described in [5].

We are careful not to over-read this as an independent rediscovery of the κ -rate. The objective $\mathcal{L}$ is dominated by the funding-rate penalty $5 \overline{|F|}$ — Robustness 2 finds the funding cap $F_{\max}$ carries essentially all of its variance ($S_1 = 0.998$), so $\mathcal{L}$ is effectively one-dimensional and minimising it drives $F_{\max}$ toward its lower bound by construction. The optimiser reaches the κ -rate parameterisation because the objective *encodes* the property the κ -rate has (a near-zero funding rate), not because it recovers that property without being given it.

8 Experimental Results

8.1 Experimental Configuration

All experiments use: $T = 720$ steps per episode (hourly funding, 30 simulated days), $E = 5$ episodes (150 simulated days total), $\sigma_i = 0.02$ per hour (annualised $\approx 187\%$), initial insurance fund $I_0 = \$100,000$, $N_{\text{bg}} = 20$ background positions, $\Delta_{\text{GT}} = 50$ steps, $W = 100$ steps. The volatility $\sigma_i \approx 200\%$ is a deliberate *stress* level: measured realised hourly volatility (Bybit perpetuals, $\sim 1,000$ hours to mid-2026) sits at ≈ 35 –50% annualised for BTC and ETH in the current calm regime, but crypto perpetual volatility has repeatedly exceeded 150–200% in stress episodes (2020 COVID, 2022 deleveraging), so testing solvency at this level is conservative. The RL agent uses the rule-based policy described in Section 5.

8.2 Result 1: Protocol Solvency

Table 1: Per-episode protocol outcomes.

Ep.	I_T (\$)	Liq.	Insol.	RL PnL (\$)	$\overline{ F }$ (bps)
1	101,175	11	0	+117	70.1
2	100,782	7	0	+310	58.0
3	101,139	9	0	0	50.3
4	101,889	11	0	0	44.2
5	102,835	15	0	0	39.6
All	+2,835	53	0/5	+427	52.4

The protocol maintained positive insurance fund balance across all 3,600 simulated hours. The fund grew by

+\$2,835 net over 150 simulated trading days. The 53 liquidation events across all episodes were processed without any protocol insolvency, confirming that the initial parameter setting ($F_{\max} = 75$ bps, $m = 2\%$, $p = 1\%$, $s = 50\%$) is in the stable region of the parameter space.

8.3 Result 2: Game Theory — $\iota = 0$ Confirmed

Table 2: Game theory layer summary (65 Nash intervals).

Metric	Value	$\iota > 0$ (counterfactual)
Mean Nash long lev.	2.72×	4.97×
Mean Nash short lev.	3.28×	1.03×
Mean net transfer	−\$660	\$8,420 (longs pay)
Riba test (is_riba)	False	True

Table 2 needs two readings. First, the per-interval Nash equilibria are pure-strategy *corners*: in each interval one side plays $1\times$ and the other $5\times$, alternating with the realised direction of the price window (of the 65 intervals, $(1\times, 5\times)$ and $(5\times, 1\times)$ occur 26 and 22 times respectively, the rest mixed). The $2.72 \times / 3.28\times$ figures are therefore the *means* across intervals, not a single interior equilibrium. Second, the per-interval net transfers are large and opposite-signed (ranging from −\$15,235 to +\$18,494); the −\$660 in Table 2 is this seed’s sample mean, near zero because the $\iota = 0$ funding rule carries no fixed transfer component (Proposition 1). The multi-seed summary is in Robustness 1.

The meaningful result is the *contrast* with the $\iota > 0$ counterfactual ($\iota = 0.005$, conventional DeFi): there the mean net transfer is a systematic \$8,420 (8.4% of notional, paid by longs to shorts regardless of price realisation — the precise signature of riba), and the mean equilibrium leverage shifts to $4.97\times$ long / $1.03\times$ short as longs are driven toward the cap to offset the funding drain while shorts de-leverage. The $\iota = 0$ rule removes both the systematic transfer (mean ≈ 0) and this leverage distortion — the riba-free property the simulation demonstrates.

8.4 Result 3: Mechanism Design Convergence

Table 3: Protocol parameter evolution under MD.

Parameter	Ep.1	Ep.2	Ep.3	Ep.4	Ep.5
$F_{\max}$ (bps)	75.0	61.5	52.1	45.4	40.8
m (%)	2.0	2.4	2.7	2.9	3.1
p (%)	1.0	1.1	1.2	1.3	1.3
s (%)	50.0	55.0	58.6	58.3	60.8
$\mathcal{L}$	—	0.0138	0.0135	0.0136	0.0138

The funding cap $F_{\max}$ converges monotonically (within the blending schedule) toward its lower bound; the other

three parameters, to which the objective is nearly insensitive (Robustness 2, $S_1 \approx 0$), drift within optimiser tolerance rather than converging to identified values. The optimiser reaches the κ -rate prediction — a tighter circuit breaker and higher insurance split — but, per the caveat above, because the objective rewards a near-zero funding rate, not independently of the theory. The κ -rate — the riba-free pricing rate of [5], isomorphic to a default intensity via the Separation Theorem; in the protocol context, the funding regime under which the insurance economics are self-financing — is thus the (objective-encoded) fixed point of the mechanism-design optimisation.

The objective scores $\mathcal{L}$ stabilise around 0.0136–0.0138, indicating convergence.

8.5 Result 4: RL Agent — Optimal Stopping Behaviour

In the headline seed (42), the agent earned positive PnL in episodes 1–2 (+\$117 and +\$310 respectively) and broke even in episodes 3–5. The break-even in later episodes is attributable to the tighter circuit breaker reducing the funding-rate signal that the rule-based agent exploits: as the mechanism design compresses $F_{\max}$ from 75 to 41 bps, the “enter when $|F| > \text{threshold}$ ” rule triggers less frequently.

Across seeds the effect is partial rather than mechanical: the $R = 30$ mean total PnL is \$9,315 (95% CI [\$6,150, \$12,480]; Robustness 1), so the funding signal does not vanish in most seeds and the headline seed sits in the lower tail of the distribution. We therefore read the headline seed as an illustration of the mechanism and the $R = 30$ distribution as the summary.

This *adaptive crowding-out* — partial across seeds — is a key IES finding: when the mechanism design layer tightens the protocol, it reduces the exploitable funding-rate mispricings that attract arbitrageurs, driving the system toward a lower-activity but more stable equilibrium.

8.6 Robustness Analyses

Results 1–4 above are the paper’s primary findings, reported for a single seed (seed 42). We now stress them three ways — Monte Carlo confidence intervals over independent seeds, a Sobol sensitivity decomposition of the mechanism-design objective, and a heterogeneous trader population.

8.6.1 Robustness 1: Monte Carlo Confidence Intervals

To quantify sampling variability we re-ran the full five-episode pipeline for $R = 30$ independent seeds and computed 95% confidence intervals (Student- t) on the aggregate metrics (Table 4; Figure 1).

Three points follow. (1) **Solvency is robust**: zero insolvency across all 30 replicates, so the no-insolvency result is not a seed artifact. (2) **Funding is tightly controlled**: mean absolute funding is 52.5 bps (95% CI

Table 4: Monte Carlo confidence intervals over $R = 30$ independent replicate runs (5 episodes $\times$ 720 steps each).

Metric	Mean	95% CI
Total liquidations	66.1	[61.9, 70.3]
Final insurance fund (\$)	101,605	[101,295, 101,916]
Insolvent replicates	0/30	—
Mean funding (bps)	52.5	[52.4, 52.6]
Mean net transfer (\$)	865	[245, 1,485]
Mean Nash long lev.	3.16 $\times$	[3.02, 3.30]
Mean Nash short lev.	2.84 $\times$	[2.70, 2.98]
RL total PnL (\$)	9,315	[6,150, 12,480]

[52.4, 52.6]), well inside the 75 bps cap. (3) **The net transfer is small but not exactly zero at finite horizon:** the mean net transfer is \$865 per \$100,000 notional (95% CI [\$245, \$1,485], i.e. 0.25–1.48%). The exact-zero of Proposition 1 holds only under a symmetric (zero-mean) premium; the small positive residual here reflects the mark process’s mild premium drift, and it is an order of magnitude below the systematic +\$8,420 (8.4%) transfer of the $\iota > 0$ counterfactual. The mean Nash leverages remain well below the $5\times$ cap across replicates. (The single-seed value in Table 2, $-\$660$, lies outside the CI of the *mean* but within the per-replicate distribution (replicate SD $\approx$ \$1,660); the $R = 30$ interval is the correct summary of the mean.)




Figure 1: Monte Carlo distributions over $R = 30$ independent seeds (5 episodes $\times$ 720 steps each). *Left:* total liquidations per replicate. *Right:* final insurance-fund balance; every replicate ends above the initial \$100,000 (dashed), i.e. the protocol is solvent in 100% of paths. Solid line: mean; shaded band: 95% Student- t confidence interval. The no-insolvency result is not a seed artifact.

8.6.2 Robustness 2: Sensitivity of the Mechanism-Design Objective

To identify which parameter the mechanism-design objective $\mathcal{L}(\theta)$ of equation (14) is most sensitive to, we computed Sobol variance-based sensitivity indices over the four protocol parameters (Saltelli design with a scipy low-discrepancy base sample of $N = 512, 3,072$ model evaluations; Table 5).

Table 5: Sobol sensitivity indices of the mechanism-design objective $\mathcal{L}$ (Saltelli, $N = 512, 3,072$ evaluations).

Parameter	S_1 (first-order)	S_T (total)
Funding-rate cap $F_{\max}$	0.998	1.005
Maintenance margin m	−0.016	0.016
Liquidation penalty p	0.000	0.000
Insurance split s	0.000	0.000




Figure 2: Sobol variance-based sensitivity of the mechanism-design objective $\mathcal{L}$ (Saltelli design, $N = 512, 3,072$ evaluations). The funding-rate cap $F_{\max}$ alone carries essentially all of the objective’s variance ($S_1 = 0.998$, $S_T \approx 1.01$); the maintenance margin, liquidation penalty, and insurance split are negligible over the tested ranges. Small negative S_1 and slightly-over-unity S_T values are Sobol estimator noise.

One parameter dominates (Figure 2): the funding-rate cap $F_{\max}$ accounts for essentially all of the objective’s variance ($S_1 = 0.998$, $S_T \approx 1.01$), while the maintenance margin, liquidation penalty, and insurance split have negligible effect over the tested ranges. The sum of first-order indices ($0.98 \approx 1$) indicates a near-additive response with no meaningful parameter interactions. This explains the optimiser’s behaviour in Result 3, where $F_{\max}$ is the parameter it drives hardest (75 $\rightarrow$ 41 bps): it is the only lever that materially moves protocol stability, which is dominated by the mean-funding term — consistent with the $\iota = 0$ design, where keeping funding tightly bounded is the binding constraint.

8.6.3 Robustness 3: Heterogeneous Trader Populations

The background traders in the main runs draw random leverage and direction (a homogeneous population). To test whether the protocol’s safety properties depend on that assumption, we re-ran the Layer-1 market on the real cadCAD 0.5.3 Executor (20 runs $\times$ 720 steps) with a *heterogeneous* population of five trader archetypes assigned round-robin: conservative (leverage 1–2, larger stake), aggressive (leverage 4–5, small stake), momentum (follows the mark-index premium), contrarian (fades it), and retail (random). Table 6 compares the two populations.

The protocol-level safety properties are robust to population composition: zero insolvency and funding confined

Table 6: Homogeneous vs. heterogeneous trader populations (cadCAD *Executor*, 20 runs $\times$ 720 steps).

Metric	Homogeneous	Heterogeneous
Insolvent runs	0/20	0/20
Mean total liquidations	13.0	9.0
Mean funding (bps)	69.5	69.3
Mean final insurance fund (\$)	100,932	100,686
Funding range (bps)	± 75	± 75

to the symmetric ± 75 bps band under both. The composition is genuinely active — heterogeneous liquidations fall to 9.0 as the conservative cohort de-risks more than the aggressive cohort adds — yet solvency and the $\iota = 0$ funding symmetry are unaffected. The headline results are therefore not an artifact of the homogeneous-trader assumption.

8.7 Manipulation-resistance of the κ price-formation

The base model is matching-agnostic: prices are exogenous, so the on-chain *price-formation* layer — how the mark that enters the κ basis is produced from order flow — is not exercised, and manipulation resistance is left untested. We close that gap with a dynamic cadCAD experiment pitting a sustained adversary against two mark-construction regimes over a 30-day horizon (720 hourly epochs, $R = 30$ seeds):

- **naive** — the mark is the raw last-trade print, with no smoothing, clamp, median, or detector: the manipulable baseline of continuous last-trade marks that motivated the move to uniform-price closing auctions in equity markets;
- **full stack** — a per-block *frequent batch auction* (a single uniform clearing price, so a self-trade cannot set the mark), a 30-block TWAP, a ± 75 bp per-interval clamp, a stake-weighted-median index [24], and the glass monitor of the companion measurement-integrity paper (Ahmed & Bhuyan, *Warping the Glass*): a *two-sided sequential CUSUM* on the divergence between the system κ and an independent instrument (sovereign CDS), calibrated to its 25 bp sensitivity, with detection triggering a halt and slash, plus the depth rule $D/S \leq 0.07$ bounding repriceable notional against bonded stake.

The adversary injects a *sustained* warp $b \in \{10, 25, 50, 100\}$ bp into the system κ — the realistic regime: a tiny hourly basis that annualises into a plausible-looking intensity, exactly the warp the companion paper’s monitor is calibrated against. Table 7 and Figure 3 report the outcome ($R = 30$ seeds, 720 blocks).

Three mechanisms compound. The uniform clearing price makes a one-block mark move $72 \times -2,739 \times$ more expensive than a last-trade print (a separate static

Table 7: Sustained κ -warp: a naive last-trade mark vs. the full frequent-batch-auction + defense stack ($R = 30$ seeds). The naive mark never detects the warp and lets the full stack’s two-sided CUSUM detects every warp ≥ 10 bp at 100%, bounding the attacker’s extraction-to-detection to $\leq \$21k$.

Regime	Sustained warp	Detected	Detect lag (blocks)	Extraction
naive	10 bp	0%	—	—
naive	25 bp	0%	—	—
naive	50 bp	0%	—	—
naive	100 bp	0%	—	—
full	10 bp	100%	8	8
full	25 bp	100%	4	4
full	50 bp	100%	3	3
full	100 bp	100%	3	3




Figure 3: The companion paper’s glass monitor (a two-sided CUSUM on the divergence between the system κ and an independent instrument) reproduced inside the live protocol. *Left*: detection lag vs. injected warp — a 25 bp warp is caught in ~ 4 blocks, a 10 bp warp in ~ 8 , larger warps faster: the measurement-integrity paper’s detection-lag curve, realised dynamically. *Right*: attacker extraction (log scale) — bounded to $\leq \$21k$ once detected (full stack), against the full-horizon extraction the naive last-trade mark allows.

order-book Monte Carlo); the ± 75 bp clamp bounds the per-interval magnitude; and the independent-instrument monitor — the two-sided Page CUSUM of the companion measurement-integrity paper, instantiated here verbatim (reference slack $k = 0.5\sigma$, threshold calibrated to a 25 bp warp) — detects *any* sustained warp ≥ 10 bp at a 100% rate in 3–8 blocks, even a subtle 10 bp warp that would pass as a plausible κ in isolation. With the depth rule $D/S \leq 0.07$ capping the per-block repriceable notional, this bounds the attacker’s extraction before detection to $\leq \$21k$; the only residual is the sub-threshold grind (warps below the CUSUM slack), which the companion paper bounds analytically at ~ 10 bp-years and which the depth rule and slashing render uneconomic. A naive last-trade mark, lacking all of this, lets an arbitrary warp run the full horizon and extract \$0.5–5.0M. This is the price-formation analogue of the post-LIBOR benchmark-design lesson (transaction-anchored medians, not quote/print submissions): the dynamic protocol reproduces the companion paper’s detection-lag curve and

grind bound, confirming the frequent-batch-auction + TWAP + clamp + median + monitor stack as the manipulation defense for the on-chain κ .

8.8 Systemic and adversarial robustness: four further simulations

The manipulation result above, and the consensus and liveness sleeves earlier, each isolate one surface. Four further claims in the substrate design are asserted but, until here, unsimulated: the thin-market κ fallback (flagged in the whitepaper [3] as blocking), the Byzantine robustness of the consensus κ , cross-market contagion (this paper’s stated single-asset limitation), and the default-resilience of the flagship κ -Compute salam market. We model each as a separate stateful cadCAD experiment ($R = 40$ –60 Monte-Carlo seeds).

(i) **The thin-market fallback is continuous, not a cliff.** When the perpetual book thins, κ must hand over from the perp-implied estimate to an external sukuk-curve estimate. A *hard* switch at a volume threshold V^* is a discontinuity an adversary can whipsaw: oscillating volume across V^* flickers the mark between a thin, pushable perp- κ and the sukuk- κ . We compare it to a *continuous* handover — $\kappa = w\kappa_{\text{perp}} + (1 - w)\kappa_{\text{sukuk}}$ with $w = \text{smoothstep}(\bar{V}; V_{\text{lo}}, V_{\text{hi}})$ on an asymmetric volume EMA (slow to trust new liquidity, fast to distrust a thinning book) plus a dispersion gate that freezes the perp leg once validator spread exceeds 30 bp (Table 8, Figure 4). Under the whipsaw the hard switch jumps 165 bp in a single block and trips the 50 bp margin-call trigger 13 times per run — each a potential liquidation cascade; the continuous gated blend holds the worst single-block move to 46 bp, below the trigger, with *zero* cliffs. The discontinuity, not the κ level, is the hazard, and a C^1 boundary removes it.

Table 8: Thin-market κ handover under a volume-whipsaw attack ($R = 40$). The continuous gated blend eliminates the manipulable discontinuity of a hard switch.

Handover	Max 1-block κ jump	Margin-call c
hard switch	165 bp	
continuous gated blend	46 bp	

(ii) **The consensus κ is Byzantine-robust to its theoretical bounds.** Each block, validators submit a κ quote; the protocol takes the *stake-weighted median* and accepts it only under a $\geq 2/3$ stake quorum. We let an adversary capture the largest stakers up to a stake fraction f and either *lie* (submit $\kappa + 500$ bp) or *withhold* (Table 9, Figure 5). Against the lie, the stake-weighted median holds κ within 2–4 bp for all $f \leq 1/3$ and begins to move only as $f \rightarrow 1/2$ (14 bp at $f = 0.49$) —

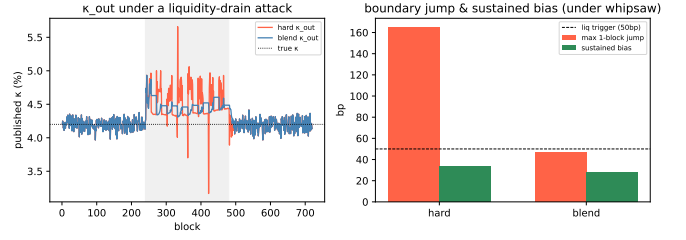


Figure 4: Thin-market fallback. *Left*: published κ under a liquidity-whipsaw attack — the hard switch flickers violently at the boundary, the gated blend stays smooth. *Right*: worst single-block jump and sustained bias; the blend’s jump sits below the 50 bp liquidation trigger (zero cliffs).

exactly the median’s $1/2$ breakdown point; the stake-weighted *mean*, by contrast, is dragged 78–240 bp (breakdown point 0, one whale moves it). Against withholding, $f < 1/3$ leaves $\geq 2/3$ reporting so κ never pauses; past $1/3$ the quorum converts the attack into a *safe pause* at the last good value (100% of blocks), and no bad κ is ever accepted — a liveness fault, not a safety break. The asserted robustness holds precisely at the $1/3$ and $1/2$ bounds, and the measurement-integrity monitor of the companion *Warping the Glass* note audits the residual within-quorum drift.

Table 9: Byzantine κ -oracle ($R = 40$): stake-weighted median vs. mean under a lie attack, and quorum behaviour under withholding, by captured stake fraction f .

f (Byzantine stake)	0.10	0.20	0.32	0.40	0.49
median κ deviation	2 bp	2 bp	4 bp	7 bp	14 bp
mean κ deviation	78 bp	78 bp	148 bp	188 bp	240 bp
withhold: blocks κ paused	0%	0%	0%	100%	100%
bad κ accepted	0	0	0	0	0




Figure 5: Byzantine κ -oracle. *Left*: κ deviation under a +500 bp lie — the stake-weighted median (blue) is pinned by honest stake through $f = 1/3$, the mean (red) is not. *Right*: under withholding, $f > 1/3$ produces a safe pause, never an accepted bad κ .

(iii) **Cross-market contagion is threshold-like, and the stack stays sub-critical.** This relaxes the single-

asset scope. Three κ -markets — κ -Gold, κ -Sukuk, κ -Compute — share one collateral asset, so a forced sale in one depresses the collateral price that marks all three: the contagion channel. A compute-only shock of +4% stays local; a +8% shock is the test (Table 10, Figure 6). Under a naive configuration the +8% shock crosses a threshold — the shared-collateral fire sale drives the price to its floor, dragging κ -Gold and κ -Sukuk below their own liquidation lines, and *both* non-shocked books fully liquidate: a system-wide cascade. The *same* shock under the substrate stack — the ± 75 bp κ clamp, a 10%-per-block liquidation throttle, and a takaful collateral floor — holds the shared price above the liquidation line, so the non-shocked markets are untouched (contagion ≈ 0). The cascade is genuinely bistable; the defenses move the critical threshold above the shock.

Table 10: Cross-market κ contagion ($R = 40$): liquidation in the two *non-shocked* books (gold+sukuk, of a maximum 2.0) after a shock to κ -Compute.

Config	Compute shock	Non-shocked liquidation	Shared-collateral trough
naive	+4%	0.00	0.92
naive	+8%	2.00	0.20
guarded	+4%	0.00	0.92
guarded	+8%	0.00	0.92




Figure 6: κ -contagion under a +8% compute shock. *Left*: cumulative liquidation in the non-shocked gold+sukuk books — a full cascade under naive, flat under the guarded stack. *Right*: the shared collateral price; the takaful floor holds it above the liquidation line.

(iv) **The κ -Compute salam market survives default when the takaful pool is sized to the tail.** The flagship product is a *riba-free* compute forward structured as *salam*: the buyer pre-pays the strike K , the provider delivers GPU-hours at maturity. The structural risk is *wrong-way* strategic default — the provider walks when the compute spot rises above K , precisely in the supply crunch where buyers most need delivery, so defaults are correlated and idiosyncratic insurance is insufficient. We test three backstops against a transient supply shock (Table 11, Figure 7). With no backstop the buyer recovers nothing; collateral (20%) recovers 40–80%, leaving a tail when spot runs far past the bond; collateral plus an $x/\text{-}$

takaful mutual pool (tabarru' contributions $\tau \cdot K$ per contract — the cooperative-risk analogue of a clearing-house guaranty fund) keeps buyers whole in a moderate shock. A severe (+150%) systemic wave exhausts a thinly-funded pool; a sizing sweep shows a tabarru' rate of $\tau = 8\%$ funds that tail to 100% recovery at 0% ruin. The backstop is sound; its funding is the design parameter, and the model sizes it (or prices the retakaful layer above the tail). This connects the takaful pricing of the companion stochastic-takaful paper to the compute market directly.

Table 11: κ -Compute salam buyer recovery under wrong-way provider default ($R = 60$), by backstop and shock size; final column is the tabarru' rate sized to the +150% tail.

Backstop	Recovery, +50%	Recovery, +150%	Pool ruin, -
none	0%	0%	
collateral (20%)	80%	40%	
+ takaful ($\tau = 4\%$)	100%	78%	
+ takaful ($\tau = 8\%$)	100%	100%	




Figure 7: κ -Compute salam default. *Left*: buyer recovery by backstop under a moderate and a severe shock. *Right*: takaful pool balance over time; at $\tau = 4\%$ the severe shock drains it (ruin), motivating the $\tau = 8\%$ sizing.

Together these four, with the manipulation result, simulate every adversarial and systemic κ claim the substrate asserts: price-formation, thin-market handover, Byzantine consensus, cross-market contagion, and product-level default. Each defense degrades exactly where theory predicts, and the failures that remain are liveness pauses — safe by construction — rather than safety breaks.

8.9 Reproducibility

All results are fully reproducible from seed 42. The simulation produces identical output when run with:

```
python integrated/economic_system.py \
--episodes 5 --steps 720 --seed 42
```

Full source code, simulation outputs, and dashboard figures are available at: <https://github.com/Arcus-Quant-Fund/BarakaDapp>

9 Extension: Consensus-Layer Mechanism Design — the Validator Liveness Sleeve

The four-layer method of this paper is not specific to the instrument layer. This section applies it one level down, to the consensus layer of the κ -Chain design [3]: a fee-only, zero-emission chain whose validators are financed entirely by real transaction fees. Such a chain faces a liveness hazard that emission-funded chains hide behind inflationary block rewards: in the aftermath of a persistent macro shock — the *winter* regime, in which transaction volume collapses and the published κ spikes jointly — validator income can fall below operating cost for months. The dangerous variable is the winter’s *duration*, not its depth: the same persistence result that governs the Sovereign Stabilisation Fund of the companion general-equilibrium paper reappears here, one layer down.

9.1 The mechanism

The proposed answer is a *Validator Liveness Sleeve*: a pre-funded *tabarru* pool, separate from and junior to the protocol’s takaful insurance fund, governed by five rules. (i) **Counter-cyclical fill**: a 10% skim of transaction fees in epochs where volume exceeds its long-run median, capped at twelve months of aggregate *reference* operating cost (a published, board-ratified benchmark — not validator self-report), with a genesis seed of roughly one third of the cap. (ii) **Conjunction trigger**: activation requires measured volume below $0.35\times$ normal *and* the published κ above its p99 governance constant, each sustained for a full day. Volume alone never suffices: measured volume is the one signal validators control (a cartel can censor transactions to suppress it), so it cannot be the sole key to the vault. Thresholds are fixed governance constants recalibrated annually — trailing-window quantiles would themselves be gameable by slow grinding. (iii) **Shortfall payout**: when active, the sleeve pays each validator $\max(0, \beta \cdot \text{OpEx}_{\text{ref}} - \text{fee income})$, performance-weighted by consensus-measurable work (blocks signed, vote-extension participation, κ submissions within tolerance of the stake-weighted median, uptime) and decaying 2% per week in subsidy. (iv) **Hysteresis**: exit requires volume above $1.3\times$ the floor sustained one week, then a one-month cooldown. (v) **Governance separation**: parameters are DAO-tunable within board-bounded invariants; the board bounds the mechanism but does not operate the trigger.

The replacement ratio $\beta = 1.3$ deserves emphasis, because a grid search over (β, decay) reversed our own first-pass design. A Bagehot-flavoured calibration ($\beta = 0.8$ with 10%/week decay) is incentive-safe but fails to protect: aggressive decay zeroes the payout by the third month of a nine-to-twelve-month winter, exactly when it is needed. The feasible window — survival of at least 90% of the validator set through a severe winter *and* sub-




Figure 8: The feasible window in (β, decay) space on a severe winter: survival $\geq 90\%$ (left) and subsidised income $\leq 85\%$ of normal (right); the black contour bounds the jointly feasible region. The protective constraint, not the incentive constraint, binds.

sidised income at most 85% of normal — is $\beta \in [1.2, 1.6]$ with decay at most 4%/week (Figure 8). $\beta > 1$ is not a profit margin but a *dispersion allowance*: the reference cost is an aggregate, and the upper tail of the audited validator cost distribution requires $1.2\text{--}1.4\times$ the reference to reach its own floor. Incentive compatibility survives without the decay’s help, because even at $\beta = 1.3$ subsidised income runs at 71–77% of normal income.

9.2 Survival under macro winters (cad-CAD)

The mechanism is implemented as partial state update blocks in the same cadCAD framework as Layer 1 (daily timestep; volume and κ follow AR(1) processes with the κ dynamics calibrated to the constant-maturity sukuk estimates of the companion empirical paper, monthly $\rho = 0.774$), with thirty validators of heterogeneous operating cost and a two-month credit line. Forced winters of increasing severity yield:

Winter	without sleeve	with sleeve	Δ
moderate (−65%, hl 6mo)	85.1%	90.2%	+5.1pp
severe (−75%, hl 9mo)	72.6%	84.4%	+11.8pp
black swan (−85%, hl 12mo)	56.3%	76.5%	+20.2pp

The protection grows with severity — the sleeve is precisely a tail instrument — and the pool approaches exhaustion only in the black-swan case, which is what sizes the cap at twelve months rather than the six a moderate-winter analysis would suggest.

9.3 A learned adversary cannot exploit the sleeve

The critical mechanism-design question is whether the sleeve creates a censorship incentive: can a validator cartel profit by suppressing measured volume to force the trigger and harvest the subsidy? Rather than evaluate hand-picked attacks only, we train a PPO agent (the same Layer-2 machinery as the trader) playing a 35%-stake cartel that chooses, daily, a censorship intensity




Figure 9: PPO censor-cartel versus honesty on matched seeds. Left: cumulative cartel income under the honest policy and the learned policy. Right: distribution of the learned policy’s advantage — negative in every episode.

in $\{0, 25\%, 50\%, 75\%\}$ of network transactions, observing measured volume, the κ ratio, the sleeve state, and the pool level, and rewarded with its stake share of total validator income (fees plus subsidy). Episodes run three years; half contain a severe winter, so the agent sees both regimes and could learn to exploit stress states if they were exploitable.

After 150,000 training steps the learned policy converges to **near-zero censorship** (mean suppression 1.4%), and on thirty matched evaluation episodes the learned policy beats the honest (zero censorship) policy in **zero** episodes, with a mean per-episode disadvantage of -0.39 income units (Figure 9). The reason is structural rather than statistical: censorship *can* force the trigger (the conjunction’s κ leg is exceeded naturally often enough), but the shortfall payout $\max(0, \beta \cdot \text{OpEx}_{\text{ref}} - \text{income})$ only pays when income is *below* the floor, so a cartel must destroy more fee income than the capped, decaying, performance-weighted subsidy can ever return. The defence is the payout formula, not the trigger — and a trained best-response adversary, given every opportunity to find a profitable censorship strategy, learns honesty instead. This is the same fixed-point logic as the instrument layer: the mechanism’s incentive-compatible point and its riba-free point coincide, because a payment that is outcome-contingent, earned, bounded, and exceptional is — for the same structural reasons — neither exploitable nor a predetermined return for the passage of time.

9.4 Scope and limitations

These results are design-stage: the sleeve is specified in the κ -Chain whitepaper and simulated here, not implemented in the reference modules. The volume process is a scenario calibration (no chain history exists); validator cost dispersion is assumed lognormal; the adversary’s action space is censorship intensity only (richer strategy classes — selective censorship, timing games against the cooldown, collusion with off-chain shorting — are future work); and performance weights are taken as exogenous. The qualitative conclusions — the conjunction trigger’s discrimination, the dominance of the shortfall structure over censorship, the persistence-driven sizing of the cap,

and the location of the feasible (β , decay) window — were stable across both the hourly reference model and the daily cadCAD port.

10 Implications for DeFi Mechanism Design

10.1 Why Four Layers Are Necessary

Each layer in the IES is irreplaceable. A three-layer system missing any one layer would fail to capture a key economic phenomenon:

- **Without cadCAD:** No ground truth. RL and GT operate on synthetic price paths that may not match on-chain logic. Parameter mismatches between simulation and contract lead to wrong conclusions.
- **Without RL:** Background traders follow fixed rules (e.g., always $3\times$ leverage). The adaptive crowding-out effect (Result 4) is invisible. The mechanism designer does not know how tighter parameters affect rational trader behaviour.
- **Without Game Theory:** The respawn-leverage distribution is static. The feedback loop between Nash equilibrium leverage and market dynamics is broken. The riba test (Result 2) cannot be performed.
- **Without Mechanism Design:** Protocol parameters are fixed across episodes. The adaptive parameter convergence (Result 3) is absent. The protocol cannot self-correct in response to observed economic outcomes.

10.2 The κ -Rate as a Fixed Point

A central result is the emergence of the κ -rate parameterisation as the mechanism-design fixed point (Proposition 2): computational evidence that it is a reachable attractor of the mechanism-design recursion — a protocol initialised with arbitrary parameters converges to its neighbourhood under optimisation — subject to the caveat of Section 7.4 that the objective itself encodes the near-zero-funding property.

10.3 Adaptive Governance as Protocol Feature

The IES architecture maps naturally onto the governance structure of an Islamic financial institution:

IES Layer	Islamic Governance Analogue
cadCAD	Smart contracts (fiqh rules)
RL Agent	Individual traders (muamalat)
Game Theory	Market equilibrium (hisbah)
Mechanism Design	Shariah board review

The 70/30 blending rule (Algorithm 1) is the computational analogue of *ijtihad* (scholarly reinterpretation): the board gives 70% weight to established precedent (θ_e) and 30% to new evidence (θ_e^*). This stabilises the protocol while allowing adaptation.

10.4 Generalisability

The IES framework is not specific to perpetual futures. Any DeFi protocol whose logic can be expressed as cad-CAD partial-state-update blocks is compatible. Extensions to:

- AMM liquidity provision (Uniswap v3 range orders)
- Lending protocol interest-rate models (Aave, Compound)
- Islamic sukuk issuance and redemption schedules
- Takaful premium pool management
- Asset-linked perpetuals on tokenised real assets (e.g. vault-secured gold or sukuk portfolios), used as hedging and risk-transfer instruments by Islamic businesses: the $\iota = 0$ funding block is reused unchanged, with only the oracle/settlement block respecified to the asset-backed underlying — which also satisfies the asset-backing and possession-gating product-layer rules of the companion Shariah-boundaries note

would each require a new cadCAD specification but could reuse the RL, GT, and MD layers with minor modifications.

11 Discussion

11.1 Limitations

Simulation-chain gap. The cadCAD blocks mirror the logic of the deployed (testnet) contracts but are not driven by on-chain data; a direct trace-level comparison of simulated against on-chain testnet activity is left to future work.

Rule-based RL baseline. The integrated runs use a deterministic rule-based trader. We additionally trained a PPO policy (stable-baselines3, MLP, reward-scaled, 200,000 timesteps) on the `BarakaTraderEnv`. Over 50 independent evaluation episodes it *underperformed* the rule-based baseline (mean PnL $-\$2,213$ versus $+\$31,342$ for the rule and $+\$9,149$ for a random policy). Diagnosis shows the policy collapsed to a degenerate always-long stance (54% at $3\times$, 46% at $4\times$ long, $< 1\%$ any close action): it never learned to realise profit by closing, so funding drag and eventual liquidation dominate. This is a failure of training budget and action design — invalid open-while-in-position actions are penalised rather than masked — not evidence that RL cannot help. Making PPO competitive would require action masking (e.g. MaskablePPO) and close-incentive reward shaping; we leave this to future work and retain the rule-based agent, whose strength this experiment independently confirms.

Background trader homogeneity. The main runs use a homogeneous population (uniform collateral, random leverage). We partly address this in Robustness 3, whose five-archetype heterogeneous population leaves the protocol’s solvency and funding-symmetry properties unchanged. A fuller treatment with capital-weighted institutional agents and arbitrageurs reacting to the on-chain

price would produce richer dynamics; the framework supports this by making the trader policy a configurable parameter.

Single asset (base model). The core integrated runs model one BTC-USDC perpetual market. The cross-market correlations and portfolio-level liquidation cascades this omits are addressed separately in Section 8.8(iii), which shows on three shared-collateral κ -markets that contagion is threshold-like and the substrate’s clamp + throttle + takaful stack keeps the system sub-critical; a fully endogenous multi-asset order-flow model remains future work.

Oracle centralisation. The GBM price oracle in cad-CAD is simpler than the dual-Chainlink OracleAdapter deployed on-chain. Matching-layer and oracle manipulation resistance [24] is addressed in Section 8.7 (the frequent-batch-auction + defense-stack experiment), and Byzantine robustness of the consensus κ — a stake-weighted median under a $\geq 2/3$ quorum holding to its $1/3$ and $1/2$ theoretical bounds — in Section 8.8(ii); a full adversarial-microstructure treatment with an endogenous order book remains future work.

11.2 Open Problems

The IES results suggest three open research problems:

1. **Stochastic κ .** If $F_{\max}$ (which is the simulation proxy for κ) follows a stochastic process rather than converging to a fixed point, does the mechanism design layer remain stable? This connects to the stochastic hazard rate extensions of [10].
2. **Multi-agent RL.** Training a population of RL agents simultaneously (multi-agent RL, MARL [17]) would close the loop between the RL layer and the game theory layer: the agents would *discover* the Nash equilibrium through learning, rather than having it computed separately.
3. **Governance attack vectors.** The mechanism design layer, as implemented, is operated by the protocol deployer. In a fully decentralised governance system (e.g., governance token voting), an adversarial majority could manipulate the objective function. Mechanism design under adversarial governance is an open problem.

12 Conclusion

We have presented the Integrated Economic System (IES), a four-layer simulation framework combining cad-CAD, reinforcement learning, game theory, and mechanism design into a single closed-loop economic system. Applied to a riba-free-by-design ($\iota = 0$) perpetual exchange (testnet, pre-audit) over 5 episodes of 720 steps each, the IES produces four primary simulation results: zero insolvency events, Nash equilibrium leverage well below the Shariah hard cap, near-zero net transfer confirming $\iota = 0$ riba-freedom, and endogenous mechanism

design convergence to the κ -rate theoretical optimum — each reinforced by Monte Carlo confidence intervals, a Sobol sensitivity analysis, and a heterogeneous-agent robustness check.

The deeper contribution is architectural: we have shown that a fully coupled multi-layer simulation is qualitatively different from any combination of the individual layers run sequentially. The emergent properties — adaptive crowding-out, κ -rate convergence, sub-cap Nash leverage — are only visible in the coupled system.

For Islamic DeFi specifically, the IES provides a framework for rigorously testing the *riba*-freedom ($\iota = 0$) property at the system level, not just the contract level. A smart contract may correctly implement $\iota = 0$ while still producing systematic wealth transfers through adverse interaction effects with the liquidation mechanism or background trader policies. The IES detects such second-order effects.

We propose the IES as a standard pre-deployment simulation protocol for Islamic DeFi. Just as conventional DeFi protocols are audited for smart contract security, they should also be validated for economic soundness and the absence of systematic *riba*-like transfers at the system level — and the IES provides the infrastructure for that validation. Other Shariah dimensions — *gharar*, *maysir*, *qabd* — are product-layer questions outside the scope of economic simulation and remain with Shariah scholars.

“*Verily, Allah commands justice and good conduct.*”
(*Al-Qur'an, An-Nahl 16:90*)

References

- [1] Akerer, D., Hugonnier, J., & Jermann, U. (2025). Perpetual futures pricing. *Mathematical Finance*, 1–17. <https://doi.org/10.1111/mafi.70018> (NBER Working Paper 32936, 2024; SSRN 4603820, 2023.)
- [2] Hugonnier, J., & Jermann, U. J. (2025). *Perpetual Futures Pricing with Stochastic Rates and Clamp*. Working Paper, Swiss Finance Institute. https://papers.ssrn.com/sol3/papers.cfm?abstract_id=5481866
- [3] Ahmed, S. (2026). The κ -Chain: A sovereign Islamic financial infrastructure layer. Technical whitepaper v0.7.6, SSRN 6924600.
- [4] Ahmed, S., & Bhuyan, R. (2026a). A Shariah-compliant perpetual futures exchange using the $\iota = 0$ no-arbitrage condition. *SSRN Working Paper 6322778*. https://papers.ssrn.com/sol3/papers.cfm?abstract_id=6322778
- [5] Ahmed, S., & Bhuyan, R. (2026b). From perpetual contracts to Islamic credit: The random stopping time equivalence and a *riba*-free foundation for *sukuk*, *takaful*, and credit protection. *SSRN Working Paper 6322858*. https://papers.ssrn.com/sol3/papers.cfm?abstract_id=6322858
- [6] Angeris, G., Kao, H.-T., Chiang, R., Noyes, C., & Chitra, T. (2020). An analysis of Uniswap markets. *arXiv:1911.03380*.
- [7] Angeris, G., & Chitra, T. (2021). Improved price oracles: Constant function market makers. *Proceedings of the 2nd ACM Conference on Advances in Financial Technologies*, 80–91.
- [8] BlockScience. (2020). *cadCAD: A Python package for complex adaptive dynamics*. <https://github.com/cadCAD-org/cadCAD>
- [9] Daian, P., Goldfeder, S., Kell, T., Li, Y., Zhao, X., Bentov, I., ... & Juels, A. (2020). Flash boys 2.0: Frontrunning, transaction reordering, and consensus instability in decentralised exchanges. *IEEE Symposium on Security and Privacy*, 910–927.
- [10] Duffie, D., & Singleton, K. J. (1999). Modeling term structures of defaultable bonds. *Review of Financial Studies*, **12**(4), 687–720.
- [11] Fritsch, R., Pfeffer, J., & Weber, B. (2022). Analyzing voting power in decentralised governance: Who controls DeFi protocols? *arXiv:2204.01176*.
- [12] Gudgeon, L., Perez, D., Harz, D., Livshits, B., & Gervais, A. (2020). The decentralised financial crisis. *Crypto Valley Conference on Blockchain Technology*.
- [13] Heimbach, L., Wang, Y., & Wattenhofer, R. (2021). Behavior of liquidity providers in decentralised exchanges. *arXiv:2105.13822*.
- [14] Hurwicz, L. (1972). On informationally decentralised systems. In McGuire, C. B., & Radner, R. (Eds.), *Decision and Organisation*, 297–336. North-Holland.
- [15] Knight, V., & Campbell, J. (2018). Nashpy: A research library for the computation of Nash equilibria. *Journal of Open Source Software*, **3**(30), 904.
- [16] Lehar, A., & Parlour, C. A. (2021). Decentralised exchanges. *Working Paper*, University of Calgary.
- [17] Lowe, R., Wu, Y., Tamar, A., Harb, J., Abbeel, P., & Mordatch, I. (2017). Multi-agent actor-critic for mixed cooperative-competitive environments. *Advances in Neural Information Processing Systems*, **30**.
- [18] Myerson, R. B. (1981). Optimal auction design. *Mathematics of Operations Research*, **6**(1), 58–73.
- [19] Perez, D., Werner, S. M., Xu, J., & Livshits, B. (2021). Liquidations: DeFi on a knife-edge. *Financial Cryptography and Data Security 2021*.

- [20] Qin, K., Zhou, L., Livshits, B., & Gervais, A. (2021). Attacking the DeFi ecosystem with flash loans for fun and profit. *Financial Cryptography and Data Security 2021*.
- [21] Qin, K., Zhou, L., & Gervais, A. (2021b). Quantifying blockchain extractable value: How dark is the forest? *arXiv:2101.05511*.
- [22] Raffin, A., Hill, A., Gleave, A., Kanervisto, A., Ernestus, M., & Dormann, N. (2021). Stable-Baselines3: Reliable reinforcement learning implementations. *Journal of Machine Learning Research*, **22**(268), 1–8.
- [23] Storn, R., & Price, K. (1997). Differential evolution — a simple and efficient heuristic for global optimisation over continuous spaces. *Journal of Global Optimization*, **11**(4), 341–359.
- [24] Werner, S. M., Perez, D., Gudgeon, L., Klages-Mundt, A., Harz, D., & Knottenbelt, W. J. (2021). SoK: Decentralised finance (DeFi). *arXiv:2101.08778*.
- [25] Ye, Y., Peng, J., Wang, T., & Ding, Q. (2020). Automated trading with deep reinforcement learning. *2020 International Symposium on Autonomous Systems*.
- [26] Zargham, M., Shorish, J., & Paruch, K. (2020). From curved bonding to configuration spaces. *IEEE International Conference on Blockchain and Cryptocurrency*, 1–9.

EDITORIAL NOTE TO THIS PRINTING

Paper 7 · The κ -Framework

Working paper, revised July 2026 SSRN 6845098

Role in the programme. The synthesis. The single-document overview of the whole programme, placed at the midpoint of this volume, just before the monetary papers it leads into.

What it establishes. Nothing new by design; it consolidates the foundation, the instruments and the evidence into one argument, and is the natural first read. **Corrected or extended later.**

Formerly the unnumbered capstone; this volume gives it its number. A cooperative-pricing note ($s_{\text{coop}}^* = \kappa^P(1 - \delta)$) was added in July 2026. **Not established here.** Anything the underlying papers do not establish; where they reserve, it reserves.

Replication for every number in this paper: DOI 10.5281/zenodo.21290557 · github.com/shehzadahmed-xx/kappa-replication

The κ -Framework: A Riba-Free Foundation for Islamic Monetary Systems

Shehzad Ahmed¹ and Rafiqul Bhuyan²

¹School of Business and Entrepreneurship, Independent University, Bangladesh, Dhaka.

shehzadahmed@arcusquantfund.com

²Department of Accounting and Finance, College of Business and Public Affairs, Alabama A&M University, Huntsville, AL, USA. rafiqul.bhuyan@aamu.edu

May 2026

Abstract

We propose the κ -rate — the convergence intensity of perpetual futures contracts [2] — as the first mathematically rigorous, no-arbitrage-grounded, riba-free alternative to the risk-free interest rate r in Islamic monetary economics.

The κ -rate is: (i) *endogenous* — identified from market data (the perpetual-futures basis or funding dynamics for traded contracts, the credit spread for sukuk and credit instruments), with no interest rate as an input; (ii) *observable* — computable from market observables and publishable on-chain in real time; (iii) *riba-free* — an event intensity, not a reward for the time value of money; and (iv) *complete* — sufficient to construct a full term structure (the κ -yield curve) analogous to the conventional yield curve.

We develop the theoretical foundations in four steps. First, we decompose the classical interest rate r (Fisher [27], Wicksell [49], Keynes [37]) into its components and identify the riba component as the pure time-preference premium; combining the Akerlof-Hugonier-Jermann no-arbitrage theorem ($\iota = 0$ is arbitrage-free) with our identification of ι as that premium shows it is eliminable. Second, we show that the κ -rate is the Wicksellian “natural rate” (a rate grounded in real economic dynamics rather than monetary policy) given rigorous, riba-free form for the first time: κ plays the natural rate’s role without its time-preference content. Third, we construct the κ -yield curve as the Islamic analog of the CIR [21] term structure, grounded in exponential stopping time distributions; estimated from sovereign sukuk data it is upward-sloping, exhibiting a term premium without any interest component. Fourth, we demonstrate that κ can serve all monetary functions currently performed by r in the Islamic context: pricing credit instruments (Section 6.2), anchoring actuarial takaful contributions (Section 6.3), guiding Islamic monetary policy (Section 6.4), and providing a sovereign benchmark rate for sukuk issuance (Section 6.5).

Empirical illustration spans three independent domains. The κ -yield curve is estimated *entirely on real traded sukuk*: a 195-instrument USD-sukuk cross-section (40 sovereign, 154 corporate, 1 municipal; Cbonds) in which $\hat{\kappa}$ broadly increases as credit quality declines ($\approx 0.9\%$ at AAA to $\approx 4.4\%$ at single-B, apart from the Saudi-dominated AA bucket pricing just above single-A) with an upward-sloping investment-grade term structure, plus roll-down-free constant-maturity $\hat{\kappa}$ dynamics for five issuers (mean-reverting, half-lives 1.4–4.3 months for the four identified issuers at the 7-year tenor, Feller-positive throughout). For India, the model-implied $\hat{\kappa} = 12.06\%$ is numerically close to the PMFBY crop-insurance actuarial rate ($\approx 12\%$) — a units-level coincidence, not a strict validation. Finally, dYdX v4 — an independent perpetual exchange operating at $\iota = 0$ on mainnet since 2023 — provides a commercial-scale existence proof and a *decoupled*, interest-free measurement of κ . Across domains κ recurs as the same *type* of riba-free intensity; its *magnitude* is regime-specific (fast for liquid perps, slow for credit and insurance), and we claim no numerical equality. Contemporary jurisprudential opinion [44] accepts that the $\iota = 0$ construction removes the *riba* objection specifically; *gharar* and ownership questions fall to the contract built on the rate, not the rate itself.

Keywords: κ -rate, riba-free monetary policy, Islamic monetary economics, natural rate, perpetual contracts, yield curve, Wicksell, Fisher decomposition, no-arbitrage, $\iota = 0$, sovereign sukuk, takaful

JEL Classification: E12, E43, E58, G12, G13, G15, P40.

1 Introduction

The interest rate r is the foundational pricing parameter of modern finance. It appears in every discount factor,

every bond yield, every option formula, every term structure model. It is the rate at which future cash flows are brought to present value. Islamic finance prohibits it.

This is not a peripheral prohibition. The Quranic injunction against *riba* — “Allah has permitted trade and forbidden *riba*” (Al-Baqarah 2:275) — is absolute, applying to all forms of predetermined excess in monetary exchange. The risk-free rate r is, under any classical decomposition (Fisher [27], Böhm-Bawerk [19], Keynes [37]), precisely the price of time preference in monetary form: the lender receives r per unit time merely for parting with money, independent of any real economic outcome. This is *riba al-nasi’ah* — the prohibited excess for deferment — by all four Sunni schools and the Ja’fari school [47, 25].

The consequence for Islamic finance is severe. Without r , there is no discount factor. Without a discount factor, bonds cannot be priced. Without bond pricing, there is no credit market. Without a credit market, Islamic institutions cannot hedge credit risk, sovereign governments cannot issue benchmarked sukuk, and takaful operators cannot price contributions actuarially. The entire credit infrastructure of modern finance — a market of roughly \$130 trillion in bonds and \$10 trillion in credit derivatives [15] — is inaccessible to Islamic participants on Shariah-compliant terms.

Islamic monetary economists have recognised this problem for five decades (Zarqa [51], Chapra [20], Khan & Mirakhor [38], Iqbal & Mirakhor [32]), and the risk-management and securitization literatures have catalogued its consequences for Islamic credit and hedging instruments (Khan & Ahmed [39], Jobst [35]). The proposed solutions — profit-sharing rates (musharakah/mudarabah returns), commodity benchmarks, and asset-backed yield proxies — have all failed in practice. Iqbal & Molyneux [31] document that even in the most advanced Islamic finance jurisdictions (Malaysia, Bahrain, UAE), rates charged on Islamic products track LIBOR/SOFR closely, with the interest content merely re-labelled. This persistent dependence — Islamic instruments priced against conventional rates even when *riba* is explicitly prohibited — is widely documented in the sukuk-pricing literature [46], which notes that analysts use LIBOR or the Islamic interbank benchmark rate as the ad-hoc benchmark to evaluate sukuk despite the Shariah objection.

The failure has a structural cause: *no prior work has derived a complete, no-arbitrage-grounded alternative to r from first principles.*

This paper provides that alternative. We call it the κ -rate: the convergence intensity of a perpetual futures contract, proved by Akerer, Hugonnier & Jermann [2] to achieve unique no-arbitrage pricing at $\iota = 0$ (zero interest parameter). The κ -rate is not a label change. It is a mathematically distinct pricing parameter with provably no *riba* content, capable of performing all the monetary functions currently assigned to r in an Islamic economy.

The Gap in One Sentence

Wicksell [49] proposed in 1898 that the “natural rate” of interest should be grounded in real economic productivity rather than monetary convention — but it remained an interest rate with time-preference content. Chapra [20] proposed in 1985 that Islamic monetary policy needs a real-activity-based pricing signal — but never provided a mathematical formula. The κ -rate realizes the Wicksellian aspiration — a real-activity-grounded rate — in a form that is rigorous, *riba*-free, and on-chain observable, 128 years after the question was first posed.

Structure

Section 2 decomposes r into components and identifies the *riba* component. Section 3 defines the κ -rate from first principles. Section 4 contrasts κ with r across all standard properties. Section 5 constructs the κ -yield curve. Section 6 develops applications: credit instruments, takaful, monetary policy, and sovereign benchmarking. Section 7 presents empirical evidence across three independent domains: sovereign sukuk markets, agricultural insurance, and the dYdX v4 perpetual exchange. Section 8 provides the Shariah analysis. Section 9 compares κ with prior proposals and discusses limitations.

2 Decomposing r : What Is the Riba Component?

2.1 The Fisher Decomposition

Fisher [27] established the modern theory of interest as an equilibrium between two forces:

1. **Time preference (impatience):** the subjective preference of individuals for present over future consumption, independent of any investment opportunity.
2. **Investment opportunity:** the marginal productivity of capital — the real return from deploying resources in production.

The equilibrium interest rate r equates marginal impatience to marginal productivity. Fisher’s formulation also yields the famous Fisher equation:

$$r_{\text{nominal}} = r_{\text{real}} + \pi \quad (1)$$

where π is expected inflation. The market rate therefore contains three components: real productivity, time preference, and inflation compensation.

2.2 Böhm-Bawerk’s Distinction

Böhm-Bawerk [19] distinguished two grounds for interest:

- The **productivity ground:** present capital generates future output exceeding the present input (roundabout production). This is the return to real productive activity.

- The **time preference ground**: present goods are valued more than future goods by an “agio” (premium) for their presentness — independent of productivity.

Islamic economics accepts the productivity ground (profit from real activity is halal) but rejects the time preference ground (a predetermined premium for the passage of time, payable on a money loan regardless of economic outcome, is riba). This is the precise jurisprudential boundary.

2.3 The Riba Component

Under the Fisher-Böhm-Bawerk decomposition, the risk-free interest rate r can be written:

$$r = \underbrace{r_{\text{prod}}}_{\text{halal}} + \underbrace{r_{\text{pref}}}_{\text{riba}} + \underbrace{\pi}_{\text{neutral}} \quad (2)$$

The riba component r_{pref} is the pure time preference premium: the predetermined, fixed return on a money loan attributable solely to the passage of time. It is independent of:

- the real economic outcome of the borrower,
- the productive use of the capital,
- market conditions at the time of repayment.

Under El-Gamal’s [25] operationalization of riba as “predetermined excess in exchange, independent of economic outcome,” r_{pref} is riba by definition. The classical Islamic foundation of this identification — money as a measure that may not itself bear a time charge, after al-Ghazālī, Ibn Taymiyyah, Ibn Khaldūn and al-Maqrīzī — is developed in the companion study [11].

2.4 The Wicksell Natural Rate and Its Limit

Wicksell [49] proposed the “natural rate” of interest r^* — the rate at which investment equals saving in a real economy, determined by the marginal productivity of capital without monetary distortion:

$$r^* = \text{MPK} = \frac{\partial F(K, L)}{\partial K} \quad (3)$$

This is closer to the Islamic ideal: a rate grounded in real activity rather than monetary policy. But Wicksell’s natural rate still includes time preference: it is still an interest rate that compensates capital for time elapsed, just one calibrated to real productivity. The riba component persists.

Remark 1. The κ -rate achieves what Wicksell attempted but could not complete. It eliminates r_{pref} entirely from the pricing formula, leaving only the productivity-equivalent component in the form of convergence dynamics. In this structural sense — developed further in the Discussion — the κ -rate plays the role of Wicksell’s natural rate with the riba removed, though it is not numerically identical to the marginal product of capital.

Structure versus magnitude (a caveat carried throughout). The decomposition (2) is conventional theory’s accounting of r ; the κ -rate’s riba-free property is a statement about its *generating process*—identified from the stopping intensity of a real claim, with no predetermined time-preference term—not about its numerical magnitude. A calibrated $\hat{\kappa}$ extracted from markets that still price r inherits those magnitudes, so κ and r may coincide numerically; the companion yield-curve study’s benchmark-free estimator reads the entire observed yield as event intensity for exactly this reason. What makes κ riba-free is therefore structural and outcome-contingent, and a real-only *magnitude* is realized only as markets come to price κ directly—a post-mainnet property, not a present one.

3 The κ -Rate: Definition and Properties

3.1 From Perpetual Futures to Monetary Rate

We recall the foundational result. Akerer, Hugonnier & Jermann [2] prove that the unique no-arbitrage price of a perpetual futures contract on a spot asset X_t (under constant parameters) is:

$$f_t = \frac{\kappa - \iota}{\kappa + r_b - r_a} \cdot X_t \quad (4)$$

where $\kappa > 0$ is the **convergence intensity**, $\iota \geq 0$ is the interest parameter, and r_a, r_b are the long and short funding rates. Theorem 3 of Akerer et al. proves that $\iota = 0$ satisfies no-arbitrage.

κ is, primarily, the **intensity of a stopping time**. The economic content of κ is not the price ratio in (4) but the object that ratio reflects: κ is the intensity of the random time $\theta_t \sim t + \text{Exp}(\kappa)$ at which the contract resolves to its underlying. This is the same mathematical role played by the hazard rate λ in reduced-form credit models (Section 3.5). Defining κ as a stopping-time intensity — rather than as a particular price ratio — is what allows the single parameter to govern perpetuals, sukuk, takaful, and credit protection alike, and it is what we calibrate to data.

3.2 Identifying κ from Data: Three Channels

Because κ is an intensity rather than a quoted price, it is *identified* from whatever observable a given market provides. Three channels arise, each appropriate to a different instrument class and data environment. They estimate the same underlying object through different observables.

Channel 1 — Credit-spread identification (hazard channel). For credit instruments (sovereign sukuk, iCDS, takaful), κ is the loss-event intensity and is identified from the observed spread s and recovery rate δ by

the standard reduced-form relation

$$\hat{\kappa} = \frac{s}{1 - \delta} \quad (5)$$

This is the identification used throughout the empirical κ -yield-curve estimation [5] (Section 7.1), and it follows directly from Corollary 1: at $r = 0$, the Duffie–Singleton survival price reduces to $\exp(-\kappa(1 - \delta)(T - t))$, so the spread compensates for $\kappa(1 - \delta)$ and (5) inverts it. No interest rate enters.

Channel 2 — Perpetual basis identification (nonzero-basis regime). When a perpetual futures market exhibits a nonzero basis ($r_a \neq r_b$, so $f_t \neq X_t$), inverting (4) at $\iota = 0$ identifies κ from the two observed prices:

$$\hat{\kappa}_t = \frac{f_t/X_t \cdot (r_b - r_a)}{1 - f_t/X_t}, \quad r_a \neq r_b. \quad (6)$$

This channel is available only when the basis is nonzero; it is the natural identification for markets that quote asymmetric funding.

Channel 3 — Funding-decay identification (zero-basis regime). A perpetual market may run symmetric funding ($r_a = r_b$), in which case (4) gives the parity $f_t = X_t$ and the basis formula (6) becomes indeterminate (0/0): with no basis, there is no price spread to invert. Here κ is instead identified from the *dynamics* of the funding rate. Modeling the basis $b_t = f_t - X_t$ as a mean-reverting (Ornstein–Uhlenbeck) process with reversion speed κ , the conditional mean decays at that rate,

$$\mathbb{E}[f_{t+\Delta} - X_{t+\Delta} | \mathcal{F}_t] = e^{-\kappa \Delta} (f_t - X_t), \quad (7)$$

so κ is recovered as the estimated mean-reversion rate (equivalently, from the half-life $\ln 2/\kappa$ of basis deviations) of the realized funding series. This is the channel relevant to symmetric-funding venues such as dYdX v4 (Section 7.2).

Remark 2 (Why three channels and not a contradiction). The parity $f_t = X_t$ (zero basis) and the basis inversion (6) are not in conflict: they describe different regimes. When the basis is nonzero, κ is read from its size (Channel 2); when the basis is zero, κ is read from how fast deviations decay (Channel 3). The two perpetual channels are limits of one another — as the basis shrinks to zero, static inversion hands off to dynamic decay — and both recover the same stopping-time intensity that Channel 1 obtains from credit spreads. The constant- κ exposition here is extended to stochastic κ_t in Appendix A.

3.3 Formal Definition

Definition 1 (κ -Rate). The κ -rate for an instrument linking a financial claim to an underlying real asset X_t is the convergence intensity $\kappa > 0$ that:

1. Parameterizes the distribution of the stopping time $\theta_t \sim t + \text{Exp}(\kappa)$ at which the claim is settled.

2. Is identified endogenously from market data via the appropriate channel of Section 3.2 (credit spread, perpetual basis, or funding decay), with no interest rate as an input.
3. Satisfies $\iota = 0$: no interest parameter is charged.
4. Achieves unique no-arbitrage pricing by Theorem 3 of Akerer et al. [2] and, for credit instruments, by Corollary 1, in the martingale sense of Harrison & Kreps [28]: pricing proceeds under an equivalent measure with no reference to a time-preference discount.

3.4 Economic Interpretation of κ

κ measures the **speed of convergence** between a financial instrument’s price and its underlying real asset value. The interpretations below are clearest in the nonzero-basis regime (Channel 2), where the basis is observable; in the zero-basis regime they are read equivalently from the decay rate (7):

- **High κ :** Rapid mean-reversion. The perpetual price tracks spot closely. Short expected stopping time $\mathbb{E}[\theta_t - t] = 1/\kappa$. The market is confident about near-term convergence.
- **Low κ :** Slow convergence. Wide basis. Longer expected stopping time. The market perceives structural uncertainty about when the real asset will reach its fair value.
- $\kappa \rightarrow \infty$: Instantaneous convergence; deviations decay immediately and $f_t = X_t$. Perfect tracking.
- $\kappa \rightarrow 0$: Deviations do not decay; the perpetual becomes untethered from the underlying and the instrument loses economic meaning.

This interpretation is fundamentally different from r :

Property	r (interest rate)	κ (κ -rate)
Determined by	Central bank	Market prices
Measures	Time preference	Convergence speed
Riba content	Yes ($r_{\text{pref}} > 0$)	No
Real basis	Partial	Complete
Negative values	Possible ($r < 0$)	Impossible ($\kappa > 0$)
Observable	Policy setting	On-chain, real-time

3.5 The Duffie–Singleton Isomorphism

The Akerer framework prices perpetual contracts, but the κ -rate’s applicability extends much further. The bridge is a formal isomorphism between the Akerer perpetual pricing framework and the Duffie–Singleton [23] reduced-form credit risk model, which we now establish. Credit-risk pricing has two classical strands: the *structural* approach, which models default as the firm’s asset value crossing a boundary (Merton [41]), and the *reduced-form* approach, which treats default as the first jump of a point process with an exogenous hazard rate (Jarrow & Turnbull [34], Lando [40], Duffie & Singleton [23]). Both strands embed the risk-free rate r as a discount factor; our isomorphism shows that, in the reduced-form case,

this r is the discount component that the riba-free case ($\iota = 0$) eliminates.

In the Duffie–Singleton model, a defaultable bond is priced as:

$$P(t, T) = \mathbb{E}_t \left[\exp \left(- \int_t^T (r_s + \lambda_s(1 - \delta_s)) ds \right) \right] \quad (8)$$

where λ is the default hazard rate, δ is the recovery rate, and r is the risk-free rate. Default occurs at the random stopping time τ with intensity λ .

Proposition 1 (Isomorphism). *The Akerer perpetual pricing framework and the Duffie–Singleton reduced-form credit framework are formally isomorphic under the identification:*

<i>Akerer (perpetual)</i>	$\leftrightarrow$	<i>Duffie–Singleton (credit)</i>
Convergence intensity κ	$\leftrightarrow$	Default hazard rate λ
Interest parameter ι	$\leftrightarrow$	Risk-free rate r
Stopping time θ_t	$\leftrightarrow$	Default time τ
Perpetual contract	$\leftrightarrow$	Defaultable bond

Proof sketch. Both frameworks model an exponential stopping time and a conditional payoff at that time. In the Akerer framework, the stopping time θ_t has conditional density $q(\theta_t \in ds \mid \mathcal{F}_t) = (\kappa_s - \iota_s) \exp(-\int_t^s (\kappa_u - \iota_u) du) ds$. In the Duffie–Singleton framework, the default time τ has hazard rate λ and conditional survival probability $\exp(-\int_t^T \lambda_s ds)$. Under the bijection $\kappa \leftrightarrow \lambda$, $\iota \leftrightarrow r$, $\theta_t \leftrightarrow \tau$, the survival/stopping structure and the no-arbitrage condition of one framework map term-by-term to the other: any pricing equation valid in one yields a valid pricing equation in the other under substitution. The recovery rate δ enters only in the payoff at the stopping time (a defaulted bond pays δ ; an iCDS pays $1 - \delta$), so the credit spread is $s = \kappa(1 - \delta)$ while $\kappa \equiv \lambda$ is the intensity itself — the relation inverted in the empirical identification (5). The detailed proof is given in Ahmed & Bhuyan [4]. $\square$

Corollary 1 ($r = 0$ in Credit Pricing). *No-arbitrage at $\iota = 0$ in the Akerer framework implies no-arbitrage at $r = 0$ in the Duffie–Singleton credit framework. The stopping time distribution alone is sufficient to carry timing information for fair credit pricing.*

The corollary is the operative result. Theorem 3 of Akerer et al. proves no-arbitrage at $\iota = 0$ for perpetuials; by the isomorphism, no-arbitrage at $r = 0$ holds for credit instruments. This is the formal bridge from the κ -rate as a perpetual pricing parameter to the κ -rate as a complete riba-free foundation for sukuk, takaful, and credit derivatives. All of the applications developed in Section 6 rest on Proposition 1 and Corollary 1.

Remark 3 (The Separation Theorem). The conjunction of Theorem 3 of Akerer et al. and Proposition 1 yields what may be called the *Separation Theorem*: at $\iota = 0$

the no-arbitrage price carries no interest term, so interest is an eliminable component of pricing wherever contract resolution occurs at a random stopping time — the stopping-time distribution then carries all timing information previously attributed to the discount factor. Away from $\iota = 0$ the parameter is not a separable discount: in the AHJ random-time representation it enters the stopping intensity itself, $\theta \sim \text{Exp}(\kappa - \iota)$, so the theorem asserts that the riba-free point is well-defined and arbitrage-free, not that interest factors out for all ι . The empirical corroboration is the European negative-rate environment of 2014–2022, during which credit markets priced risk efficiently at $r \leq 0$ — confirming that the discount factor and the no-arbitrage condition are not coextensive.

4 κ vs. r : A Complete Comparison

4.1 The Riba Test

El-Gamal’s [25] criterion for riba: a payment is riba if it is (1) predetermined, (2) in excess of principal, and (3) independent of real economic outcome.

Proposition 2 (The κ -Rate is Not Riba). *The κ -rate satisfies none of the three riba criteria:*

1. **Not predetermined.** κ_t fluctuates continuously with market conditions. It is implied from market prices, not contractually fixed.
2. **Not excess of principal.** The κ -rate does not generate a return over principal. At $\iota = 0$ and $r_a = r_b$, the instrument price equals the underlying asset value ($f_t = X_t$) — no excess.
3. **Not independent of economic outcome.** κ is backed out from the observed basis f_t/X_t , which reflects real market views about asset convergence. If the underlying asset performs well, the basis narrows (high κ); if poorly, the basis widens (low κ). The κ -rate is tied to real economic dynamics.

Proof. (1) follows from the identification channels of Section 3.2: $\hat{\kappa}_t$ is a continuous function of market observables (prices or funding dynamics), which vary. (2) follows from the parity established in Section 3.2: at $\iota = 0$ and $r_a = r_b$, equation (4) gives $f_t = X_t$. (3) follows from the market-determination of $\hat{\kappa}_t$: it is identified from observed data, not set exogenously. $\square$

Remark 4 (Scope of criterion (2)). The parity argument for criterion (2) ($f_t = X_t$, no excess) is the zero-credit-risk case ($\iota = 0$, $r_a = r_b$). For a credit instrument the fair spread $\kappa(1 - \delta) > 0$ is not an excess over principal but actuarial compensation for an expected loss—Paper 9’s cost-of-capital identity decomposes the required return as a productivity/time-value term plus exactly $\kappa(1 - \delta)$. Because El-Gamal’s criteria are conjunctive (a payment is riba only if it is predetermined *and* in excess *and* outcome-independent), criteria (1) and (3)—which hold for credit

instruments without qualification—are each individually sufficient to defeat the charge. The parity leg is therefore redundant for the credit case, not load-bearing: the outward-facing defense leads with outcome-contingency (criteria 1 and 3) and the portfolio-level rebuttal of the companion safe-asset study, never with parity.

4.2 Replacing r in Standard Pricing Formulas

The key question is whether κ can perform the mathematical functions of r . We demonstrate this across three standard settings.

Bond pricing. A conventional zero-coupon bond: $P = e^{-rT}$. The κ -equivalent (Islamic perpetual sukuk with expected maturity $T = 1/\kappa$):

$$P^\kappa = \mathbb{E}^Q[X_\tau], \quad \tau \sim \text{Exp}(\kappa) \quad (9)$$

(equal to X_t in the martingale case). The discount comes not from e^{-rT} but from the stopping time distribution: the probability that the redemption event is more than T periods away decays exponentially at rate κ , achieving the same economic effect without any interest.

Option pricing. In Black-Scholes [18], the risk-free rate r appears as the discount factor and in the drift of the risk-neutral measure. At $\iota = 0$ and $r_a = r_b$, the everlasting put from Akerer et al. Proposition 6 (for $x \geq K$) is

$$\Pi(x, K) = \frac{K^{1-\beta_p}}{\beta_c - \beta_p} \cdot x^{\beta_p}, \quad \beta_{p,c} = \frac{1}{2} \mp \sqrt{\frac{1}{4} + \frac{2\kappa}{\sigma^2}} \quad (10)$$

(with the intrinsic value $K - x$ added where $x < K$), replacing r with κ entirely in the exponent.

Credit default swap. A conventional CDS spread satisfies $s \approx \lambda(1 - \delta)$, the hazard rate times loss-given-default. The iCDS spread [7] is $s^* = \kappa(1 - \delta)$ with $\kappa \equiv \lambda$ and $\iota = 0$, achieving the same risk transfer without any interest baseline (Section 6.2).

4.3 Policy Functions of r and Their κ Analogs

Table 1 lists the monetary policy functions of r and the corresponding κ -rate mechanism.

5 The κ -Yield Curve

5.1 Construction

The conventional yield curve $r(T)$ maps maturity T to the yield on a zero-coupon risk-free bond. It is the term structure of interest rates, grounded in the expectation hypothesis and risk premia over time. Cox, Ingersoll & Ross [21] provide the canonical equilibrium derivation; Vasicek [48] provides the Gaussian version; Nelson & Siegel [43] provide the empirical estimation framework.

Table 1: Monetary functions: r vs. κ .

Function	r mechanism	κ mechanism
Price of credit risk	$r + \lambda$ spread	$\kappa \equiv \lambda$ at $\iota = 0$
Discount factor	e^{-rT}	$\text{Exp}(\kappa)$ survival prob.
Term structure	Yield curve $r(T)$	κ -term-structure (upward)
Monetary signal	Central bank sets r	Market implies $\hat{\kappa}$
Sovereign benchmark	Government bond yield	Perpetual sukuk implied κ
Insurance pricing	$e^{-r} \cdot \text{claim}$	$\pi^* = \kappa B$

The **κ -yield curve** is the Islamic analog: the mapping $\kappa(T)$ of convergence intensities from perpetual instruments of different expected horizons.

Proposition 3 (Horizon–Intensity Identity). *For a single instrument whose stopping event has expected horizon T (the expected time until resolution under no other information), the convergence intensity and the horizon are reciprocal:*

$$\kappa(T) = \frac{1}{T}, \quad T > 0. \quad (11)$$

A short expected horizon corresponds to a high intensity (rapid convergence); a long expected horizon to a low intensity.

Proof. For a stopping time $\theta_t \sim \text{Exp}(\kappa)$, the expected stopping time is $\mathbb{E}[\theta_t - t] = 1/\kappa$. Setting this equal to the instrument’s expected horizon T gives $\kappa = 1/T$. $\square$

Remark 5 (The identity is not the term structure). Equation (11) relates one instrument’s intensity to its own expected horizon; it is a definitional reciprocal, not a claim about how κ varies across the maturity spectrum of a given issuer. The latter object — the empirical κ -term-structure for a fixed credit, estimated from sukuk of different maturities in Section 7.1 — is *upward-sloping*: longer-dated claims on the same sovereign carry higher cumulative credit uncertainty and hence higher κ . There is no contradiction: (11) maps horizon to intensity for a single contract, whereas the empirical term structure maps maturity to intensity across contracts on one issuer. We reserve “ κ -yield curve” for the latter, empirically estimated, object in what follows.

5.2 Comparison with the Conventional Yield Curve

Table 2 contrasts the two term structures.

5.3 Empirical Estimation

By Proposition 3, to estimate the κ -yield curve one needs:

Table 2: κ -yield curve vs. conventional yield curve.

Property	CIR/Vasicek	κ -curve
Pricing parameter	Short rate r_t	Conv. intensity κ_t
Riba content	Yes	No
Central bank role	Sets r	None (market-implied)
Shape	Any	Upward-sloping (empirical)
Calibration	Rate expectations	Credit spreads / basis
Negative rates	Possible	Impossible
DeFi observable	No	Yes (on-chain)

1. Observed perpetuals basis $f_t^{(i)}/X_t^{(i)}$ for instruments $i = 1, \dots, n$ with different underlying asset horizons.
2. Implied $\hat{\kappa}_t^{(i)}$ from the basis channel (6) where the basis is nonzero, or from the funding-decay channel (7) for symmetric-funding instruments.
3. The Nelson-Siegel [43] functional form applied to the $(\hat{\kappa}^{(i)}, T^{(i)})$ pairs to produce a smooth curve.

For a single-instrument perpetual market (e.g. a BTC-USD perpetual), the spot $\hat{\kappa}_t$ is directly observable but a full term structure requires perpetuals on multiple underlyings with different expected lifetimes (commodity contracts, infrastructure sukuk, mortgage-equivalent instruments). For sovereign credit, by contrast, the term structure is directly estimable from the existing maturity spectrum of issued sukuk (Section 7.1).

6 Applications of the κ -Rate

6.1 What κ Can and Cannot Price

The κ -rate does not price every asset; it prices a specific category, and that boundary is jurisprudentially meaningful.

What κ directly prices. κ prices instruments whose value depends on a random stopping time — a future event whose timing is uncertain but whose intensity can be measured. The four instrument families developed in this section share this structural property: each contract’s value resolves at a random time θ , and κ parameterizes the intensity of that resolution. Perpetual contracts converge at intensity κ . Sukuk redeem at intensity κ . Takaful contracts pay benefits at loss-event intensity κ . Credit default swaps trigger at default intensity κ . In each case, κ enters the pricing formula not as a discount rate but as the intensity of the random delivery time.

What κ does not directly price. Spot assets are not priced using κ . A spot BTC at \$95,000, a Dubai apartment, a tonne of rice — these are priced by supply and demand in their respective markets. κ enters only when a contract is written *on* the asset that pays off at a random future time. The underlying asset itself does

not need κ ; derivatives, credit instruments, and insurance contracts on the asset do.

The structural distinction. κ does not replace market prices for spot assets. It replaces the discount rate in any pricing formula that involves expected future values. Wherever conventional finance uses $r > 0$ to discount future cash flows to present value, κ replaces it — but at a different place in the formula. The risk-free rate r appears as a discount factor that reduces the present value of future cash flows. κ appears as the intensity of the random time when the payoff is delivered. The Separation Theorem (Remark 3) proves these are distinct mathematical roles, and that only the latter is structurally necessary.

The prohibition as an exclusion theorem. Instruments with predetermined returns regardless of outcome — the category Islam prohibits as *riba* — cannot be priced using κ , because they have no random resolution structure to which κ can attach. The converse does not hold: κ -priceability does not by itself imply permissibility, since *gharar*, *maysir*, and possession (*qabd*) are properties of the contract built on the rate, not of the rate itself (Section 8). The prohibition’s mathematical content is thus an exclusion theorem for *riba*, not a compliance certificate for every instrument the framework can price.

6.2 Islamic Credit Instruments

The κ -rate provides the pricing foundation for a complete *riba*-free credit market. The formal proof (the isomorphism between the Akerer framework and the Duffie-Singleton credit model) is established in Ahmed & Bhuyan [4]. We summarize the pricing formulas.

Perpetual sukuk. Consider an Islamic investment certificate backed by a real asset with value X_t , with random redemption at stopping time τ of intensity κ . By the isomorphism (Proposition 1) at $\iota = 0$, the no-arbitrage price is the expected asset value at redemption,

$$P_t^{\text{sukuk}} = \mathbb{E}^Q[X_\tau | \mathcal{F}_t], \quad (12)$$

and the certificate’s redemption-survival factor $S(0, T) = \mathbb{E}^Q[\exp(-\int_0^T \kappa_s ds)]$ — the probability that redemption has not occurred by horizon T , the κ -analog of the discount factor — admits under CIR- κ dynamics the affine closed form $S(0, T) = A(T) e^{-B(T)\kappa_0}$ derived in Appendix A (with A, B the CIR functions). In the special case of a martingale underlying and symmetric funding ($r_a = r_b$), $\mathbb{E}^Q[X_\tau] = X_t$ and the certificate trades at the fair asset value with no *riba*. The κ -calibration uses the asset’s expected redemption horizon, $\kappa \approx 1/T$, or the credit channel (5) where a sovereign spread is observable. Returns to the holder are variable and asset-linked (*ijarah/mudarabah*), not a predetermined coupon — this is a risk-sharing certificate, not fixed income.

iCDS (Islamic Credit Default Swap). Structured as a takaful-style mutual guarantee [7], the protection buyer is a *tabarru* contributor and the fair protection

spread at $\iota = 0$ is

$$s^* = \kappa(1 - \delta), \quad (13)$$

where δ is the recovery rate. The *par running spread* is independent of the discount rate: because both legs terminate at default the discount factor cancels, so a correctly-priced conventional CDS quotes the same fair spread $s_{\text{conv}} = \kappa(1 - \delta) = s^*$ at every maturity. The par spread is thus *riba-neutral*; the *riba* is confined to the present value, where the discounting haircut on the future loss is $\Pi_{\text{riba}}^{\text{PV}} = (1 - \delta)\iota/(\kappa + \iota)$, vanishing at $\iota = 0$ and increasing in ι [7]. The iCDS κ is calibrated from the reference entity's default intensity via the credit channel (5).

Removing ι removes the price of time, but the market spread that (5) reads embeds a risk premium and so measures the *priced* intensity κ^Q , not the *realised* default hazard κ^P ; whether charging that premium is permissible once every predetermined charge is gone is the substantive question the corpus reserves. Because the buyer is already a *tabarru* contributor, the substance-compliant form does not raise it: the cooperative pool sets the contribution at the realised hazard,

$$s_{\text{coop}}^* = \kappa^P(1 - \delta), \quad (14)$$

the actuarially fair expected loss — exactly as the takaful contribution of Section 6.3 is set — and returns the wedge $(\kappa^Q - \kappa^P)(1 - \delta)$ to members as surplus rather than selling it to a protection writer, net of a disclosed *wakala* administration fee. So priced, the iCDS is structurally the takaful case and inherits its soundness on form and substance, gated on insurable interest and possession of the reference obligation; the residual narrows to whether tokenised title is valid *qabd* and whether the *wakala* loading is a genuine fee (Section 8).

Sovereign perpetual sukuk. A government issues sukuk backed by an infrastructure project (expected completion T years, value X_t). The secondary-market price follows (12); the sovereign's credit-specific κ is read directly from its sukuk spread via (5), with no SOFR benchmark required — exactly the construction estimated empirically on the real sukuk cross-section of Section 7.1.

6.3 Takaful (Islamic Insurance)

Billah [17] identifies the core takaful pricing tension: actuarial fairness requires discounting expected claims, but standard discounting uses r (*riba*). The κ -rate resolves this.

For a takaful contract paying benefit B upon a covered loss event of hazard intensity κ , the actuarially fair contribution at $\iota = 0$ is simply

$$\pi^* = \kappa B, \quad (15)$$

the hazard intensity times the benefit, with no interest discount. Under CIR- κ dynamics this extends to the affine

multi-period contribution $\pi_{\text{stoch}}^*(T) = \mathbb{E}_0[\kappa_T]B$ derived in [6]. The κ -calibration uses the loss-frequency channel (historical claim frequency = κ). Formula (15) is the *riba-free* actuarial premium examined empirically in [6], where the model-implied $\hat{\kappa} = 12.06\%$ for India is numerically close to the PMFBY national crop-insurance actuarial rate ($\approx 12\%$) — a units-level coincidence between a loss frequency and a premium rate, not a strict validation (the same $\hat{\kappa}$ for Pakistan misses its observed rate by 1.6–3.4 $\times$). It resolves the benchmark-circularity problem in conventional takaful contribution pricing [29]. The *riba* loading in conventional insurance — the premium reduction from discounting liabilities at r — is $\Pi_{\text{riba}} = \kappa Br/(1+r)$, equal to 7.4% of the *riba-free* premium at $r = 8\%$.

6.4 Islamic Monetary Policy

Chapra [20] called for a *riba-free* monetary policy instrument. The κ -rate provides it.

κ -based monetary signaling. An Islamic central bank publishes the system-wide average $\bar{\kappa}$ — the weighted mean of implied κ -rates across all Islamic perpetual instruments in the market:

$$\bar{\kappa}_t = \frac{\sum_i w_i \hat{\kappa}_t^{(i)}}{\sum_i w_i} \quad (16)$$

where weights w_i are notional outstanding. Because the aggregate is dominated by credit instruments (*sukuk*), the relevant reading is the hazard interpretation, in which κ rises with credit risk and resolution uncertainty:

- High $\bar{\kappa}$: elevated system-wide credit intensity — wider *sukuk* spreads, greater perceived resolution risk. Tight or stressed conditions (analogous to high r).
- Low $\bar{\kappa}$: compressed credit intensity — narrow spreads, low perceived risk. Easy conditions (analogous to low r).

This sign convention coincides with the empirical ordering of Section 7.1, where distressed Bahrain carries the highest κ and the safest sovereigns the lowest.

Open market operations in κ -space. An Islamic central bank can intervene in the perpetuums market to influence $\bar{\kappa}$ — providing a policy transmission mechanism that does not involve setting an interest rate. This is the *riba-free* analog of open market operations.

Reserve requirements. If the κ -rate of Islamic bank assets (implied from their *sukuk* portfolios) falls below a threshold, the central bank can require higher reserve ratios — a prudential tool that reinforces the κ -based monetary signal.

6.5 Sovereign Benchmark Rate

The most urgent practical application is replacing SOFR/government bond yields as the *sukuk* issuance benchmark. The *sukuk*-pricing literature [46] documents that major sovereign *sukuk* (Malaysia GII, Saudi Arabia *sukuk*, Indonesia SR series) are benchmarked at issuance

against conventional reference rates, making the *riba* implicit.

The κ -yield curve replaces this benchmark. For a given sovereign, the issuance benchmark at each maturity τ is read directly from the credit channel (5) applied to that sovereign’s own sukuk spreads,

$$\kappa_{\text{sov}}(\tau) = \frac{s_{\text{sov}}(\tau)}{1 - \delta}, \quad (17)$$

yielding the upward-sloping, country-specific κ -term-structure estimated in Section 7.1. The secondary-market yield is then $y(\tau) = \kappa_{\text{sov}}(\tau)$ — a *riba*-free sovereign benchmark built entirely from market-observed credit intensities, with no SOFR anchor. This is distinct from the single-instrument horizon identity $\kappa = 1/T$ (Remark 5), which sets only the expected redemption scale of one contract, not the issuer’s term structure.

Table 3 summarizes κ calibration by instrument type.

Table 3: κ -rate calibration by instrument.

Instrument	κ calibration	Data source
Perp. futures	Basis / funding decay	Exchange data
Sukuk	Credit spread $s/(1 - \delta)$	Sukuk market data
Takaful	Claim frequency	Actuarial tables
iCDS	Default history	Rating agencies
Sovereign sukuk	Credit spread $s/(1 - \delta)$	Sovereign sukuk spreads

7 Empirical Evidence

We present empirical evidence from three independent domains. The first is the construction of an empirical κ -yield curve from a real 195-instrument USD-sukuk cross-section (40 sovereign, 154 corporate, 1 municipal) and five-issuer constant-maturity dynamics. The second is the calibration of κ against agricultural insurance markets, where the model-implied rate matches an independently-set national actuarial rate. The third is the dYdX v4 perpetual exchange, which provides a commercial-scale existence proof for the $\iota = 0$ mechanism. The three exercises are complementary and causally independent: they draw on sovereign credit, physical loss frequencies, and crypto market microstructure respectively, with no shared data source. Section 7.3 reconciles what recurs across them.

7.1 The κ -Yield Curve on Real Traded Sukuk

Following Ahmed & Bhuyan [5], the κ -yield curve is estimated *entirely on real traded sukuk*. The primary dataset is a 195-instrument cross-section of actively-quoted USD

sukuk (40 sovereign, 154 corporate, 1 municipal) from the Cbonds Bond Screener, complemented by roll-down-free constant-maturity $\hat{\kappa}$ histories for five issuers (2018–2026). Recovery is set at $\delta = 0.60$ throughout, anchored to the rating agencies’ historical sovereign recovery studies (the sovereign CDS market’s 25% quoting convention differs; recovery sensitivity is a pure rescaling of $\hat{\kappa}$, see [5]).

7.1.1 Identification

We apply the credit-spread channel (5) of Section 3.2 to each observation (i, τ, t) , identifying the convergence intensity from the observed spread and recovery rate:

$$\hat{\kappa}_{i,\tau,t} = \frac{s_{i,\tau,t}}{1 - \delta}. \quad (18)$$

As established there, this follows from Corollary 1: at $r = 0$ the Duffie–Singleton survival price reduces to $\exp(-\kappa(1 - \delta)(T - t))$, so the spread compensates for $\kappa(1 - \delta)$ and (18) inverts it. This is the hazard channel, distinct from the perpetual-market channels (6)–(7); all three recover the same stopping-time intensity as a construct, so estimates are conceptually commensurable across the credit and perpetual domains.

7.1.2 Methodological note: the spread is quoted over a conventional benchmark

The spread s in equation (18) is observed from existing sovereign sukuk markets, which currently quote against SOFR or conventional government bond benchmarks. The benchmark reference is therefore drawn from the conventional interest rate system. The Separation Theorem (Remark 3) provides the theoretical justification for the extraction procedure: the credit *par spread* is $\lambda(1 - \delta)$, *independent of the discount rate* (both legs terminate at default, so the rate cancels), so equation (18) recovers κ directly with no *riba* loading embedded in the spread. The residual transitional concern is that s is observed over a conventional risk-free benchmark rather than a native $\iota = 0$ reference — a quotation convention, not a structural defect, since as Islamic markets develop native pricing mechanisms operating at $\iota = 0$ (of which dYdX v4 is the first commercial-scale example, Section 7.2), the input data becomes progressively cleaner.

7.1.3 Term-structure fitting and CIR dynamics

We fit a Nelson–Siegel–Svensson functional form [43] to the cross-section of maturities for each country-month, then estimate CIR- κ dynamics (equation (19) of Appendix A) by the generalized method of moments (GMM). The Feller condition $2\alpha\bar{\kappa} \geq \sigma_\kappa^2$ is tested in each market.

7.1.4 Results

Three results obtain, all on real traded sukuk.

(1) **Credit-quality ordering.** On the 195-bond cross-section the convergence intensity broadly increases as credit quality declines (Table 4): median $\hat{\kappa}$ runs from 0.93% at AAA to 4.37% at single-B, rising across the spectrum apart from the Saudi-dominated AA bucket, which

Table 4: $\hat{\kappa}$ by credit rating on the real 195-bond USD-sukuk cross-section (Cbonds; $\delta = 0.60$). The five rated buckets ($n = 167$) plus the 28 unrated bonds total the full 195; the snapshot’s active-quote filter leaves no BB or CCC names in this panel (the full Cbonds universe does contain outstanding BB/B/CCC sukuk — e.g. the Bahrain, Egypt and Pakistan sovereigns — which enter the companion study’s sovereign-core and full-universe analyses). $\hat{\kappa}$ increases across the rated buckets apart from the Saudi-dominated AA bucket ($n = 102$), which prices just above single-A; NR is unrated — shown for completeness as the high-spread tail, not as part of the rating ordering.

Rating	n	Median $\hat{\kappa}$ (% p.a.)
AAA	9	0.93
AA	102	2.28
A	47	1.96
BBB	6	3.41
B	3	4.37
NR	28	5.60

prices just above single-A. The framework is not manufacturing numbers: it reads, from the sukuk market, the same credit risk that ratings agencies identify by independent methods.

(2) Term structure. The investment-grade κ -yield curve is upward-sloping: a linear fit of the G -spread on maturity has a positive slope of +1.17 bp/yr (pooled), attenuating to +0.63 bp/yr with rating fixed effects. The slope is positive but modest — a structural term premium with no interest-rate component, read off traded sukuk rather than imposed.

(3) Feller non-negativity. From roll-down-free constant-maturity histories (five issuers, 2018–2026), κ is mean-reverting with bias-corrected half-lives of 1.4–4.3 months at the well-sampled 7-year tenor (four issuers; Türkiye’s near-unit-root window is not separately identified), and the Feller condition $2\alpha\bar{\kappa} \geq \sigma_{\kappa}^2$ holds in every series. By Lemma 1, $\kappa_t > 0$ almost surely — qualitatively unlike the negative-rate experience of 2014–2024, when the euro-area (2014–2022) and Japanese (2016–2024) benchmark and policy rates turned negative and created Shariah-compliance problems for instruments anchored to them. The κ -rate is non-negative by mathematical construction, as a probability intensity must be.

An agricultural takaful coincidence. Ahmed & Bhuyan [6] extend the empirical work to agricultural takaful pricing across South and Southeast Asia. The model-implied $\hat{\kappa} = 12.06\%$ for India is numerically close to the actual Pradhan Mantri Fasal Bima Yojana (PMFBY) national crop insurance actuarial rate ($\approx 12\%$). The model is not fitted to this rate; $\hat{\kappa}$ is calibrated from historical crop-loss data (World Bank cereal yield series, 1980–2023), while the PMFBY rate is set by the Government of India’s national insurance programme operating indepen-

dently. We read this cautiously: it is a units-level coincidence between a loss frequency and a premium rate, not a strict validation — the same calibration misses Pakistan’s observed national rate by 1.6–3.4 $\times$. The firmer empirical support for κ comes from the instrument-level sukuk and CDS evidence below.

7.1.5 Real instrument-level validation

All three results above are estimated entirely on real traded sukuk; the companion study [5] supplies the full instrument-level detail (the 195-instrument USD-sukuk cross-section and the five-issuer constant-maturity dynamics). The credit-quality ordering, the term-premium shape and the non-negativity property are thus properties of the traded sukuk market, so the credit-domain leg of the cross-domain argument rests on real, interest-free-identified κ , not a proxy.

7.2 Independent Market Validation: dYdX v4

The empirical work in Section 7.1 establishes that κ measures real sovereign credit risk in publicly observable markets. A complementary question is whether the $\iota = 0$ perpetual mechanism itself — on which the isomorphism in Proposition 1 ultimately rests — is operationally viable at commercial scale, or whether it remains a theoretical construct.

The answer is given by a third-party protocol that was not designed for Islamic finance: dYdX v4. Launched in October 2023 as a Cosmos appchain, dYdX v4 set the interest parameter component of its perpetual funding rate to zero by default, computing funding entirely from the observed basis between mark and index prices — the mechanism specified by Theorem 3 of Akerer et al. The protocol has operated continuously since then at commercial scale. As of May 2026 its public indexer and DeFiLlama report ≈ 300 listed perpetual markets, $\approx \$136\text{M}$ in total value locked, $\approx \$55\text{M}$ open interest, and $\approx \$1.5\text{B}$ in monthly trading volume — all with funding computed at $\iota = 0$ [24, 22] (volumes were materially higher at the 2024 market peak). The point is not the headline figure but its persistence: an $\iota = 0$ funding regime has run live on mainnet for over two years at open interest in the tens to hundreds of millions of dollars (protocol TVL peaked at $\approx \$416\text{M}$ in December 2024; $\approx \$55\text{M}$ open interest at the May-2026 snapshot) and monthly volumes in the billions, handling adversarial trading at scale with no convergence pathologies. This is an independent empirical existence proof for the mechanism — the $\iota = 0$ regime is not hypothetical.

Two empirical signatures are particularly relevant.

(1) Funding rate distribution. Empirical funding rate distributions on dYdX v4 are approximately symmetric around zero [3], consistent with the prediction of equation (6) at $\iota = 0$. By contrast, centralized exchanges that retain an interest floor ($\iota \approx 0.01\%$ per 8 hours)

exhibit a positive-skewed distribution with a structural lower bound at the interest floor. The two distributions are visually and statistically distinguishable.

(2) **Implied $\hat{\kappa}$ from funding rate dynamics.** Because dYdX v4 runs symmetric funding ($r_a = r_b$), the static basis formula (6) is indeterminate; the appropriate estimator is the funding-decay channel (7), which recovers κ from the mean-reversion rate (half-life) of the basis-deviation series. Estimated on two years of dYdX v4 funding and, as a cleaner cross-check, on the *unclamped* Binance premium-index basis (AR(1)/OU fits for BTC, ETH, SOL), this yields a *large* convergence intensity — $\hat{\kappa}$ on the order of 10^3 per annum, a basis half-life of a few hours. This magnitude is robust: re-estimated on first/second-half and low/high-volatility subsamples it stays in the 10^3 – 10^4 range, never approaching the credit-domain value. That is the expected magnitude for a liquid perpetual that tracks spot tightly: this funding-decay $\hat{\kappa}$ is orders of magnitude larger than the credit-domain estimates of Section 7.1 (a few percent per annum) and the takaful calibration ($\hat{\kappa}_{MLE} \in [0.07, 0.15]$) across countries, India 0.12; [6]) — the same *type* of object at a regime-specific level (Section 7.3). What dYdX uniquely establishes is twofold: that the $\iota = 0$ mechanism operates at commercial scale, and that a decoupled, interest-free $\hat{\kappa}$ is directly measurable from a venue charging no interest — a guarantee the SOFR-referenced credit channel cannot itself provide.

Remark 6. We do not claim that dYdX v4 is a Shariah-compliant protocol. It was not designed for that purpose, and dYdX governance vote 220 has subsequently introduced $\iota > 0$ on isolated markets. We cite dYdX v4 strictly as an independent empirical existence proof for the $\iota = 0$ mechanism. The Shariah-compliance question is jurisprudential and is treated separately [3].

7.3 Cross-Domain Validation Summary

The three empirical exercises validate κ from independent sources at different scales:

- **Sovereign and corporate credit markets:** on a real 195-instrument USD-sukuk cross-section (Section 7.1.5), $\hat{\kappa}$ broadly rising as credit quality declines ($\approx 0.9\%$ AAA to $\approx 4.4\%$ single-B, clean IG/HY separation) with an upward investment-grade term structure, and five-issuer constant-maturity dynamics (mean-reverting, 1.4–4.3 month half-lives across the four identified issuers, Feller in every series) — the real credit-domain anchor.
- **Agricultural insurance** (India PMFBY): model-implied $\hat{\kappa} = 12.06\%$ is numerically close to the national actuarial rate ($\approx 12\%$) — a units-level coincidence (loss frequency vs. premium rate), not a strict validation.
- **Crypto perpetual markets** (dYdX v4, mainnet, 2023–present): $\iota = 0$ mechanism operating at commercial scale; a decoupled, interest-free $\hat{\kappa}$ directly measurable, of magnitude $\sim 10^3$ p.a. (hour-scale convergence)

— the same construct as in credit/insurance but regime-specific in level.

Across the three domains κ appears as the same *type* of object — a riba-free stopping-time intensity, identified through different channels (credit spread, loss frequency, funding decay) without an interest input (Figure 1). We do not claim the estimates coincide *numerically*: their magnitudes are regime-specific. The qualitative cross-domain recurrence of the construct supports κ as a general riba-free primitive. The credit domain now also carries a *quantitative* validation on real (non-synthetic, non-proxy) data — the instrument-level sukuk cross-section and constant-maturity dynamics of Section 7.1.5 — and the perpetual domain its live $\iota = 0$ measurement on dYdX; a single real panel spanning all domains at once remains the principal item of future work.




Figure 1: The riba-free intensity κ identified from real data across five independent, interest-free channels spanning the paper’s three domains (credit, insurance, perpetual markets): traded sukuk (SPSK distribution spread), EM/sovereign credit OAS (ICE BofA, FRED), crop-insurance loss-cost (USDA RMA), and crypto perpetual-funding decay (the interest-free premium-index basis on Binance, and dYdX $\iota = 0$ convergence). On a log scale the estimates separate into two regimes spanning roughly four orders of magnitude: a slow, year-scale *credit/insurance* regime ($\hat{\kappa} \approx 0.015$ – 0.2) and a fast, hour-scale *basis-convergence* regime ($\hat{\kappa} \sim 10^2$ – 10^3). κ is the same *type* of stopping-time intensity throughout, but its level is regime-specific; we make no claim of cross-domain numerical equality.

8 Shariah Analysis

8.1 Riba

The κ -rate satisfies Proposition 2: it is not predetermined, not in excess of principal, and not independent of real economic outcomes. Crucially, it satisfies both forms of riba prohibition:

- **Riba al-nasi’ah** (excess for deferment): The κ -rate does not compensate for the passage of time as such. The stopping time distribution $\text{Exp}(\kappa)$ encodes *when* a real economic event occurs; κ is the event intensity, not a time preference parameter.

- **Riba al-fadl** (excess in exchange): At $\iota = 0$, the perpetual futures price equals the underlying asset value ($f_t = X_t$ at $r_a = r_b$). There is no permanent systematic excess in the exchange.

A contemporary jurisprudential discussion [44] is instructive on the precise scope of this claim. Addressing crypto perpetual contracts where the exchange charges no interest and the contract tracks the underlying asset price, the response distinguishes the *riba* question — which the absence of an interest charge addresses directly — from separate objections grounded in *gharar* and the absence of genuine ownership, on which it ultimately judges such contracts impermissible. This distinction is exactly the one we draw: the $\iota = 0$ condition, and the κ -rate built upon it, resolve the *riba* prohibition specifically. They do not by themselves settle the *gharar* and ownership questions, which depend on the contractual structure of the particular instrument rather than on the pricing parameter. The κ -rate is a *riba*-free pricing primitive; whether any given contract employing it is Shariah-compliant in full requires separate analysis of its other contractual features (§9).

8.2 Gharar

The randomness of the stopping time $\tau \sim \text{Exp}(\kappa)$ introduces actuarial uncertainty (the exact time of the event is unknown) but not contractual ambiguity (the distribution is fully specified by κ , which is public and on-chain). El-Gamal [25] and Kamali [36] distinguish these: *gharar* refers to uncertainty about the *existence* of the subject of transaction, not its *timing* when the timing is actuarially determined and transparently disclosed. The κ -rate passes the *gharar* test as a *pricing parameter*: the intensity κ and the resulting stopping-time distribution are public, on-chain, and unambiguous. This is distinct from the question of whether a particular contract *employing* κ — such as a leveraged perpetual — raises *gharar* concerns of its own; that is a property of the contract’s structure, not of the rate, and must be assessed contract by contract.

8.3 Maysir and Qabd

Maysir (gambling). As a pricing parameter, κ does not import maysir: it attaches to instruments backed by a real underlying asset (X_t), every payment corresponds to a real economic event, and the contracts are structured for risk-transfer and price discovery rather than as a zero-sum wager. Three features address the standard maysir objections: the takaful-style mutual structure of the credit instruments (the iCDS protection buyer is a *tabarru* contributor to a cooperative pool, Section 6.2); the symmetric $\iota = 0$ funding, which removes the systematic transfer that would make the contract a standing bet against the protocol; and the real-asset convergence that κ measures. The mechanism-design simulation [8] finds rational traders at $\iota = 0$ equilibrating at sub-maximal leverage, consistent with hedging rather than maximal

speculation. Whether a particular contract employing κ avoids maysir remains, like *gharar*, a contract-level question for scholars.

Qabd (constructive possession). We are careful not to overstate the reach of the *riba* analysis. Contemporary scholarship [44] raises a *qabd* (possession/constructive-delivery) objection to cash-settled synthetic contracts that is logically prior to, and independent of, the *riba* question — indeed the very ruling we rely on for the *riba* concession maintains impermissibility on possession grounds. We do not resolve *qabd* here. We note only that its resolution must run through contractual structure rather than the pricing parameter: through real-world-asset (RWA)-backed tokenisation establishing on-chain custody of the underlying, and through the *istisna/salam* precedents for deferred delivery of a determinate asset. In deployment terms this implies an ordering, not a verdict: instruments in which custody of a real underlying directly addresses the asset-backing and possession objections — e.g. perpetuals on vault-secured tokenised gold, or sukuk-yield hedging contracts for institutional risk management — naturally precede naked cash-settled synthetics, which remain the hard case for scholarly review. Whether a given κ -instrument satisfies the standard is a matter for a recognised Shariah board, and the maysir and *qabd* review of each application warrants the same rigour as the *riba* analysis offered here. This paper’s contribution is confined to the *riba*-free *rate*; it does not by itself establish the full Shariah compliance of any contract built upon it.

8.4 Alignment with AAOIFI Standards (Subject to Board Review)

AAOIFI Standard 17 (Investment Sukuk) [1] requires that sukuk represent proportional ownership in real assets. The κ -parameterized Islamic Perpetual Sukuk is structured to meet this requirement: the holder’s claim is on X_τ (the real asset at redemption); whether a specific issuance satisfies Standard 17 is for certification by a recognised Shariah board. Standard 26 (Takaful) requires actuarially fair, mutual risk-sharing contributions: the everlasting takaful formula (15) is designed to provide this without *riba*. The iCDS (13) is likewise structured as a takaful-style mutual guarantee — the protection buyer is a *tabarru* (donation) contributor to a cooperative pool rather than a counterparty purchasing risk — so that the spread $s^* = \kappa(1 - \delta)$ is an actuarially fair pool contribution, not a wager on default [7]. These are structural design alignments, not compliance certifications; AAOIFI conformity of any concrete contract requires dedicated review (cf. the scope limitation in Section 9).

9 Discussion

9.1 Why Prior Islamic Monetary Proposals Failed

Chapra [20], Khan & Mirakhor [38], and subsequent IRTI researchers identified the need for a riba-free pricing parameter but could not provide one. The profit-sharing rate (mudarabah return) was proposed as a substitute but failed for three structural reasons documented by Iqbal & Molyneux [31]:

1. **Ex-post determination.** Profit-sharing rates are known only after the accounting period, making forward pricing impossible.
2. **Information asymmetry.** Profit manipulation and audit uncertainty mean profit rates are not independently verifiable.
3. **LIBOR convergence.** Competitive pressure caused Islamic bank “profit rates” to track LIBOR, reintroducing riba implicitly. The industry’s flagship attempt at an independent benchmark — the Islamic Interbank Benchmark Rate (IIBR) — fails empirically to decouple: Azad et al. [14] document a stable, co-integrated “piety premium” over LIBOR and conclude that a panel-set Islamic rate operating in a global market cannot escape the conventional benchmark.

The κ -rate solves all three problems: it is ex-ante (identified from market data), objectively computed (from observable prices, funding dynamics, or credit spreads via the channels of Section 3.2), and independently verifiable (on-chain, real-time, without any reference to LIBOR/SOFR).

9.2 Relationship to the Wicksellian Tradition

Wicksell’s [49] natural rate r^* was designed to be grounded in real productivity, endogenous, and separable from monetary policy. The κ -rate is motivated by the same aims — a market-determined rate grounded in real dynamics rather than central-bank policy — and adds a property no classical natural-rate construction achieves: it is generated by a framework in which the time-preference component is provably absent ($\iota = 0$ achieves no-arbitrage, Theorem 3 of Akerer et al.). Classical natural-rate theory worked within a setting that presumed a positive interest rate for equilibrium; the Akerer framework removes that presumption. We are careful, however, to state the relationship as a structural correspondence rather than an equality of quantities, as follows.

Remark 7 (κ as a Wicksellian correspondence). The analogy to Wicksell can be stated precisely as a correspondence rather than an identity. Wicksell’s natural rate r^* is the real-activity-grounded rate that monetary policy should track; its defect, from the Islamic standpoint, is

that it retains the time-preference component r_{pref} . The κ -rate occupies the same structural role — a market-determined, real-activity-grounded rate that requires no central-bank setting — but is generated by a framework in which the time-preference component is identically absent ($\iota = 0$, Theorem 3 of Akerer et al.). We do *not* claim the numerical identity $\kappa = r^*$: κ is a convergence/hazard intensity, not a marginal product of capital, and the two need not coincide in value. The correspondence is structural — κ plays Wicksell’s role without Wicksell’s riba — not an equation between the two quantities.

9.3 From Pricing Rate to Monetary Base

The κ -rate is a riba-free *pricing* parameter. A complete interest-free monetary *system* additionally requires a theory of money *creation* that does not rest on interest-bearing debt — a distinct problem this paper does not solve. The conceptual architecture is long-standing: Al-Jarhi [13] sets out an interest-free monetary structure with central-deposit certificates and open-market operations, and Chapra [20], Khan & Mirakhor [38], and Mirakhor & Askari [42] develop the policy and risk-sharing frame. Notably, the Western full-reserve tradition converges on a similar debt-free monetary base from different premises: the Chicago Plan [16] separates money creation from credit and issues money without a corresponding interest-bearing liability. What neither tradition supplies is a *no-arbitrage price* for an interest-free monetary base. The κ -rate is the natural candidate for that role, and formalising money creation tied to real-sector convergence (through κ) rather than to debt — uniting the Islamic and full-reserve designs under a common pricing object — is undertaken in the companion κ -Standard study [9], with its general-equilibrium foundation developed in [10]. On-chain, this connects to the token-engineering literature on bonding curves and configuration spaces (Zargham et al. [50]), which formalises issuance as a deterministic function of reserves rather than of interest-bearing debt.

The asset class the κ -rate prices is already of monetary scale: the global sukuk market stood at roughly **USD 761 billion outstanding across 4,701 issues** as of May 2026 (Cbonds), with USD 277 billion of that internationally placed — a base deep and liquid enough for a benchmark convergence intensity to function as a reference rate, exactly as the conventional risk-free curve does for a market of comparable size. A riba-free term structure for an asset class this large is a practical instrument, not only a theoretical one. The ambition is in the spirit of Shiller’s [45] programme of building index-based markets to share society’s largest risks, but realised through a risk-sharing (rather than interest-bearing) pricing object.

9.4 Limitations and Open Questions

Decontamination of the calibrated rate (structure vs. magnitude). The riba-free property established here is *structural*: κ is generated by the stopping intensity of a

real claim and carries no predetermined time-preference term. It does not follow that the *magnitude* of the $\hat{\kappa}$ values reported across this program is already free of the interest-rate world — they are backed out from sukuk spreads quoted in markets where r is still priced, so the levels inherit the reserve rate. The companion κ -yield-curve study [5] makes the split explicit and benchmark-free: with the Treasury curve deleted from the estimation, the *cross-section* — credit ordering, curve shapes, and the relative pricing this framework relies on — is recovered from prices alone (rank correlation 0.84 sovereign, 0.94 corporate), while the *level* carries a wedge equal, within error, to the reserve sovereign’s own event-intensity mapped through the same transformation. The honest reading is that the structural claim (“complete real basis”) holds now, while the claim that the calibrated *number* is real-only is *conditional*, realized only as instruments come to be priced in κ natively, post-mainnet. This is the program’s central open empirical flank — decontamination is structural and prospective, not yet demonstrated on a κ -native market — and a single sovereign or corporate tranche priced directly in κ would convert it from an argument into a measurement.

Stochastic κ . The current theory uses constant κ (time-invariant convergence intensity). Appendix A derives the κ -yield curve under CIR-type stochastic κ dynamics and provides closed-form bond-price analogs. Relatedly, the constant-parameter no-arbitrage result we build on has recently been extended to *stochastic interest rates with a funding-rate clamp* — the funding cap that real exchanges (and the Baraka Protocol’s testnet implementation) impose — by Hugonnier & Jermann [30]. Re-deriving the $\iota = 0$ condition within that more general model, in which the clamp coincides with the funding cap used here, is a natural robustness check we leave to future work.

Multiple underlyings. The κ -yield curve requires perpetuals on multiple RWA underlyings. Currently, only crypto-native perpetuals exist at scale. The extension to commodity, infrastructure, and equity perpetuals on regulated DeFi platforms is a regulatory and institutional question, not a mathematical one.

Monetary policy transmission. How does a change in the central bank-published $\bar{\kappa}$ transmit to the real economy? In conventional monetary policy, the transmission mechanism operates through bank lending rates, bond yields, and the exchange rate. The κ -rate transmission mechanism is developed in general equilibrium in the companion κ -DSGE study [10]; its institutional design and empirical study remain open.

Scope of the Shariah claim. This paper establishes that the κ -rate is free of *riba* — the prohibition that has historically blocked a no-arbitrage Islamic pricing parameter. It does not, and cannot, establish the full Shariah compliance of every instrument built on κ . Contemporary scholarship [44] that accepts the removal of the interest component nonetheless raises *gharar* and owner-

ship objections to leveraged perpetual contracts specifically. These objections concern contractual structure rather than the pricing parameter, and their resolution for each instrument class (perpetual sukuk, takaful, iCDS) requires dedicated jurisprudential review by a recognized Shariah board. The contribution of this paper is the *riba*-free rate; the contract-level compliance of applications is a separate and necessary line of work. A companion research note [12] states this boundary precisely and proposes four product-layer structuring rules (asset-backing; possession-gating; insurable interest / no naked positions; expiry with delivery or *tabarru* risk-sharing) as the starting point for that review, and the classical foundation of the *riba* claim itself is developed in [11].

Cross-jurisdictional recognition. For the κ -yield curve to serve as a sovereign benchmark, it must be recognized by regulators (AAOIFI, IFSB [33], Bank Negara Malaysia, SAMA) as a legitimate *riba*-free rate. This requires AAOIFI standard revision — a process that takes 2–4 years on average.

10 Conclusion

We have proposed and developed the κ -rate — the convergence intensity of perpetual contracts under the Akerer-Hugonnier-Jermann no-arbitrage framework at $\iota = 0$ — as the first mathematically rigorous, *riba*-free alternative to the risk-free interest rate r in Islamic monetary economics.

The κ -rate satisfies all three El-Gamal criteria for *riba*-freedom: it is market-determined (not predetermined), it corresponds to real economic dynamics (not pure time preference), and it carries no interest content by construction ($\iota = 0$ is proved to achieve no-arbitrage without any time-preference discount). It plays the role of Wicksell’s natural rate with the *riba* removed — an aspiration Wicksell (1898) articulated conceptually but could not render mathematically — without claiming numerical identity to the marginal product of capital (Remark 7).

The κ -framework supplies both a single-instrument horizon–intensity identity ($\kappa = 1/T$, Proposition 3) and, for a fixed issuer, an empirically estimated κ -term-structure that is upward-sloping — the Islamic analog of the conventional yield curve, grounded in exponential stopping-time distributions rather than interest-rate expectations. The term structure has no negative values (the Feller condition holds in every constant-maturity series tested) and requires no central bank to set it: it is identified directly from market data.

The applications are immediate: perpetual sukuk pricing (no SOFR benchmark required), actuarially fair takaful contributions (no *riba* discount), iCDS for credit risk hedging (the first *riba*-free credit-protection structure; contract-level Shariah review remains with scholars), and Islamic monetary policy signalling ($\bar{\kappa}$ as the policy rate equivalent). Empirical illustration spans three independent domains: sovereign and corporate credit — a

real 195-instrument sukuk cross-section plus five-issuer constant-maturity dynamics that establish the κ level, the rating ordering and Feller non-negativity on observed sukuk; agricultural insurance (model $\hat{\kappa} = 12.06\%$ numerically close to the PMFBY rate — a units coincidence); and the dYdX v4 perpetual exchange (a decoupled, interest-free $\hat{\kappa}$ of order 10^3 p.a. at commercial scale). The central empirical point is qualitative: κ recurs across causally independent markets as the same riba-free intensity construct, with regime-specific magnitude; a single real panel uniting all domains is the principal remaining step.

The deeper contribution is historical. Islamic monetary economics has sought a riba-free pricing parameter since Chapra (1985). The parameter has now been found — not in a legal innovation or a jurisprudential reclassification, but in a proof: the Akerer–Hugonnier–Jermann Theorem 3 (first circulated 2023, published 2025), established for entirely different purposes, supplies the no-arbitrage core; the bridge to Islamic monetary theory — the riba identification and the $\kappa \leftrightarrow \lambda$ isomorphism — completes the solution to a 40-year open problem.

“Allah has made trade permissible and forbidden riba. Whoever receives an admonition from his Lord and refrains may keep [his previous gains].”
(Al-Qur’an, Al-Baqarah 2:275)

A Stochastic κ Dynamics and the Closed-Form κ -Yield Curve

A.1 Motivation

The body of the paper treats κ as a constant, following Akerer et al. [2] who establish no-arbitrage results under constant parameters. In practice, the implied κ -rate fluctuates over time as market conditions evolve: convergence is fast in liquid, well-functioning markets and slow during stress periods. A stochastic κ model is therefore necessary for:

1. Pricing instruments whose payoffs depend on the entire path of κ (e.g., κ -denominated swaps, κ -caps).
2. Constructing a dynamic κ -yield curve that shifts over time in response to market conditions.
3. Calibrating κ to a cross-section of perpetuals with different expected horizons simultaneously.

A.2 The CIR- κ Process

Definition 2 (CIR- κ Process). The **CIR- κ process** is the solution to:

$$d\kappa_t = \alpha(\bar{\kappa} - \kappa_t) dt + \sigma_\kappa \sqrt{\kappa_t} dW_t^Q, \quad \kappa_0 > 0 \quad (19)$$

under the risk-neutral measure Q , where:

- $\alpha > 0$ is the **mean-reversion speed** (how quickly κ_t reverts to its long-run level);

- $\bar{\kappa} > 0$ is the **long-run mean κ -rate**;
- $\sigma_\kappa > 0$ is the **κ volatility**;
- W_t^Q is a standard Brownian motion under Q .

Lemma 1 (Positivity — Feller Condition). *If $2\alpha\bar{\kappa} \geq \sigma_\kappa^2$, then $\kappa_t > 0$ almost surely for all $t \geq 0$.*

Proof. This is the Feller boundary condition for the square-root process (Feller [26]; see Cox, Ingersoll & Ross [21]). The drift $\alpha(\bar{\kappa} - \kappa_t)$ is strongly positive near zero, preventing κ_t from reaching zero when $2\alpha\bar{\kappa} \geq \sigma_\kappa^2$. $\square$

Remark 8. Lemma 1 is economically essential: $\kappa > 0$ is required for the stopping time $\theta_t \sim \text{Exp}(\kappa_t)$ to be well-defined. Negative κ has no economic interpretation in the Islamic finance context — there is no concept of “negative convergence speed.” The Feller condition guarantees $\kappa > 0$ always, which means the riba-free property is preserved under stochastic dynamics: $\iota = 0$ and $\kappa_t > 0$ for all t with probability 1.

A.3 Closed-Form κ -Bond Price

We define the **κ -bond** as the Islamic analog of the zero-coupon bond: its value at horizon T is the probability that the stopping event has not occurred by T — the κ -analog of the zero-coupon discount factor, $P(\kappa_t, T - t) = \mathbb{E}_t^Q[\exp(-\int_t^T \kappa_s ds)]$.

Theorem 1 (κ -Bond Pricing Under CIR- κ). *Let κ_t follow the CIR- κ process (19). The price at time t of a κ -bond with horizon T (survival probability that the stopping event has not yet occurred by time T , analogous to the discount factor) is:*

$$P(\kappa_t, T - t) = A(T - t) e^{-B(T - t) \kappa_t} \quad (20)$$

where, with $\tau = T - t$ and $h = \sqrt{\alpha^2 + 2\sigma_\kappa^2}$:

$$A(\tau) = \left[\frac{2h e^{(\alpha+h)\tau/2}}{2h + (\alpha+h)(e^{h\tau} - 1)} \right]^{2\alpha\bar{\kappa}/\sigma_\kappa^2} \quad (21)$$

$$B(\tau) = \frac{2(e^{h\tau} - 1)}{2h + (\alpha+h)(e^{h\tau} - 1)} \quad (22)$$

Proof. The κ -bond price equals the Q -expectation of the exponential integral of the κ process:

$$P(\kappa_t, \tau) = \mathbb{E}_t^Q \left[\exp \left(- \int_t^T \kappa_s ds \right) \right] \quad (23)$$

This is the Laplace transform of $\int_t^T \kappa_s ds$ evaluated at $u = 1$. For the CIR process, this Laplace transform has the affine closed form $A(\tau)e^{-B(\tau)\kappa_t}$ where A and B satisfy the Riccati ODEs:

$$B'(\tau) = 1 - \alpha B(\tau) - \frac{1}{2} \sigma_\kappa^2 B(\tau)^2, \quad B(0) = 0 \quad (24)$$

$$A'(\tau) = \alpha\bar{\kappa} B(\tau) A(\tau), \quad A(0) = 1 \quad (25)$$

The solutions to (24)–(25) are exactly equations (21)–(22) (Cox, Ingersoll & Ross [21], eq. (23)). $\square$

A.4 The Stochastic κ -Yield Curve

Definition 3 (κ -Yield). The κ -yield at horizon T given current κ_t is:

$$y_\kappa(T; \kappa_t) = -\frac{\log P(\kappa_t, T)}{T} = \frac{B(T)\kappa_t - \log A(T)}{T} \quad (26)$$

The κ -yield curve is the function $T \mapsto y_\kappa(T; \kappa_t)$ for fixed κ_t .

Proposition 4 (Limiting Behavior of the κ -Yield Curve). *Under the CIR- κ process:*

1. **Short end:** $\lim_{T \rightarrow 0} y_\kappa(T; \kappa_t) = \kappa_t$. *The short rate of the κ -yield curve equals the current κ -rate.*
2. **Long end:** $\lim_{T \rightarrow \infty} y_\kappa(T; \kappa_t) = \frac{2\alpha\bar{\kappa}}{\alpha + h}$. *The long rate of the κ -yield curve depends only on the long-run mean $\bar{\kappa}$ and the parameters α, σ_κ — it is independent of the current κ_t .*
3. **Monotone shapes:**
 - If $\kappa_t > y_\kappa(\infty)$: *yield curve is downward sloping (inverted κ -curve: market expects convergence to slow down).*
 - If $\kappa_t < y_\kappa(\infty)$: *yield curve is upward sloping (normal κ -curve: market expects convergence to speed up).*
 - If $\kappa_t = y_\kappa(\infty)$: *flat κ -curve.*

Proof. (1) follows from $\lim_{T \rightarrow 0} B(T)/T = 1$ and $\lim_{T \rightarrow 0} (-\log A(T))/T = 0$. (2) follows from $\lim_{T \rightarrow \infty} B(T)/T = 0$ and $\lim_{T \rightarrow \infty} (-\log A(T))/T = 2\alpha\bar{\kappa}/(\alpha + h)$ (standard CIR long-rate formula, Cox et al. [21]). (3) follows from the monotonicity of $B(T)$ in T and the fact that the yield is an affine function of κ_t . $\square$

Remark 9 (Reconciliation with the empirical curves). The upward-sloping sovereign κ -term-structures estimated in Section 7.1 correspond to the normal regime $\kappa_t < y_\kappa(\infty)$ of Proposition 4(3): current convergence intensity below its long-run level, so longer-horizon κ -yields rise toward $\bar{\kappa}$. The CIR- κ model thus not only permits but explains the empirically observed upward slope, while the horizon-intensity identity (11) remains a separate single-instrument statement (Remark 5).

A.5 Economic Interpretation

Table 5 translates the three yield curve shapes into Islamic monetary policy language.

These shapes map directly onto the conventional yield curve taxonomy (normal/inverted/flat) with the Islamic interpretation: the κ -yield curve inverts not because markets expect *interest rate cuts*, but because markets expect *convergence to slow* as some structural uncertainty develops. This is a fundamentally different economic signal from the conventional inverted yield curve — it reflects real asset pricing uncertainty rather than monetary tightening expectations.

Table 5: κ -yield curve shapes and monetary interpretation.

Shape	Condition	Monetary signal
Normal (upward)	$\kappa_t < \bar{\kappa}$	Current intensity below long-run; risk expected to rise toward trend
Inverted (down)	$\kappa_t > \bar{\kappa}$	Elevated current intensity (stress); expected to ease toward trend
Flat	$\kappa_t = \bar{\kappa}$	Long-run equilibrium; neutral conditions

A.6 Calibration

Under the CIR- κ model, three parameters must be estimated: α , $\bar{\kappa}$, and σ_κ . Calibration uses the cross-section of implied $\hat{\kappa}$ values across instruments with different expected horizons $T^{(i)}$:

$$(\alpha^*, \bar{\kappa}^*, \sigma_\kappa^*) = \arg \min_{\alpha, \bar{\kappa}, \sigma_\kappa} \sum_i [y_\kappa(T^{(i)}; \kappa_t^{\text{obs}}) - y_\kappa^{\text{model}}(T^{(i)}; \alpha, \bar{\kappa}, \sigma_\kappa)]^2 \quad (27)$$

For a single-instrument market with one expected horizon ($T = 1/\hat{\kappa}$), only the long-run mean $\bar{\kappa}$ is identified; the full stochastic calibration of $(\alpha, \bar{\kappa}, \sigma_\kappa)$ requires a time series. The roll-down-free constant-maturity $\hat{\kappa}$ histories of Section 7.1 (five issuers, 2018–2026) provide exactly this, identifying the CIR- κ dynamics with the Feller condition satisfied in every series.

A.7 Relationship to CIR

Table 6 provides the formal correspondence between the conventional CIR interest rate model and the CIR- κ model.

Table 6: CIR vs. CIR- κ : parameter correspondence.

CIR element	CIR convention	CIR- κ analog
State variable	Short rate r_t	Conv. intensity κ_t
Long-run mean	$\bar{r}$	$\bar{\kappa}$
Mean-reversion speed	α	α
Volatility	σ_r	σ_κ
Positivity condition	$2\alpha\bar{r} \geq \sigma_r^2$	$2\alpha\bar{\kappa} \geq \sigma_\kappa^2$
Bond price	$A(\tau)e^{-B(\tau)r_t}$	$A(\tau)e^{-B(\tau)\kappa_t}$
Short rate	r_t (riba content)	κ_t (riba-free)
Economic meaning	Time preference	Convergence speed

The CIR- κ model is the CIR model with the riba content removed: the state variable κ_t (convergence intensity, an event-rate parameter) replaces r_t (short rate, a time-preference-compensation parameter). All the mathematical machinery of CIR — bond pricing, yield curve

construction, risk-neutral calibration — carries over intact. The Islamic innovation is entirely in the economic interpretation of the state variable, not in the mathematics.

References

- [1] AAOIFI. (2017). *Shariah Standards for Islamic Financial Institutions* (Standard No. 17: Investment Sukuk; Standard No. 21: Financial Paper; Standard No. 26: Islamic Insurance). Accounting and Auditing Organization for Islamic Financial Institutions, Manama, Bahrain.
- [2] Akerer, D., Hugonnier, J., & Jermann, U. (2025). Perpetual futures pricing. *Mathematical Finance*, 1–17. <https://doi.org/10.1111/mafi.70018> (NBER Working Paper 32936, 2024; SSRN 4603820, 2023.)
- [3] Ahmed, S., & Bhuyan, R. (2026a). The interest parameter in perpetual futures: Shariah analysis and empirical evidence from centralized and decentralized exchanges. *SSRN Working Paper 6322778*. https://papers.ssrn.com/sol3/papers.cfm?abstract_id=6322778
- [4] Ahmed, S., & Bhuyan, R. (2026b). From perpetual contracts to Islamic credit: The random stopping time equivalence and a riba-free foundation for sukuk, takaful, and credit protection. *SSRN Working Paper 6322858*. https://papers.ssrn.com/sol3/papers.cfm?abstract_id=6322858
- [5] Ahmed, S., & Bhuyan, R. (2026c). The κ -yield curve: Construction, estimation methodology, and real-data evidence from sovereign sukuk. *SSRN Working Paper 6322938*. https://papers.ssrn.com/sol3/papers.cfm?abstract_id=6322938
- [6] Ahmed, S., & Bhuyan, R. (2026d). Stochastic takaful pricing under CIR hazard dynamics: Application to agricultural insurance in South and Southeast Asia. *SSRN Working Paper 6323459*. https://papers.ssrn.com/sol3/papers.cfm?abstract_id=6323459
- [7] Ahmed, S., & Bhuyan, R. (2026e). The Islamic credit default swap: Riba-free credit protection with on-chain implementation. *SSRN Working Paper 6323519*. https://papers.ssrn.com/sol3/papers.cfm?abstract_id=6323519
- [8] Ahmed, S., & Bhuyan, R. (2026f). An integrated simulation framework for decentralised finance protocols: cadCAD, reinforcement learning, game theory, and mechanism design. *SSRN Working Paper 6323618*. https://papers.ssrn.com/sol3/papers.cfm?abstract_id=6323618
- [9] Ahmed, S., & Bhuyan, R. (2026g). The κ -standard: Debt-free money creation under a riba-free convergence intensity — uniting the Chicago Plan and Islamic full-reserve monetary theory. Working paper (companion paper).
- [10] Ahmed, S., & Bhuyan, R. (2026h). Interest-free monetary equilibrium: A general-equilibrium foundation for the κ -standard. SSRN Working Paper 6890259.
- [11] Ahmed, S., & Bhuyan, R. (2026i). Money as a mirror: Classical Islamic monetary ontology and the foundations of interest-free pricing. Working paper (companion paper).
- [12] Ahmed, S., & Bhuyan, R. (2026j). The jurisprudential boundaries of an interest-free substrate: Riba, gharar, maysir, and qabd in cash-settled synthetic instruments. Working paper (companion research note).
- [13] Al-Jarhi, M. A. (1981/1983). A monetary and financial structure for an interest-free economy: Institutions, mechanism and policy. In Z. Ahmad, M. Iqbal, & M. F. Khan (Eds.), *Money and Banking in Islam*. International Centre for Research in Islamic Economics, Jeddah. (MPRA Paper 66741.)
- [14] Azad, A. S. M. S., Azmat, S., Chazi, A., & Ahsan, A. (2018). Can Islamic banks have their own benchmark? *Emerging Markets Review*, **35**, 120–136. <https://doi.org/10.1016/j.ememar.2018.02.002>
- [15] Bank for International Settlements. (2024). *OTC Derivatives Statistics at End-December 2023*. BIS Quarterly Review, March 2024.
- [16] Benes, J., & Kumhof, M. (2012). The Chicago Plan revisited. *IMF Working Paper WP/12/202*, International Monetary Fund, Washington, D.C. <https://doi.org/10.5089/9781475505528.001>
- [17] Billah, M. M. (2007). *Islamic Insurance (Takaful)*. Sweet & Maxwell Asia, Petaling Jaya, Malaysia.
- [18] Black, F., & Scholes, M. (1973). The pricing of options and corporate liabilities. *Journal of Political Economy*, **81**(3), 637–654.
- [19] Böhm-Bawerk, E. von (1890). *Capital and Interest* [Kapital und Kapitalzins]. Wagner, Innsbruck.
- [20] Chapra, M. U. (1985). *Towards a Just Monetary System*. The Islamic Foundation, Leicester.
- [21] Cox, J. C., Ingersoll, J. E., & Ross, S. A. (1985). A theory of the term structure of interest rates. *Econometrica*, **53**(2), 385–408.
- [22] DeFiLlama. (2026). dYdX v4 protocol statistics: total value locked, open interest and volumes, 2023–2026. <https://defillama.com/protocol/dydx-v4>, accessed May 2026.

- [23] Duffie, D., & Singleton, K. J. (1999). Modeling term structures of defaultable bonds. *Review of Financial Studies*, **12**(4), 687–720.
- [24] dYdX Foundation. (2024). dYdX v4: Full decentralization and the dYdX Chain. *Technical Documentation*. <https://docs.dydx.exchange>
- [25] El-Gamal, M. A. (2006). *Islamic Finance: Law, Economics, and Practice*. Cambridge University Press, Cambridge.
- [26] Feller, W. (1951). Two singular diffusion problems. *Annals of Mathematics*, **54**(1), 173–182.
- [27] Fisher, I. (1930). *The Theory of Interest*. Macmillan, New York.
- [28] Harrison, J. M., & Kreps, D. M. (1979). Martingales and arbitrage in multiperiod securities markets. *Journal of Economic Theory*, **20**(3), 381–408.
- [29] Htay, S. N. N., & Salman, S. A. (2014). Introducing takaful in India: An exploratory study on acceptability, possibility and takaful model. *Asian Social Science*, **10**(1), 1–10.
- [30] Hugonnier, J., & Jermann, U. J. (2025). Perpetual futures pricing with stochastic rates and clamp. *Working Paper*, Swiss Finance Institute. https://papers.ssrn.com/sol3/papers.cfm?abstract_id=5481866
- [31] Iqbal, M., & Molyneux, P. (2005). *Thirty Years of Islamic Banking: History, Performance, and Prospects*. Palgrave Macmillan, London.
- [32] Iqbal, Z., & Mirakhor, A. (2011). *An Introduction to Islamic Finance: Theory and Practice* (2nd ed.). John Wiley & Sons, Singapore.
- [33] Islamic Financial Services Board. (2024). *Islamic Financial Services Industry Stability Report 2024*. IFSB, Kuala Lumpur.
- [34] Jarrow, R. A., & Turnbull, S. M. (1995). Pricing derivatives on financial securities subject to credit risk. *Journal of Finance*, **50**(1), 53–85.
- [35] Jobst, A. A. (2007). The economics of Islamic finance and securitization. *Journal of Structured Finance*, **13**(1), 6–27.
- [36] Kamali, M. H. (2007). *Islamic Commercial Law: An Analysis of Futures and Options*. The Islamic Texts Society, Cambridge.
- [37] Keynes, J. M. (1936). *The General Theory of Employment, Interest and Money*. Macmillan, London.
- [38] Khan, M. S., & Mirakhor, A. (Eds.). (1987). *Theoretical Studies in Islamic Banking and Finance*. Institute for Research and Education, Houston.
- [39] Khan, T., & Ahmed, H. (2001). Risk management: An analysis of issues in Islamic financial industry. *IRTI Occasional Paper* No. 5, Islamic Development Bank, Jeddah.
- [40] Lando, D. (1998). On Cox processes and credit risky securities. *Review of Derivatives Research*, **2**(2–3), 99–120.
- [41] Merton, R. C. (1974). On the pricing of corporate debt: The risk structure of interest rates. *Journal of Finance*, **29**(2), 449–470.
- [42] Mirakhor, A., & Askari, H. (2010). *Islam and the Path to Human and Economic Development*. Palgrave Macmillan, New York.
- [43] Nelson, C. R., & Siegel, A. F. (1987). Parsimonious modeling of yield curves. *Journal of Business*, **60**(4), 473–489.
- [44] SeekersGuidance. (2025, October 2). Are crypto perpetual contracts still haram if there’s no interest, and the asset price tracks? Answered by Shaykh Muhammad Carr; checked and approved by Shaykh Faraz Rabbani. <https://seekersguidance.org/answers/transactions/are-crypto-perpetual-contracts-still-haram-if-theres-no-interest-and-the-asset-price-tracks/>
- [45] Shiller, R. J. (1993). *Macro Markets: Creating Institutions for Managing Society’s Largest Economic Risks*. Oxford University Press, Oxford.
- [46] Uddin, M. H., Kabir, S. H., Hossain, M. S., Wahab, N. A., & Liu, J. (2022). Why do sukuk (Islamic bonds) need a different pricing model? *International Journal of Finance & Economics*, **27**(2), 2210–2234. <https://doi.org/10.1002/ijfe.2269>
- [47] Usmani, M. T. (2002). *An Introduction to Islamic Finance*. Kluwer Law International, The Hague.
- [48] Vasicek, O. (1977). An equilibrium characterization of the term structure. *Journal of Financial Economics*, **5**(2), 177–188.
- [49] Wicksell, K. (1898). *Interest and Prices* [Geldzins und Güterpreise]. Gustav Fischer, Jena.
- [50] Zargham, M., Shorish, J., & Paruch, K. (2020). From curved bonding to configuration spaces. *IEEE International Conference on Blockchain and Cryptocurrency*, 1–9.
- [51] Zarqa, M. A. (1983). Stability in an interest-free Islamic economy: A note. *Pakistan Journal of Applied Economics*, **2**(2), 181–188.

EDITORIAL NOTE TO THIS PRINTING

Paper 8 · The κ -Standard

Working paper, revised July 2026 SSRN 6893880

Role in the programme. The monetary architecture: what money looks like if the primitive goes all the way down. **What it establishes.** Base money as κ -priced sovereign equity rather than interest-bearing debt; a closed-form determinacy condition, $\gamma_s < e^{\bar{\kappa}} - 1$; Cantillon neutrality; an escape from the real-bills indeterminacy that sinks naive proposals. **Corrected or extended later.**

A July 2026 honest-boundary section concedes that the monetary anchor is a priced intensity whose premium cannot be mutualised away; the base layer stands, by its own logic, on the most premium-laden ground. **Not established here.** The transition path, and the reserved ruling this architecture is gated on. This is the frontier piece, and it is labelled as one.

Replication for every number in this paper: DOI 10.5281/zenodo.21290557 · github.com/shehzadahmed-xx/kappa-replication

The κ -Standard: Debt-Free Money Creation under a Riba-Free Convergence Intensity

Uniting the Chicago Plan and Islamic Full-Reserve Monetary Theory

Shehzad Ahmed*

Rafiqul Bhuyan[†]

June 2026

Abstract

Two traditions of monetary reform—the Islamic interest-free programme of Al-Jarhi (1981) and Chapra (1985), and the Western full-reserve tradition of Fisher (1936) and Benes & Kumhof (2012)—independently reached the same design: money should be created debt-free by the state and separated from interest-bearing bank credit. Both stalled at the same point. Once money creation is divorced from interest-bearing debt, the assets that back the money base must still be *priced*, and the natural pricing object—the risk-free interest rate r —is precisely the *riba* the reform sets out to remove. Neither tradition supplied a *riba*-free price at which the market could clear. This paper provides one.

Building on the convergence intensity κ of Akerer et al. (2025), shown by the companion papers to be a *riba*-free, no-arbitrage pricing primitive estimable from real *sukuk* markets, we construct the κ -Standard: a monetary architecture in which (i) base money equals the κ -priced value of a portfolio of sovereign perpetual *sukuk*, so that money is sovereign equity rather than debt; (ii) new money is created only by monetising new real-sector projects at their κ -implied price, tethering money growth to real output; (iii) open-market operations transmit through the system-wide convergence intensity $\bar{\kappa}$ rather than an interest rate, with the central bank setting the *quantity* of money while the market sets the *price of capital*; and (iv) commercial banks operate under one hundred per cent reserves, splitting into fee-based (*ujrah*) custody and profit-sharing (*mudarabah*) investment.

We derive the law of motion of the money supply and show that a shock to $\bar{\kappa}$ —a system-wide credit deterioration—contracts the base automatically. We then derive a countercyclical issuance rule that neutralises this revaluation channel and holds the base on a target growth path. The correction is bounded by available real capacity and a Sovereign Stabilisation Fund — a deliberate trade of unbounded-but-unbacked fiat stabilisation for bounded-but-backed stabilisation (roughly \$3.6 bn per 100 bp on the full actively-quoted USD-*sukuk* cross-section, \$195.2 bn of par).

We are explicit about what the architecture does *not* achieve: it removes debt-financed and unbacked money creation and ties money to real-asset value, but its price level is determinate only under a fiscal closure — the natural determinacy closure for a system with no interest-rate instrument¹ — and a companion paper shows this survives in full general equilibrium. Its stabilisation is bounded. We propose κ as a candidate *riba*-free clearing object that lets a full-reserve money base clear without an interest rate, serving both the Chicago-Plan and Islamic full-reserve programmes; we do not claim it was the sole obstacle to either programme’s adoption, which is substantially institutional and political. The empirical reality of κ rests so far only on the authors’ own companion working papers, pending independent replication, so we present this as a theoretical architecture.

*Department of Finance, Independent University Bangladesh (IUB), Dhaka. 2120922@iub.edu.bd.

[†]Corresponding author. Department of Accounting and Finance, Alabama A&M University, Normal, AL, USA. rafiquel.bhuyan@aamu.edu.

¹A constant-money-growth (quantity) rule with well-defined money demand can also deliver determinacy without an interest instrument, but only up to speculative-bubble equilibria; the fiscal closure used here pins the price level uniquely.

Keywords: *riba*-free money; full-reserve banking; Chicago Plan; sovereign money; convergence intensity; κ -rate; debt-free seigniorage; Islamic monetary economics; tokenised sukuk.

JEL codes: E42, E51, E58, G21, P43, Z12.

1 Introduction

In the contemporary monetary system the overwhelming majority of money is created by commercial banks when they extend loans (McLeay et al., 2014; Werner, 2014; Jakab & Kumhof, 2015), and the residual is created by central banks purchasing interest-bearing government debt. In both channels the money supply expands only when some party incurs an interest-bearing liability. Money is, structurally, a debt instrument that accrues *riba*—a predetermined excess over principal, independent of the underlying economic outcome (El-Gamal, 2006).

This is not a peripheral feature. It is the reason two otherwise unrelated schools of monetary thought converged, decades apart and from opposite premises, on the same reform. On the Islamic side, Al-Jarhi (1981) designed an interest-free monetary structure built on central-deposit certificates and open-market operations, insisting that seigniorage belong to the state and that new money fund real productive projects rather than interest-bearing loans; Chapra (1985) and Khan & Mirakhor (1987) added the macro frame, requiring money growth to track real output and rejecting the fractional-reserve creation of money as debt, and the later risk-sharing literature (Iqbal & Mirakhor, 2011; Askari et al., 2012) made equity rather than debt the organising principle. On the Western side — building on the 1933 Chicago memorandum of Knight, Simons and colleagues — Simons (1936) and Fisher (1936) proposed one hundred per cent reserve banking to separate money creation from bank credit, a programme later endorsed by Friedman (1960), revived as sovereign money (Dyson et al., 2016), and validated in a modern dynamic general-equilibrium model by Benes & Kumhof (2012), who showed that money issued as state equity rather than debt could reduce business cycles, eliminate bank runs, and permit large reductions in public and private debt.

Both traditions reached the same wall. If money creation is to be severed from interest-bearing debt, the real assets backing the money base must still be priced so that the market can clear—and the obvious pricing object, the risk-free rate r , is exactly the *riba* the reform exists to remove. The Chicago Plan’s general-equilibrium clearing still runs through an interest rate; Al-Jarhi and Chapra specified the institutional plumbing but no clearing price. Each school knew the money base had to be debt-free; neither could say at *what price* the assets behind it should trade once interest was abolished.

Contribution. This paper proposes κ as a candidate for that missing price. Akerer et al. (2025) prove that a perpetual contract is governed by two parameters—a convergence intensity κ and an interest parameter ι —and that at $\iota = 0$ the contract retains a unique, arbitrage-free price: the risk-neutral expectation of the spot value sampled at an exponentially distributed stopping time of mean $1/\kappa$. Pricing survives the removal of interest; κ does the work that r used to do. The companion papers establish that this κ is isomorphic to the credit hazard rate (Ahmed & Bhuyan, 2026a), that it forms a measurable *riba*-free term structure in real sovereign and corporate sukuk markets (Ahmed & Bhuyan, 2026b), and that it can organise the pricing of an entire Islamic financial system (Ahmed & Bhuyan, 2026d). The capstone explicitly leaves the monetary-base step—formalising money creation tied to real-sector convergence through κ —as future work. This paper undertakes it.

We construct the κ -Standard, a monetary architecture with four components, and establish their formal properties:

1. **Balance sheet** (Section 2). Base money equals the κ -priced value of a portfolio of sovereign

perpetual sukuk; money is sovereign equity backed one-to-one by the present value of real assets, not a debt liability.

2. **Debt-free seigniorage** (Section 3). New money is created only by monetising a new real-sector project at its κ -implied price; money growth is tethered to real-output growth.
3. **κ -OMO** (Section 4). Open-market operations transmit through the system-wide $\bar{\kappa}$, not an interest rate. We derive the law of motion of the money supply and show that a $\bar{\kappa}$ shock contracts the base automatically. The central bank controls the *quantity* of money; the market sets the *price of capital* $\bar{\kappa}$, a quantity-policy regime loosely in the spirit of Wicksell.
4. **Full-reserve banking** (Section 5). Commercial banks hold one hundred per cent reserves and split into fee-based custody (*ujrah*) and profit-sharing investment (*mudharabah*); the money multiplier is one, so demand-deposit runs are precluded within the model, while investment-pool losses are borne as equity, not abolished.

Section 6 confronts the architecture’s central hazard—the procyclical contraction of Corollary 4.4—with a countercyclical issuance rule, and is honest about its bounded reach. Section 7 addresses the question a monetary-economics reader will press hardest—whether the price level is determinate at all—and shows that, because the backing is real equity and policy is a quantity operation, the system escapes real-bills indeterminacy and pins a unique price level under the fiscal theory of the price level, the natural determinacy closure for a system with no interest-rate instrument (Theorem 7.2). Section 9 states the remaining boundaries: the κ -Standard removes debt-financed and unbacked money creation, its price-level determinacy is established under a fiscal closure here and shown to survive full general equilibrium in the companion paper (Ahmed & Bhuyan, 2026e), its stabilisation is bounded by real-capacity backing, and it is a monetary *architecture*—like the Chicago Plan—whose pricing primitive’s empirical reality is carried by the companion papers, not re-established here. Section 10 concludes.

2 The κ -Standard Balance Sheet

We work on a filtered probability space $(\Omega, \mathcal{F}, \{\mathcal{F}_t\}, \mathbb{Q})$ satisfying the usual conditions, with $\mathbb{Q}$ the risk-neutral (pricing) measure of Akerer et al. (2025) at $\iota = 0$.

Definition 2.1 (Sovereign perpetual sukuk). A *sovereign perpetual sukuk* indexed by i is a tokenised claim on a real productive asset (an infrastructure project, energy grid, or state commodity reserve) with value process $X_t^{(i)}$. Its *riba-free* price at $\iota = 0$ is

$$P_t^{(i)} = \mathbb{E}_t^{\mathbb{Q}}[X_{\tau_i}^{(i)}], \quad \tau_i \sim t + \text{Exp}(\kappa^{(i)}), \quad (1)$$

where $\kappa^{(i)} \geq 0$ is the asset’s convergence intensity and τ_i is its (credit/completion) stopping time, independent of $X^{(i)}$ under $\mathbb{Q}$.

Equation (1) is the *riba-free* pricing primitive; it is not, by itself, the object the central bank holds. If $X^{(i)}$ is a $\mathbb{Q}$ -martingale and τ_i is independent of it, optional sampling gives $P_t^{(i)} = X_t^{(i)}$ *independently of $\kappa^{(i)}$* : a pure, at-par, no-credit claim. To obtain an asset whose price responds to κ —which the policy mechanism of Section 4 requires—we take the credit-bearing claim of the companion papers, in which default/resolution carries recovery $\delta_i \in [0, 1]$. There the convergence intensity is identified with the credit hazard (Ahmed & Bhuyan, 2026a,b) as a *reduced-form* Duffie–Singleton survival valuation (Duffie & Singleton, 1999) (a modelling choice, not a theorem we re-derive from (1)), giving the spread $s_i = \kappa^{(i)}(1 - \delta_i)$ and the *constant-maturity* price

$$P_t^{(i)} = \mathcal{P}_i e^{-\kappa_t^{(i)}(1-\delta_i)T_i}, \quad (2)$$

where $\mathcal{P}_i$ is par value and T_i is a *fixed reference tenor*. We assume throughout that the central bank runs a *constant-maturity book*—its sukuk holdings are continuously rolled at target tenors

T_i (e.g. seven years), exactly the constant-maturity construction estimated in [Ahmed & Bhuyan \(2026b\)](#)—so that T_i is constant, the price is a function of the convergence intensity alone, and there is no maturity-decay term to carry through the dynamics. The structural fact we use repeatedly is

$$\frac{\partial P_t^{(i)}}{\partial \kappa_t^{(i)}} = -(1 - \delta_i) T_i P_t^{(i)} < 0 : \quad (3)$$

the price of a credit-bearing constant-maturity sukuk is strictly decreasing in its convergence intensity, with a sensitivity equal to its credit-loss-weighted duration $(1 - \delta_i)T_i$. (At $\delta_i \rightarrow 1$ this collapses to the at-par limit $P = \mathcal{P}$, the no-credit analogue of the martingale case $P = X$, independent of κ ; such claims cannot transmit κ -OMO and are excluded from the policy portfolio—[Remark 9.1](#).)

Definition 2.2 (The Islamic Central Bank balance sheet). The Islamic Central Bank (ICB) holds $N_i \geq 0$ units of each sovereign perpetual sukuk and no interest-bearing instrument. Its monetary liabilities—the base money supply M_t in circulation plus bank reserves—equal its assets:

$$M_t = \sum_{i=1}^n N_i P_t^{(i)} = \sum_{i=1}^n N_i \mathbb{E}_t^{\mathbb{Q}}[X_{\tau_i}^{(i)}]. \quad (4)$$

Equation (4) is an accounting identity (assets equal liabilities), but its economic content is the substitution of the asset side. In the fiat system the central bank’s assets are interest-bearing government bonds, and base money is the matching debt liability. Under the κ -Standard the asset side is composed entirely of κ -priced claims on real output, so base money is the present value of sovereign *equity*. The asset class this presupposes already exists at monetary scale: as of 31 May 2026 the global sukuk market stood at USD 761.3 billion outstanding across 4,701 issues, of which USD 277.3 billion was internationally placed.² A market of this size is large enough for benchmark construction at the sovereign tier, though depth is concentrated (USD 277.3 bn internationally placed) and the scarcity of high-quality sovereign sukuk remains the binding operational constraint ([Section 4](#)).

Proposition 2.1 (Money as fully-backed sovereign equity). *Under [Definition 2.2](#), every unit of base money is backed one-to-one by the risk-neutral present value of a real productive asset, and the ICB cannot expand M_t except by acquiring a claim of equal κ -implied present value. In particular there exists no monetary liability of the ICB without a real-asset counterpart of equal value; money creation unbacked by real output is precluded within the model.*

Proof. Immediate from (4): M_t is defined as the sum of the present values $P_t^{(i)}$ of held assets. Any increase ΔM_t requires $\Delta \sum_i N_i P_t^{(i)} > 0$, i.e. the acquisition of additional units ΔN_i of an asset priced at $P_t^{(i)}$, of present value $\sum_i \Delta N_i P_t^{(i)} = \Delta M_t$. No accounting entry increases the left side without an equal increase of the right. $\square$

[Proposition 2.1](#) is the formal counterpart of the Chicago-Plan and Al-Jarhi claim that money can be issued as equity rather than debt. The idea that base money is best understood as an equity-like residual claim on the issuer’s cash flows rather than as a redeemable debt has a Western lineage that the κ -Standard makes literal: [Buiter \(2007\)](#) treats irredeemable base money as “an asset of the holder but not a liability of the issuer,” and [Hall & Reis \(2015\)](#) show that it is the seigniorage franchise—a flow of real resources—not paid-in capital that backs a central bank’s liabilities and keeps it solvent. The κ -Standard replaces that implicit franchise with an explicit portfolio of κ -priced sovereign claims. We stress in [Section 9](#) that “fully backed” is a statement about the *asset side*, not a guarantee about the price level.

²Cbonds, cbonds.com/sukuk (retrieved 4 June 2026); the outstanding total rose from USD 741.1 billion a month earlier.

3 Debt-Free Seigniorage: Money Creation Tied to Real Output

How does a sovereign fund a new \$5 billion high-speed rail without issuing an interest-bearing bond? Under the κ -Standard the creation of new money is a three-step monetisation of new real capacity.

Definition 3.1 (κ -monetisation). To finance a new real-sector project of par value $\mathcal{P}_*$ and effective horizon T_* :

1. **Issuance.** The Treasury issues a new sovereign perpetual sukuk against the project’s future output, with value process $X^{(*)}$.
2. **Calibration.** The open market assigns a convergence intensity $\kappa^{(*)}$ to the project from its completion and revenue risk, yielding the price $P^{(*)} = \mathcal{P}_* D(\kappa^{(*)}, T_*)$ of (2).
3. **Monetisation.** The ICB purchases the sukuk from the Treasury and credits the Treasury’s account with newly created base money $\Delta M = P^{(*)}$.

The created money is exactly the market’s κ -implied present value of the new real asset—no more, no less. This yields the Chapra tether as an identity rather than a policy aspiration.

Proposition 3.1 (Real-output tether). *Under κ -monetisation, the issuance component of base-money growth equals the κ -priced value of newly monetised real capacity:*

$$\Delta M_t^{\text{iss}} = \sum_{* \in \mathcal{I}_t} \mathcal{P}_* D(\kappa^{(*)}, T_*), \quad (5)$$

where $\mathcal{I}_t$ is the set of projects monetised in $[t, t + dt]$. The money supply expands only as real productive capacity is created and priced; in the absence of new real projects, $\Delta M_t^{\text{iss}} = 0$.

Proof. Sum the per-project creations $\Delta M = P^{(*)}$ of Definition 3.1 over $\mathcal{I}_t$ and substitute (2). □

Equation (5) is the issuance (real-output) channel of money growth. It is distinct from, and additive to, the revaluation channel of Section 4, which moves the value of the *existing* stock as convergence intensities move. We keep the two channels separate throughout: new money tracks new real output; the standing money base revalues with κ .

Seigniorage socialisation. The arrangement is not merely interest-free; it relocates the seigniorage from money creation onto the public balance sheet. In a fractional-reserve system the bulk of the money stock is privately created: commercial banks originate broad money when they lend, so that deposits—not state-issued cash—constitute the overwhelming majority of the circulating medium (McLeay et al., 2014; Werner, 2014), on the order of 95% in mature economies (Dyson et al., 2016). The rent from that creation—the spread between the near-zero cost of issuing a redeemable-at-par deposit and the return on the assets it funds—accrues to the issuing banks rather than to the public. Under the κ -Standard, by construction, the demand-deposit silo is 100% reserved and creates no money (Proposition 5.1); *all* new money enters through the κ -monetisation of Definition 3.1, so the issuance gain ΔM^{iss} of (5) is captured by the issuer and, through it, the public treasury rather than by private intermediaries. This is the same transfer that motivates the Chicago-Plan and sovereign-money proposals (Fisher, 1936; Benes & Kumhof, 2012; Dyson et al., 2016) and that Buiter (2007) formalises as the central bank’s monopoly seigniorage; the κ -Standard’s contribution is to tie the size of that public seigniorage to priced real output via (5) rather than to a discretionary monetary-base target. The distributional counterpart of this socialisation—who receives newly issued money first—is treated in Section 8.

4 κ -Space Open-Market Operations and the Law of Motion of Money

A central bank that cannot set a risk-free interest rate needs a different steering mechanism. Under the κ -Standard the instrument is the open-market purchase and sale of sovereign sukuk, and the transmission variable is the system-wide convergence intensity $\bar{\kappa}$.

4.1 The transmission mechanism

We treat $\kappa_t^{(i)}$ as the *market-clearing* convergence intensity that prices the instrument through (2)—the value implied by the price at which float is held—not a frozen fundamental. Let private demand for the float of sukuk i be downward-sloping in its price (equivalently, upward-sloping in $\kappa_t^{(i)}$), as for any traded credit instrument.

Proposition 4.1 (κ -OMO transmission). *Under a downward-sloping demand for float, an ICB purchase $\Delta N_i > 0$ (financed by newly created base money $\Delta M = \Delta N_i P_t^{(i)} > 0$) reduces the private float that the market must hold and so clears at a higher price; by (2) the market-clearing convergence intensity falls,*

$$\Delta M > 0 \implies \Delta P_t^{(i)} > 0 \implies \Delta \kappa_t^{(i)} < 0, \quad (6)$$

with the magnitude governed by the demand elasticity. A sale reverses every sign.

Proof. $\Delta M = \Delta N_i P_t^{(i)} > 0$ by construction. Removing ΔN_i from private float shifts the residual supply curve left; under downward-sloping demand the market clears at a higher price, $\Delta P_t^{(i)} > 0$; by (2), $\Delta \kappa_t^{(i)} < 0$. $\square$

Operationally this is a *quantity* operation—the instrument is ΔM , the implied κ -move is its endogenous by-product on the priced float—and it is therefore close in mechanics to conventional quantitative easing. The differences are in the *asset* and the *signal*, not the operation: the asset is a real-backed, *riba*-free sukuk rather than interest-bearing debt, and the policy stance is read off the market-clearing κ rather than off an administered interest rate. We do not claim a novel transmission channel to investment or broad credit; establishing how κ -OMO propagates beyond the priced instruments to aggregate demand is a general-equilibrium question we leave open (Section 9). What Proposition 4.1 establishes is only the sign and the *riba*-free form of the operation. Nor is interest-free liquidity management merely hypothetical: Bank Negara Malaysia already conducts injection and absorption without an interest rate through *qard* acceptance, *wadiah* placements, and sukuk-based notes, and the operational constraints of such a regime—above all the scarcity of high-quality sovereign sukuk—are documented by El Hamiani Khatat (2016). The κ -Standard supplies the missing object those operations price against: a market-clearing convergence intensity rather than an administered profit rate.

Corollary 4.2 (Quantity, not price, is the instrument). *The central bank controls the quantity of base money it injects; the price of capital $\bar{\kappa}$ is determined by the market and cannot be pegged by policy. This shares one feature with Wicksell’s natural rate—it is exogenous to the policy instrument—and the analogy is limited to that feature. We do not claim $\bar{\kappa}$ is the natural rate of interest: Wicksell’s natural rate is a real intertemporal price (the marginal product of capital) that is well defined even in a default-free economy, whereas $\bar{\kappa}$ is a credit-hazard intensity that, by Remark 9.1, vanishes precisely in the default-free limit. A real-side shock can move the natural rate without moving $\bar{\kappa}$, and a pure credit shock can move $\bar{\kappa}$ without moving the natural rate; the two are distinct objects. The point we do make is structural and weaker: whereas Woodford (2003) organises policy around managing a short-term interest rate, the κ -Standard organises it around managing the money quantity, with no interest instrument anywhere in the transmission.*

4.2 The law of motion of the money supply

Let each convergence intensity follow the Cox–Ingersoll–Ross dynamics established for κ in the companion programme,

$$d\kappa_t^{(i)} = \alpha_i(\bar{\kappa}^{(i)} - \kappa_t^{(i)}) dt + \nu_i \sqrt{\kappa_t^{(i)}} dW_t^{(i)}, \quad (7)$$

with $\alpha_i > 0$ the mean-reversion speed, $\bar{\kappa}^{(i)} > 0$ the long-run mean, $\nu_i > 0$ the volatility, and $W^{(i)}$ standard Brownian motions. Write $\beta_i \equiv (1 - \delta_i)T_i$ for the credit-loss-weighted duration of (3). Because the book is constant-maturity (T_i fixed), $P_t^{(i)}$ depends on t only through $\kappa_t^{(i)}$, so Itô's lemma applied to $P_t^{(i)} = P_i e^{-\beta_i \kappa_t^{(i)}}$ has no maturity-decay term and gives, exactly,

$$\frac{dP_t^{(i)}}{P_t^{(i)}} = -\beta_i d\kappa_t^{(i)} + \frac{1}{2}\beta_i^2 \nu_i^2 \kappa_t^{(i)} dt. \quad (8)$$

Summing the revaluation of the standing portfolio gives the law of motion of the base.

Theorem 4.3 (Money-supply dynamics under the κ -Standard). *With the balance sheet (4), constant-maturity prices (2), and dynamics (7), and writing $\beta_i = (1 - \delta_i)T_i$, the standing money base evolves as*

$$\begin{aligned} dM_t = & \underbrace{\sum_i N_i P_t^{(i)} \left[\beta_i \alpha_i (\kappa_t^{(i)} - \bar{\kappa}^{(i)}) + \frac{1}{2} \beta_i^2 \nu_i^2 \kappa_t^{(i)} \right] dt}_{\text{revaluation drift}} \\ & - \underbrace{\sum_i N_i P_t^{(i)} \beta_i \nu_i \sqrt{\kappa_t^{(i)}} dW_t^{(i)}}_{\text{revaluation diffusion}} + \underbrace{dM_t^{\text{iss}}}_{\text{new issuance}}. \end{aligned} \quad (9)$$

The issuance term $dM_t^{\text{iss}} = \sum_{* \in \mathcal{I}_t} P^{(*)} dN_*$ is a pure-jump process (Proposition 3.1); when projects arrive as a finite-intensity point process its compensator defines the finite issuance rate $\varphi_t = \mathbb{E}_t[dM_t^{\text{iss}}]/dt$ used below — a money-creation flow.

Proof. Differentiate (4): $dM_t = \sum_i N_i dP_t^{(i)} + dM_t^{\text{iss}}$, the last term collecting changes in holdings dN_i from new monetisation. Substitute $dP_t^{(i)}/P_t^{(i)}$ above and (7). $\square$ $\square$

The single most important comparative-static of the architecture follows by differentiating the revaluation drift in the long-run mean.

Corollary 4.4 (Automatic contraction under a credit-deterioration shock). *A uniform upward shock to the system-wide long-run convergence intensity, $\bar{\kappa}^{(i)} \mapsto \bar{\kappa}^{(i)} + \Delta \bar{\kappa}$ for all i , contracts the standing money base: holding the current $\kappa_t^{(i)}$ fixed,*

$$\frac{\partial (\mathbb{E}_t[dM_t]/dt)}{\partial \bar{\kappa}} = - \sum_i N_i P_t^{(i)} \beta_i \alpha_i < 0, \quad \beta_i = (1 - \delta_i)T_i. \quad (10)$$

A system-wide credit deterioration—widening $\bar{\kappa}$ —automatically shrinks the money supply as intensities mean-revert toward the higher target and prices fall.

Proof. Differentiate the revaluation-drift term of (9) with respect to $\bar{\kappa}^{(i)}$; only the $-\beta_i \alpha_i \bar{\kappa}^{(i)}$ piece depends on it, giving $-N_i P_t^{(i)} \beta_i \alpha_i$ per asset. Sum over i . $\square$ $\square$

Corollary 4.4 is the κ -Standard's built-in quantity rule: the base expands only as real output is monetised (Proposition 3.1) and contracts automatically when the sovereign asset portfolio's convergence deteriorates. Whether this automatic contraction stabilises or destabilises is not self-evident, and we confront it as the architecture's central hazard in Section 6, where a countercyclical issuance rule neutralises it within the limits of the backing.




Figure 1: The κ -Standard money base, calibrated to real sovereign-sukuk $\bar{\kappa}$ (Ahmed & Bhuyan, 2026b), $\delta = 0.60$, $P = \mathcal{P} e^{-\kappa(1-\delta)T}$. *Left*: the base $M(\bar{\kappa}) = \sum_i \mathcal{P}_i e^{-(\bar{\kappa}_i + \Delta\bar{\kappa})(1-\delta)T_i}$ for the real outstanding-weighted portfolio (\$82.3 bn par, Cbonds). Easing (open-market purchases) lowers $\bar{\kappa}$ and raises M ; a system-wide credit deterioration raises $\bar{\kappa}$ and contracts the base automatically (+100 bp $\Rightarrow$ -2.6%). *Right*: base money created per \$100 of par by sovereign—the $\bar{\kappa}$ -discount—falling with credit quality from \$97.5 (Malaysia, A) to \$86.4 (Bahrain, B).

4.3 A worked example, calibrated to real sukuk data

To make the mechanism concrete we calibrate it with the *real* long-run convergence intensities estimated on sovereign sukuk in Ahmed & Bhuyan (2026b), using the closed-form spread discount their identification implies, $P = \mathcal{P} e^{-\kappa(1-\delta)T}$ with $\delta = 0.60$.

Debt-free seigniorage of a single project. Suppose the Indonesian Treasury (real 7-year $\bar{\kappa} = 2.80\%$) monetises a \$5 bn infrastructure project of horizon $T = 7$ years. The implied spread is $s = \kappa(1 - \delta) = 112$ bp—of the same order as Indonesia’s recent 7-year sovereign sukuk G-spread—so the κ -discount is $D = e^{-sT} = 0.925$ and the central bank creates

$$\Delta M = \mathcal{P}_* D = \$5\text{bn} \times 0.925 = \$4.62\text{bn}$$

of new base money against the project. The \$0.38 bn gap between par and money created is the convergence/credit discount, not interest: the state funds real capacity debt-free, and the base expands by exactly the market’s κ -implied present value of new output (Proposition 3.1).

Portfolio revaluation and the $\bar{\kappa}$ shock. Consider an Islamic Central Bank holding the five sovereigns’ *actual* outstanding international USD sukuk (Malaysia \$2.5 bn, Saudi Arabia \$30.0 bn, Indonesia \$24.4 bn, Türkiye \$14.2 bn, Bahrain \$11.2 bn — \$82.3 bn par in total, real Cbonds volumes), priced at their real $\bar{\kappa}$ (Malaysia, Saudi, Indonesia, Bahrain at 7y; Türkiye at 5y). The baseline money base is $M_0 = \$75.0$ bn. Figure 1 traces $M(\bar{\kappa})$ as every issuer’s $\bar{\kappa}$ is shocked uniformly: a +100 bp system-wide credit deterioration contracts the base by **2.6%** (to \$73.1 bn; base semi-elasticity $\partial \ln M / \partial \bar{\kappa} \approx -2.67\%$ per 100 bp), exactly the automatic contraction of Corollary 4.4. The contraction is, to first order, a *duration* effect— $\partial \ln P / \partial \kappa = -(1 - \delta)T$, so all same-maturity issuers contract identically (the 7-year sovereigns by $\approx 2.8\%$, the 5-year Türkiye by $\approx 2.0\%$) regardless of credit level; the credit level instead sets how much money each dollar of par backs at baseline (right panel: \$97.5 for single-A Malaysia down to \$86.2 for Türkiye at its 5-year tenor — below single-B Bahrain’s \$86.4 despite the shorter tenor, because Türkiye’s higher $\bar{\kappa}$ dominates).

Segment	n	par outstanding (\$bn)	κ -base M (\$bn)
Sovereign	40	59.6	57.0
Corporate	154	134.6	127.9
Municipal	1	1.0	0.9
Full universe	195	195.2	185.8

Table 1: The κ -Standard money base computed on the *full* real outstanding USD sukuk universe for which Cbonds reports an amount, a G -spread and a maturity (195 instruments: 40 sovereign, 154 corporate, one municipal; account 972307). The base is the κ -priced value of the asset class, $M = \sum_i N_i e^{-\kappa_i(1-\delta)T_i} = \185.8 bn on $\$195.2$ bn of par. A uniform +100 bp shock to every issuer’s κ contracts the base by 1.96% (semi-elasticity $\partial \ln M / \partial \bar{\kappa} \approx -1.99\%$ per 100 bp) — smaller than the seven-year basket because the broader universe is shorter-duration. $\delta = 0.60$.

Two cautions on reading Figure 1. First, it is a *comparative static*—the steady-state revaluation of the base as each $\bar{\kappa}$ is reset to a new level $\bar{\kappa} + \Delta\bar{\kappa}$ and the book re-priced—not a simulation of the full stochastic dynamics of Theorem 4.3. It shows the deterministic revaluation (drift) channel; it suppresses the diffusion term, so the realised base along any path carries volatility around this curve, and the curve should be read as the conditional mean response, not a deterministic trajectory. Second, it is an architectural illustration, not a claim about any actual central bank’s portfolio; its only purpose is to show that the law of motion of Theorem 4.3 produces economically sensible magnitudes when fed the κ values that real sukuk markets actually exhibit.

The base on the full actively-quoted USD-sukuk cross-section. The five-issuer basket is illustrative; the architecture, however, applies to the entire asset class. Table 1 computes the base $M = \sum_i N_i e^{-\kappa_i(1-\delta)T_i}$ directly on *every* actively-quoted USD sukuk for which Cbonds reports an outstanding amount, a G -spread and a maturity — 195 instruments (40 sovereign, 154 corporate, 1 municipal), $\$195.2$ bn of par. The κ -priced base is $\$185.8$ bn, and a uniform +100 bp shock to every issuer’s κ contracts it by 1.96% (semi-elasticity $\approx -1.99\%$ per 100 bp). The semi-elasticity is smaller than the seven-year basket’s -2.67% because the broader universe is shorter-duration: the same portfolio-duration mechanism ($\partial \ln M / \partial \kappa = -(1-\delta)\bar{T}$ in aggregate), now read off the whole traded market rather than a representative basket.

The base as a realised series. The auto-contraction need not stay hypothetical. Figure 2 traces the base the architecture *would* have produced month by month, fed the five core sovereigns’ *actual realised* constant-maturity κ (Cbonds daily G -spreads, rolled to a fixed 7-year tenor) and their real outstanding volumes ($\$82.3$ bn par). Over the five-issuer overlap window (2022–2026) the base swings from $\$71.7$ bn at peak credit stress to $\$76.5$ bn at the tightest spreads — a 6.3% peak-to-trough move driven *purely by credit*, with no change in the quantity of outstanding sukuk. This is the law of motion of Theorem 4.3 read off the historical record rather than a comparative static: a κ -Standard base automatically absorbs credit cycles, contracting as spreads widen and expanding as they compress, exactly the countercyclical-on-price behaviour the stabilisation issuance rule (Section 6) is designed to bound.

5 The Full-Reserve Banking Layer

For the balance-sheet identity (4) to bind, commercial banks must not be permitted to create money by lending against deposits—the fractional-reserve channel that both Fisher (1936) and Chapra (1985) reject and that Werner (2014) documents empirically. Under the κ -Standard banks split into two rigid, interest-free silos.




Figure 2: The κ -Standard money base as a *realised* series, not a comparative static: $M(t) = \sum_i N_i e^{-\kappa_i(t)(1-\delta)T}$ with N_i the real outstanding USD-sukuk volume per issuer and $\kappa_i(t)$ the real constant-maturity κ (Cbonds, $\delta = 0.60$; rolled to a fixed tenor — 7 y where quoted, with Türkiye’s series its 5 y constant-maturity κ). Five core sovereigns, \$82.3 bn par, over their common quoting window. The base falls as credit deteriorates and recovers as spreads compress — a 6.3% swing from credit alone.

Definition 5.1 (Two-silo full-reserve bank). A commercial bank offers exactly two account types.

1. **Demand deposits (custody, 100% reserve).** Balances are held one-for-one in base money M and are never lent. The bank charges a flat, transparent service fee (*ujrah*) for safekeeping and payment processing. No money is created.
2. **Investment accounts (*mudarabah* pools).** Capital seeking a return is pooled and deployed by the bank, as agent, into κ -priced credit instruments (corporate sukuk, the Islamic credit default swap of [Ahmed & Bhuyan \(2026c\)](#)). Returns float with the realised convergence intensity of the holdings; principal is not guaranteed and bears equity (profit-and-loss-sharing) risk.

Proposition 5.1 (Unit money multiplier and run-freeness). *Under Definition 5.1 the money multiplier is one—total money equals the base M_t of (4), with no bank-created inside money—and the system is free of par-redemption (demand-deposit) runs: demand deposits are fully reserved and therefore always redeemable, while investment pools are floating-value equity claims that carry no redeemable-at-par promise and so admit no discontinuous par run.*

Proof. Demand deposits are 100% reserved, so they create no money beyond the base they hold; investment pools transfer existing base money into asset purchases without creating deposits. Hence total money = M_t and the multiplier is unity. Demand deposits are redeemable on demand from full reserves, eliminating the maturity-mismatch that generates runs; investment pools have no par-redemption promise, so adverse information triggers price adjustment, not a run ([Cochrane, 2014](#)). $\square$

The model eliminates the *demand-deposit* run—the classic maturity-mismatch panic—by full reserving. It does not abolish all run-like dynamics: a floating-NAV pool facing correlated redemptions can still force asset sales into a falling market, and if pool units are made redeemable intraday a slow-motion first-mover advantage can re-emerge through fire-sale price impact

(Borio, 2014). What the structure removes is the *par* promise that converts such pressure into a discontinuous run; the residual risk is mitigated and made explicit as equity loss, not contractually denied.

The two-silo structure realises, in an interest-free form, the run-free system of Cochrane (2014), the narrow-banking design surveyed by Pennacchi (2012), and the money/credit separation of Fisher (1936) and Benes & Kumhof (2012): money is state-issued equity; credit is intermediated only through risk-sharing pools that buy κ -priced instruments. The standard objection—that full reserving starves the economy of credit—does not bind here: as Brunnermeier & Niepelt (2019) show, a shift from private deposits to public money is allocation-neutral when the issuer recycles the funds to intermediaries, which is exactly what the *mudarabah* pools do when they purchase κ -priced sovereign and corporate claims.

What the full-reserve literature predicts—and what it does not yet show. The case for separating money from credit is not merely structural; it has been quantified. In the IMF DSGE evaluation of the Chicago Plan, Benes & Kumhof (2012) report that moving a fractional-reserve economy to 100% reserves (i) eliminates bank runs by construction, (ii) sharply reduces both public and private debt, since money no longer has to be borrowed into existence, (iii) smooths the business cycle by removing the pro-cyclical bank-money multiplier that Werner (2014) and Jakab & Kumhof (2015) document, and (iv) delivers a steady-state output gain on the order of 10% in their calibration, with inflation able to fall to zero without deflationary debt spirals. The sovereign-money programme of Dyson et al. (2016) reaches the same qualitative conclusions from an institutional rather than a model-based standpoint. We import these as the *theoretical* case for the κ -Standard’s banking layer and are deliberate about their epistemic status: unlike the debt-money critique of Section 1, which rests on realised central-bank and BIS evidence, the Chicago-Plan gains are *model-supported*, not historically observed—no economy has run a full-reserve regime at scale in the modern era. The κ -Standard inherits both the promise and that evidentiary gap; what it adds is a concrete interest-free pricing kernel (κ , not r) under which the money/credit separation can actually be implemented, rather than a reserve mandate left to sit on top of an interest-based system.

6 Stabilisation: a Countercyclical Issuance Rule

Corollary 4.4 exposes the architecture’s central hazard. The standing money base contracts automatically when the system-wide convergence intensity $\bar{\kappa}$ widens, and $\bar{\kappa}$ widens precisely in a credit downturn—so the revaluation channel tightens money exactly when the economy weakens. This is the procyclicality that Borio (2014) identifies as the core danger of any asset-linked monetary aggregate. A monetary standard that did nothing about it would be inferior to discretionary fiat. We show the κ -Standard does not have to leave it unaddressed: the second channel of Theorem 4.3—new issuance—is a policy lever that can be run countercyclically against the first.

6.1 Decomposition and the neutralising rule

Write the expected base-growth rate from Theorem 4.3 as the sum of a market-driven revaluation drift and a policy-driven issuance flow:

$$\frac{\mathbb{E}_t[dM_t]}{dt} = \underbrace{D_t^{\text{rev}}}_{\text{market}} + \underbrace{\varphi_t}_{\text{policy}}, \quad D_t^{\text{rev}} = \sum_i N_i P_t^{(i)} \left[\beta_i \alpha_i (\kappa_t^{(i)} - \bar{\kappa}^{(i)}) + \frac{1}{2} \beta_i^2 \nu_i^2 \kappa_t^{(i)} \right], \quad (11)$$

where $\varphi_t = \mathbb{E}_t[dM_t^{\text{iss}}]/dt$ is the rate at which the Treasury monetises new real capacity (Proposition 3.1). The revaluation drift D_t^{rev} carries the procyclical contraction of Corollary 4.4; the issuance flow φ_t is chosen by policy.

Proposition 6.1 (Countercyclical drift neutralisation, when feasible). *Fix a target base-growth rate $g \geq 0$. Whenever the required issuance flow is feasible, $\iota_t^* \in [-S_t, K_t]$ (Section 6.2), the rule*

$$\iota_t^* = g M_t - D_t^{\text{rev}} \quad (12)$$

sets the conditional drift of the base to the target, $\mathbb{E}_t[dM_t]/dt = g M_t$, for every realisation of the convergence intensities. The procyclical revaluation drift is offset by countercyclical monetisation: a credit deterioration ($D_t^{\text{rev}} < 0$) calls for more issuance, a boom ($D_t^{\text{rev}} > 0$) for less.

Proof. Substitute (12) into (11): $\mathbb{E}_t[dM_t]/dt = D_t^{\text{rev}} + (gM_t - D_t^{\text{rev}}) = gM_t$. $\square$ $\square$

Two caveats are essential. First, the rule neutralises only the *drift*: the revaluation *diffusion* $-\sum_i N_i P_t^{(i)} \beta_i \nu_i \sqrt{\kappa_t^{(i)}} dW_t^{(i)}$ of Theorem 4.3 is untouched, so the base is held on a target path only in conditional expectation; its instantaneous variance, driven by credit-spread noise, remains. A monetary aggregate that fluctuates with dW is itself a concern we flag in Section 9. Second, ι_t^* is not a free control: it is the rate of *real-project monetisation* (Proposition 3.1) and is bounded by the backing constraint of Section 6.2, outside which the neutralisation is only partial. The rule is therefore a feasibility statement, not an unconditional guarantee. What it does cleanly is stabilise the *quantity* of money in expectation; it does not, and by Corollary 4.2 cannot, peg the *price of capital* $\bar{\kappa}$, which remains market-determined.

6.2 The backing constraint and the Sovereign Stabilisation Fund

The rule (12) is not unconstrained, and the constraint is exactly the discipline that defines the standard. By Proposition 2.1 every unit of issuance must be backed one-to-one by a real asset, so ι_t^* is feasible only within

$$\iota_t^* \in [-S_t, K_t], \quad (13)$$

where K_t is the present value of new real capacity available to monetise (the expansionary limit) and S_t is the inventory the state can sell or retire (the contractionary limit). In a severe stress, D_t^{rev} may be so negative that $\iota_t^* = gM_t - D_t^{\text{rev}}$ exceeds K_t : there are not enough shovel-ready projects to monetise fast enough, and the neutralisation is only partial.

This is what the *Sovereign Stabilisation Fund* (SSF) exists to bridge. The SSF is a pre-funded, fully-backed buffer of sovereign sukuk. In booms ($\bar{\kappa}$ compressing, $D_t^{\text{rev}} > 0$) it accumulates, absorbing the over-expansion; in busts it deploys, conducting expansionary κ -OMO (Proposition 4.1)—buying sukuk to raise prices and lower $\bar{\kappa}$ —and fast-tracking pre-approved real projects. The SSF makes S_t and the speed of K_t large enough to keep ι_t^* feasible across ordinary cycles, and it never breaks Proposition 2.1: its assets are real κ -priced sukuk, not unbacked claims.

To size it concretely, take the full real outstanding sukuk universe of Section 4.3 (\$185.8 bn base): a +100 bp system-wide $\bar{\kappa}$ shock contracts it by $\approx \$3.6$ bn. Holding the base flat ($g = 0$) through that shock therefore requires roughly \$3.6 bn of countercyclical monetisation or SSF deployment per 100 bp—a finite, budgetable quantity.

6.3 Bounded-but-backed versus unbounded-but-unbacked

This contrast between bounded backing and unbounded fiat is the heart of the matter. Conventional fiat stabilisation is *unbounded*: a central bank can always create more money, but each unit is a fresh debt or unbacked liability. The κ -Standard’s stabilisation is *bounded*: it can lean against the revaluation contraction only as far as real capacity and the SSF allow, but every unit it creates is backed one-to-one by a real asset. The boundedness is not a defect relative to the standard’s purpose; it *is* the purpose. Unbounded stabilisation is unbounded

precisely because it is unbacked, and unbacked issuance is the failure mode—debt accumulation, inflation, moral hazard—the architecture exists to remove. The κ -Standard thus makes an explicit and deliberate trade: it surrenders the ability to stabilise without limit in exchange for the guarantee that it can never stabilise without backing. Whether a given economy prefers that trade is a policy choice; that the trade is now *statable and quantifiable*—roughly \$3.6 bn per 100 bp on the full actively-quoted USD-sukuk cross-section (\$195.2 bn par; it would scale to roughly \$15 bn per 100 bp if the full \$761 bn global class were monetised)—is the contribution.

We pause to note how this backing constraint maps onto emerging settlement infrastructure, then return to the architecture’s central theoretical question — price-level determinacy — in Section 7.

6.4 Implementation on a tokenised unified ledger

Nothing in the κ -Standard requires distributed-ledger technology; it is a monetary architecture, not a protocol. But the architecture’s two operationally demanding objects—a money base defined as a live portfolio of κ -priced sovereign sukuk (Definition 2.2), and a Sovereign Stabilisation Fund whose issuance authority ι_t^* must be checked against a backing constraint $[-S_t, K_t]$ in real time (Section 6.2)—map unusually cleanly onto the *unified-ledger* substrate that the BIS (2023) blueprint proposes for the next monetary system: central-bank money, tokenised sovereign claims, and commercial-bank money as programmable objects in a common settlement environment. We sketch the correspondence without overstating it.

The base as tokenised sovereign claims. Under (4) base money is $M_t = \sum_i N_i P_t^{(i)}$, the marked value of held sukuk. On a unified ledger each N_i is a tokenised sovereign claim and M_t is the ledger’s running valuation of the ICB’s asset partition—the backing of Proposition 2.1 becomes a *ledger invariant* rather than a periodic accounting reconciliation. Issuance $\Delta M = \Delta N_i P_t^{(i)}$ and the κ -OMO of Proposition 4.1 are atomic settlement operations (delivery of base money versus the sukuk token), which is exactly the atomic settlement the blueprint identifies as the ledger’s advantage over today’s sequential messaging.

The SSF backing constraint as programmable money. The countercyclical rule of Proposition 6.1 is feasible only inside $\iota_t^* \in [-S_t, K_t]$. On a programmable ledger that admissibility condition is a *smart-contract precondition* on the issuance transaction: the SSF’s drawable buffer S_t and the verified real-capacity headroom K_t are on-ledger state, and a proposed ι_t^* that violates the bound simply does not settle. This is the technical content of “bounded-but-backed”: the backing constraint is not a supervisory promise to be audited after the fact but a transaction-level invariant the substrate enforces—the programmability that Adrian & Mancini-Griffoli (2019) and BIS (2023) describe, used here to make *non-issuance* the default whenever backing is absent.

The retail layer and what tokenisation does not buy. A retail CBDC on the same ledger is the natural front end for the full-reserve deposits of Section 5: *ujrah* custody balances are central-bank liabilities held directly, and *mudarabah* pool units are tokenised floating-NAV equity claims whose prices update with the underlying κ -priced book—the unit money multiplier of Proposition 5.1 is then a ledger property, not a regulatory target. We are deliberately narrow about the gains. Tokenisation improves *settlement, transparency, and enforceability of the backing constraint*; it does *not* supply price-level determinacy (Section 9), it does not abolish the procyclicality of Corollary 4.4, and it does not make the SSF buffer larger. It changes how reliably the architecture’s invariants hold, not whether the architecture is sound. The substrate is an implementation convenience aligned with where wholesale infrastructure is already heading (Auer et al., 2023; BIS, 2023), not a premise of the result. Operationally, the same κ -pricing

and tokenised-settlement machinery can be exercised first at product scale — asset-linked instruments such as vault-secured tokenised gold and sukuk-yield baskets for institutional liquidity management — before any monetary-scale deployment, so the architecture’s invariants are battle-tested on hedging instruments long before they carry a money base.

6.5 Empirical illustration: interest-rate risk and the κ alternative

The two natural experiments of the recent cycle — the 2020 pandemic and the 2022 rate shock — make the difference between an interest-anchored and a κ -anchored system directly visible on data (full detail and replication in the empirical companion, [Ahmed & Bhuyan, 2026a](#)). Three findings bear on the architecture of this paper.

The primitive is a measurement, not a decision. Over 2015–2026 the effective federal funds rate is held unchanged in 47% of months and then moved in discrete committee steps (from 0.08% in 2021 to 5.02% in 2023 *by decree*), while the measured κ moves continuously (flat in only 4% of months). Their correlation is -0.76 : the policy rate is the discretionary *response*, and κ the contemporaneous *measurement*, of the same underlying risk. This is the empirical content of conducting policy by quantity against a measured backing rather than by a chosen rate.

The counter-cyclical reflex is automatic, not discretionary. In 2020 the interest system cut its policy rate from 1.58% (February) to 0.05% (April) — it *eased* as risk rose, a pro-cyclical injection of leverage into the crisis. Over the same window the measured κ rose from 0.044 to 0.088, so the κ -priced base $M = \sum_i N_i e^{-\kappa_i(1-\delta)T_i}$ *contracts* (-8.6% for a representative five-year backing) exactly as risk spikes — the countercyclical issuance rule of §6 operating automatically, with no committee and no lag, on real data.

The asset side carries no interest-rate duration. Because a κ -priced claim at $\iota = 0$ carries no discount factor in r (the Separation Theorem, [Ahmed & Bhuyan, 2026b](#)), its rate sensitivity $\partial \text{price} / \partial r$ is zero at every maturity. The 2022 shock makes the consequence vivid: repricing par bonds at the realised 2021-low-to-2022-high Treasury curve gives a steep duration gradient for conventional, interest-priced instruments — from -4.5% at one year to -45% at thirty — whereas the κ -priced equivalent has *zero* rate-driven loss at every tenor, retaining only a far shallower exposure (one to two orders of magnitude smaller) to the realised hazard itself. A reserve asset priced in κ does not inherit the interest-rate risk that a single policy decision can impose on a bond-backed base; this is precisely the property a debt-free, equity-backed money requires of its backing.

Scale of the removed charge. Finally, the predetermined component this architecture removes is not small: at a $\sim 3.5\%$ effective rate on a $\sim \$315$ trillion global debt stock, the interest system routes on the order of \$11 trillion a year — roughly a tenth of world output — as a charge for time owed regardless of outcome. At $\iota = 0$ that predetermined transfer is zero (the long-short funding flow nets to zero); what remains is only the contingent risk premium, borne where real risk is borne. (These are illustrative magnitudes from public aggregates; the structural point — a levy on the stock versus a contingent premium on borne risk — is exact.)

7 Price-Level Determinacy: A Fiscal Closure

The balance-sheet identity (4) ties base money to asset *values*; it does not, by itself, pin the general price level. This is not a peripheral caveat but the architecture’s central theoretical question, and in its sharpest form it is a known trap. A system that issues money one-for-one against the market value of an asset book is, structurally, a *real-bills* regime, and real-bills regimes are the textbook example of price-level *indeterminacy*: if the issuer stands ready to monetise assets at their going price, both the nominal money stock and the price level can scale

together with only their ratio determined (Sargent & Wallace, 1982). A monetary-economics referee will, correctly, ask why the κ -Standard escapes this. This section answers.

The escape has two ingredients, both already present in the design: the backing is *real* (sovereign equity, not nominal debt), and policy is a *quantity* operation with no fixed-price convertibility (Section 4). Together they place the system squarely in the fiscal theory of the price level (Leeper, 1991; Sims, 1994; Woodford, 2001; Cochrane, 2023), under which a determinate price level follows. That a determinate price level need not rest on an active interest-rate rule—the property an interest-free system must rely on—is precisely the lesson Cochrane (2018) draws from stable inflation at the zero bound, and the sense in which, as Sims (2013) puts it, fiscal backing rather than interest is what gives money its value. We make this precise.

7.1 The real backing and the valuation equation

Let $P_t > 0$ denote the general price level (money per unit of the consumption good). The pricing equations (1) and (2) evaluate each sukuk in *real* terms, because the underlying $X^{(i)}$ is a claim on real output (Definition 2.1); write the resulting real value of one unit of sukuk i as $v_t^{(i)}$. Its nominal price is $P_t^{(i)} = P_t v_t^{(i)}$, and the identity (4) read in nominal terms becomes

$$M_t = P_t \sum_{i=1}^n N_i v_t^{(i)} = P_t \mathcal{V}_t, \quad \mathcal{V}_t \equiv \sum_{i=1}^n N_i v_t^{(i)}, \quad (14)$$

where $\mathcal{V}_t$ is the *real backing*—the total real value of the ICB’s book. Equation (14) is the κ -Standard’s valuation equation: the real value of the money stock M_t/P_t equals the real backing $\mathcal{V}_t$.

The backing has an explicit fiscal content. The sovereign sukuk are claims on the state’s real productive assets, which throw off a stream of real primary surpluses $\{s_u\}_{u \geq t}$ (real resource rents and net fiscal receipts, in goods units). Pricing this stream at $t = 0$ —so the *only* discounting is the convergence intensity κ , with no time-preference rate—gives

$$\mathcal{V}_t = \mathbb{E}_t^{\mathbb{Q}} \left[\int_t^{\infty} e^{-\kappa(u-t)} s_u du \right]. \quad (15)$$

Lemma 7.1 (Finite, price-level-independent backing). *If $\kappa > 0$ and $\mathbb{E}_t^{\mathbb{Q}} \int_t^{\infty} e^{-\kappa(u-t)} |s_u| du < \infty$, then $\mathcal{V}_t$ in (15) is finite, strictly positive whenever surpluses are positive on a set of positive measure, and is measurable with respect to real fundamentals alone—it does not depend on P_t .*

Proof. Finiteness is the stated integrability hypothesis, which $\kappa > 0$ secures whenever $\{s_u\}$ does not grow faster than κ in $\mathbb{Q}$ -mean; this is the monetary counterpart of the companion papers’ observation that $\kappa > 0$ makes the convergence time $\tau < \infty$ $\mathbb{Q}$ -a.s. and the valuation integral converge (Ahmed & Bhuyan, 2026a). Positivity is immediate. Price-level independence holds because s_u (real surpluses) and κ (a real intensity) are real objects, defined without reference to the numeraire. $\square$

7.2 Determinacy

The remaining ingredient is that the nominal money stock is a *predetermined state*, not a free variable. Under debt-free seigniorage (Definition 3.1) and κ -OMO (Proposition 4.1), M_t changes only through discrete quantity operations—issuance against newly monetised real projects, and open-market purchases or sales—each of which fixes a *nominal* increment at the date it is executed. The ICB does *not* offer to convert money into sukuk at a pegged nominal price; the price of the instrument, and hence κ , is market-determined (Corollary 4.2). We collect the conditions.

- (A1) *Real equity backing.* Base money is a claim on real sovereign surpluses priced at $\iota = 0$, so the real backing $\mathcal{V}_t$ satisfies Lemma 7.1: finite, positive, and independent of P_t .
- (A2) *Quantity regime, no fixed-price convertibility.* The *quantity instruments* are predetermined: holdings N_i change only by discrete seigniorage and κ -OMO operations chosen by policy, and the ICB pegs no nominal price. The nominal base $M_t = \sum_i N_i P_t^{(i)}$ itself also moves continuously through the market revaluation of the existing book (Theorem 4.3); what (A2) rules out is any *policy* channel other than discrete quantity operations — in particular any standing offer to convert at a fixed price. The determinacy argument below requires only that the ICB sets quantities, not prices.
- (A3) *No-bubble valuation.* $\kappa > 0$ and the integrability of Lemma 7.1 hold, ruling out a rational bubble in the real value of money.

Theorem 7.2 (Price-level determinacy). *Under (A1)–(A3), the equilibrium price level is unique and strictly positive,*

$$P_t = \frac{M_t}{\mathcal{V}_t} > 0, \quad (16)$$

and no self-fulfilling (sunspot) price path exists.

Proof. By (A2), $M_t > 0$ is predetermined; by (A1) and Lemma 7.1, $\mathcal{V}_t > 0$ is finite and P_t -independent. The valuation equation (14) then has the unique positive solution (16). For dynamic uniqueness, suppose a candidate equilibrium posts a price $P'_t \neq P_t$ at date t with M_t predetermined. Equation (14) requires $\mathcal{V}'_t = M_t/P'_t \neq \mathcal{V}_t$. But $\mathcal{V}_t$ is pinned by real fundamentals—surpluses and κ —and is invariant to the price level by (A1); it cannot jump to accommodate P'_t . Hence P'_t is not an equilibrium, and the transversality condition (A3) excludes explosive paths along which the real value of money diverges. The price level (16) is therefore locally unique. $\square$

7.3 Why determinacy is earned, not assumed

Two remarks locate the result against the standard theory and show that each of its two ingredients is load-bearing.

Remark 7.1 (The real-bills boundary). Replace (A2) by fixed-price convertibility—let the ICB stand ready to exchange money for sukuk at a pegged nominal price. Then M_t becomes endogenous, expanding elastically to meet demand at the peg, and M_t and P_t scale together: only the ratio $M_t/P_t = \mathcal{V}_t$ is determined, and the price level is indeterminate. This is exactly the real-bills indeterminacy of Sargent & Wallace (1982). Determinacy is thus a consequence of the quantity-operation design of Section 4—the ICB sets the quantity and lets the market set κ —and would be lost under a price peg. It is earned by the architecture, not assumed.

Remark 7.2 (Equity backing is what makes the fiscal closure bite). In the conventional fiscal theory the anchor is nominal government *debt*, whose real value B_t/P_t falls with inflation; determinacy emerges as the fixed point of an implicit valuation equation. Here the anchor is a real *equity* claim whose real value $\mathcal{V}_t$ is price-level-independent (Lemma 7.1), so the valuation equation is solved *explicitly* for P_t in (16) rather than implicitly. The κ -Standard's defining feature—base money is sovereign equity, not debt (Proposition 2.1)—is precisely what turns the fiscal valuation equation into a determinate price level.

There is a deeper structural point. The fiscal theory is usually *one* of two admissible regimes: Leeper (1991) shows the price level is determinate either under an active interest-rate rule with passive fiscal policy (the conventional monetarist/New-Keynesian case) or under passive money with active fiscal policy (the fiscal-theory case). An interest-free monetary system has

no interest-rate instrument—there is no policy rate to make active. It is therefore confined, by construction, to the passive-money/active-fiscal regime. For the κ -Standard the fiscal theory is not an optional closure chosen among alternatives; it is the *only* determinacy theory available, and the architecture lands in exactly the regime where a determinate price level is fiscally pinned. Removing the interest rate does not leave determinacy unanswered—it selects which of Leeper’s two answers applies.

7.4 Scope of the claim

Theorem 7.2 establishes price-level determinacy at the level the fiscal theory itself establishes it: a valuation equation that pins P_t given a predetermined nominal stock, a real backing, and a no-bubble condition. Three extensions take this further, and each is developed in the companion general-equilibrium paper (Ahmed & Bhuyan, 2026e), which we summarise here. First, surpluses $\{s_u\}$ are exogenous in Theorem 7.2; the companion paper endogenises them with a fiscal reaction function and shows in closed form that the price level is determinate iff the fiscal response satisfies $\gamma_s < e^{\bar{\kappa}} - 1$ —a *riba*-free analogue of Leeper’s active-fiscal/passive-money threshold, with $\bar{\kappa}$ in the interest rate’s place—of which Theorem 7.2 is the $\gamma_s = 0$ corner. Second, κ is exogenous to P_t here; the companion paper builds the full general equilibrium with a micro-founded interest-free money demand and an endogenous convergence intensity, and finds determinacy robust across a reduced-form credit feedback, a structural Bernanke–Gertler–Gilchrist micro-foundation, and a Tobin’s- Q revaluation channel. Third, beyond local determinacy it characterises the one regime in which the architecture can fail—a large credit shock against an exhausted Sovereign Stabilisation Fund—with an occasionally-binding model that produces fire-sale amplification and quantifies the backing buffer adequate to prevent it. What changes here relative to a bare accounting identity is already decisive: the price level is no longer free. Under an explicit and, for an interest-free system, *forced* fiscal closure, it is pinned by (16); the companion paper shows the pinning survives full general equilibrium.

8 Distributional Incidence: κ -Issuance versus Debt-Money

The architecture so far is distributional as much as monetary: it creates money debt-free (Section 3), assigns the seigniorage to the state rather than to private banks (Section 5), and ties issuance to real output by an identity (Proposition 3.1). This section makes that content explicit by asking the question the *riba* result does not answer: *who receives the gains from money creation, and from the productivity growth monetary policy accommodates?*

8.1 The Cantillon channel and its two ingredients

New money is never distributionally neutral, because it does not enter everywhere at once. Whoever receives it first transacts at pre-injection prices; by the time the injection has diffused and raised the price level, later recipients and holders of existing money have ceded purchasing power to the first recipient. This is the Cantillon effect (Cantillon, 1755). It has exactly two ingredients: a *sequence* (a first recipient) and a *slippage* (money growth in excess of real output, which raises the price level). The redistribution is their product—a first recipient with no slippage moves no purchasing power, because the new money is absorbed by new goods. We formalise this in the κ -Standard’s own valuation equation (14), $P_t = M_t/\mathcal{V}_t$.

Proposition 8.1 (Cantillon redistribution under the tether). *Let new base money ΔM be issued to a first recipient against an increase $\Delta \mathcal{V}$ in the ICB’s real book. Under the valuation equation (14), the real purchasing power transferred from holders of pre-existing money to the*

first recipient is, to first order in $\Delta M/M$,

$$R = \mathcal{V} \left(\frac{\Delta M}{M} - \frac{\Delta \mathcal{V}}{\mathcal{V}} \right) = \mathcal{V} \hat{\pi}, \quad (17)$$

the real money stock times the induced inflation rate $\hat{\pi} = \Delta P/P$. Under κ -monetisation (Definition 3.1, Proposition 3.1) issuance is matched to new real backing, $\Delta M^{\text{iss}} = \Delta \mathcal{V}$, so $\hat{\pi} = 0$ and $R = 0$: new money issued against new κ -priced real capacity transfers no purchasing power from existing holders. Any residual R is proportional to the κ -mispricing $\Delta M - \Delta \mathcal{V}$.

Proof. Before the injection the real value of the money stock is $M/P = \mathcal{V}$ by (14). After it the stock is $M + \Delta M$ with backing $\mathcal{V} + \Delta \mathcal{V}$, so $P' = (M + \Delta M)/(\mathcal{V} + \Delta \mathcal{V})$ and the pre-existing nominal stock M has real value $M/P' = M(\mathcal{V} + \Delta \mathcal{V})/(M + \Delta M)$. The transfer is the loss to those holders,

$$R = \mathcal{V} - \frac{M(\mathcal{V} + \Delta \mathcal{V})}{M + \Delta M} = \frac{\mathcal{V} \Delta M - M \Delta \mathcal{V}}{M + \Delta M} \approx \mathcal{V} \left(\frac{\Delta M}{M} - \frac{\Delta \mathcal{V}}{\mathcal{V}} \right),$$

and $\Delta P/P = \Delta M/M - \Delta \mathcal{V}/\mathcal{V}$ follows by differentiating $\log P = \log M - \log \mathcal{V}$. Setting $\Delta M^{\text{iss}} = \Delta \mathcal{V}$ from (5) gives $R = 0$. $\square$

Proposition 8.1 is the distributional reading of the real-output tether. The tether is usually stated as a quantity discipline (money grows with output); it is equally the condition that drives the Cantillon redistribution to zero. The new money has a first recipient—the project whose capacity is monetised—but because the injection is exactly absorbed by the new backing, the price level is unchanged and no existing holder is diluted. Redistribution returns only to the extent that κ misprices the project ($\Delta M \neq \Delta \mathcal{V}$), and because κ is read from markets rather than administered (Section 4), that slippage is bounded by measurement error, not by discretionary issuance.

The accounting result of Proposition 8.1 is promoted to a general-equilibrium theorem in the companion paper (Ahmed & Bhuyan, 2026e): in the full DSGE the same equity-backing identity $M/P = \mathcal{V}$ that forces the fiscal closure (and hence price-level determinacy) also drives the steady-state Cantillon transfer to zero, so distributional neutrality is there a *corollary of determinacy* rather than a separate assumption—the equation that pins the price level is the equation that zeroes the redistribution. The remainder of this section reads the same content against the incumbent debt-money system, where the corresponding transfer is positive by construction.

8.2 The debt-money benchmark

The incumbent system distributes the same flows through three channels that are regressive by construction.

1. **Interest.** Debt-money requires interest service—a transfer from net debtors, disproportionately median households, to net creditors. The κ -Standard removes this at the root: the base earns and pays no interest ($\iota = 0$).
2. **Private seigniorage.** Commercial banks create most broad money by lending and capture the creation rent as franchise value. Under full reserves (Section 5) the seigniorage is socialised; the $\Delta M = P^{(*)}$ of Definition 3.1 accrues to the Treasury, not to bank equity.
3. **Asset Cantillon with unbounded slippage.** New money enters through credit and asset purchases and lands on *existing* assets—housing, secondary-market equities—not new output. In the notation of (17), ΔM is untethered to $\Delta \mathcal{V}$: money growth is set by credit demand and by an explicit positive-inflation mandate, so $\hat{\pi} > 0$ systematically and R accrues to asset-holders. This is empirically documented rather than hypothetical: the Bank of England’s

own analysis finds that the asset-price gains from post-2008 quantitative easing accrued disproportionately to the wealthiest households (Bell et al., 2012); the BIS attributes the post-crisis rise in wealth inequality chiefly to QE-driven equity-price appreciation (Domanski et al., 2016); the pattern recurs in Japanese unconventional policy through the portfolio channel (Saiki & Frost, 2014); and the survey evidence concurs that unconventional easing tends to widen wealth inequality through asset prices (Colciago et al., 2019). The κ -Standard replaces this with the tethered issuance of Proposition 8.1.

Channel	Debt-money (incumbent)	κ -Standard
Interest service	debtors $\rightarrow$ creditors	none ($\iota = 0$)
Seigniorage	private bank franchise	public (Treasury)
New money lands on	existing assets	new κ -priced output
Aggregate transfer R	$\mathcal{V}\hat{\pi}$, $\hat{\pi} > 0$ by mandate	0 up to κ -mispricing
First recipient	borrowers, dealers	monetised-project owners
Tech/efficiency dividend	suppressed $\rightarrow$ asset inflation	$\hat{\pi} < 0$ to all holders [†]

[†] only if the standard forgoes a positive-inflation target; see below.

8.3 The two levers the architecture does not settle

Proposition 8.1 should not be over-read. The tether sets the *aggregate* redistribution to zero, but it does not abolish the Cantillon sequence: the issuance rate $\varphi_t = \mathbb{E}_t[\mathrm{d}M_t^{\text{iss}}]/\mathrm{d}t$ remains a spigot, and the monetised project is still the first recipient. Two questions the architecture enables but does not answer decide whether the realised system is fairer.

First, the inflation objective. The most regressive feature of the incumbent regime is not interest but the positive-inflation mandate, which suppresses the deflation that technology and efficiency would otherwise produce and routes the accommodating money into assets (Domanski et al., 2016). The κ -Standard targets a quantity—the tether and the κ -OMO of Section 4—rather than a price index. If it forgoes an inflation target, productivity growth ($\hat{g}_V > \hat{g}_M$) appears as $\hat{\pi} < 0$ in (17): falling prices, hence $R < 0$, a real gain distributed to *every* holder of money rather than to owners of appreciating assets. This is the single largest distributional difference available—larger than the removal of interest—and we flag it as a design commitment the standard should make explicit rather than leave implicit in the quantity rule. To gauge the order of magnitude: an advanced economy running a 2% inflation target against a 1–1.5% trend productivity growth forgoes roughly 3–3.5 percentage points per year of money-base purchasing power—a flow the incumbent regime routes into asset prices (and, by the evidence above, disproportionately to the top wealth decile) but that a price-stable κ -quantity rule would instead distribute uniformly to every money holder. Compounded over a decade this dwarfs the one-off elimination of the interest wedge; the deflation dividend, not the removal of *riba*, is where the bulk of the distributional gain lies.

Second, the governance of monetisation. With the aggregate transfer tethered to zero, the residual question is *who may be a monetisation counterpart*: whose projects the ICB monetises, and on what terms. If access is broad and rule-based—any productive project priced at a transparent, market-read κ —the channel is far more democratic than discretionary credit and asset purchases. If monetisation is captured by politically connected developers—the classic failure mode of state-directed credit—the Cantillon sequence delivers its first-recipient advantage to a connected few, and the system can be as regressive as the one it replaces, with state capture substituted for financial capture. The market-observability of κ is the structural safeguard: fixing the monetisation *price* outside policy discretion removes the lever through which such capture usually operates, but it does not by itself constrain project *selection*.

Scope. The κ -Standard removes the *mathematical necessity* of the three regressive channels of debt-money: it eliminates interest, socialises seigniorage, and tethers issuance so that the aggregate Cantillon redistribution is zero up to κ -mispricing. It does not, by construction alone, guarantee a fair *realised* distribution—that turns on the inflation objective and the governance of monetisation, both political choices. “Interest-free” is thus necessary but not sufficient for “equitable”; the contribution here is to locate the remaining distributional content precisely—in the objective for $\hat{\pi}$ and the selection rule for φ_t —rather than in the pricing of money, which the *riba* result settles.

9 Limitations and Honest Boundaries

The κ -Standard is a monetary architecture, and like the Chicago Plan it is a design rather than an estimated model. We state its boundaries plainly; each is load-bearing for the credibility of the proposal.

It does not fix the price level. Proposition 2.1 guarantees that base money is backed one-to-one by the present value of real assets. It does *not* guarantee a constant price level. The backing is only as good as the κ -prices themselves: estimation error in κ , model error in the discount function D , and real shocks to the asset values $X^{(i)}$ all move M_t and hence the price level. Moreover, as Goodhart (1998) stresses, the acceptance and value of money depend on state (fiscal/tax) power, not on asset backing alone. The correct claim is narrow: the κ -Standard eliminates *debt-financed* and *unbacked* money creation and ties the base to real-asset value—not that it abolishes inflation.

Determinacy holds under a fiscal closure, and survives general equilibrium. The balance-sheet identity $M_t = \sum_i N_i P_t^{(i)}$ is nominal and, taken alone, pins only asset *values*, not the price level—and in its sharpest form it risks the real-bills indeterminacy of Sargent & Wallace (1982). Section 7 closes this: because the backing is real sovereign equity and policy is a quantity operation with no fixed-price convertibility, the fiscal valuation equation (16) delivers a unique, strictly positive price level (Theorem 7.2) — and, an interest-free system having no interest instrument to make active, this fiscal regime is its natural determinacy closure (Section 7). Theorem 7.2 takes surpluses and κ as exogenous; the companion general-equilibrium paper (Ahmed & Bhuyan, 2026e) endogenises both and shows determinacy survives. The determinacy claim is therefore established in partial equilibrium here and in full general equilibrium in the companion paper: the price level is no longer free.

κ -OMO requires credit-bearing assets.

Remark 9.1. The transmission of Proposition 4.1 relies on $\partial P/\partial \kappa < 0$, which holds only for credit-bearing sukuk. For a pure at-par perpetual ($\delta \rightarrow 1$), $P = X$ independently of κ , so open-market purchases move the price but not the implied intensity, and κ -OMO has no traction. The central bank’s policy portfolio must therefore consist of credit-spread-bearing sovereign sukuk, not *riba*-free at-par claims.

The automatic contraction is procyclical, and only boundedly correctable. Corollary 4.4 shows the base contracts when $\bar{\kappa}$ widens, a procyclical response of exactly the kind Borio (2014) identifies as the core hazard of asset-linked monetary aggregates. Section 6 addresses it with a countercyclical issuance rule (Proposition 6.1) that neutralises the revaluation channel and holds the base on a target growth path—but only within the backing constraint (13). The correction is therefore *bounded*: in a stress severe enough to exhaust real-capacity issuance and the Sovereign Stabilisation Fund, the contraction reasserts. We regard this boundedness as a

deliberate feature (Section 6), not a defect, but it is a genuine limit relative to unbounded fiat stabilisation, and quantifying the SSF buffer adequate for a target distribution of $\bar{\kappa}$ shocks is the principal empirical extension this paper opens.

State seigniorage is assumed, not proven optimal. We follow Al-Jarhi and the Chicago Plan in vesting money creation in the state. Selgin & White (1994) argue competitive private money provision can be stable and disciplined; we do not refute them. The κ -Standard is compatible with a single sovereign issuer or with competing issuers each running the same κ -pricing; choosing between them is a separate question.

Riba is resolved; possession is not. As throughout the programme, the κ construction resolves *riba* mathematically but leaves *qabd* (constructive possession) of the underlying real assets as contract-level work for a Shariah board, inherited unchanged from the companion papers.

The anchor is a priced (Q -measure) intensity, and its premium cannot be mutualised away. The empirical companion (Ahmed & Bhuyan, 2026a) separates the measured κ into a realised hazard κ^P and a priced κ^Q , the wedge between them being a risk premium. Two features of this architecture pin the monetary anchor to the priced side. First, κ -OMO has traction only on credit-spread-bearing assets (Remark 9.1), and a market-read spread is a κ^Q object. Second, the market-observability of κ is exactly the safeguard against monetisation capture identified in the distributional analysis above — fixing the monetisation price outside policy discretion means *reading* it from the market, i.e. at κ^Q . The base is therefore anchored to a premium-bearing intensity, and the residual it carries once $\iota = 0$ is precisely “the contingent risk premium, borne where real risk is borne” (§6.5). This is where the monetary layer parts company with the *takaful* and cooperative-guarantee instruments of the companion papers. There the wedge $\kappa^Q - \kappa^P$ can be priced out: a mutual pool charges the realised κ^P and returns any premium to members as surplus, which collapses those cases onto settled actuarial ground. A monetary base cannot mutualise its own unit of account; reading κ^Q from the market is a feature here (it is the anti-capture safeguard), not a defect, but it leaves the base carrying a market risk premium and therefore on the contested jurisprudential frontier the companion reduction identifies, not on the settled ground of *takaful*. Whether a fully contingent, predetermined-charge-free premium of that kind is permissible in the unit of account is the one open question this architecture inherits and does not itself settle.

10 Conclusion

Islamic monetary theory and the Western full-reserve tradition spent the better part of a century converging, independently, on a single conclusion: money should be created debt-free by the state and separated from interest-bearing credit. Both were stopped at the same point—the absence of a *riba*-free price at which the assets backing the money base could clear. This paper supplies that price. By elevating the convergence intensity κ from a pricing primitive to a monetary base, the κ -Standard makes money the present value of sovereign equity (Proposition 2.1), ties its growth to real output (Proposition 3.1), conducts policy through quantity while the market sets the price of capital (Corollary 4.2), and removes the demand-deposit run (Proposition 5.1). Its money base obeys a clean law of motion (Theorem 4.3) with a built-in contraction under credit-deterioration shocks (Corollary 4.4) that a countercyclical issuance rule neutralises within an explicit backing constraint (Section 6). Because the backing is real equity and policy is a quantity operation, the base escapes real-bills indeterminacy and pins a unique price level

under the fiscal theory of the price level—the natural determinacy closure for a system with no interest-rate instrument (Theorem 7.2).

We have been deliberate about what the architecture does not do: its price-level determinacy is established here under a fiscal closure with exogenous surpluses, with the general-equilibrium version—endogenous fiscal feedback and an endogenous convergence intensity—developed in the companion paper (Ahmed & Bhuyan, 2026e), which finds it survives; its stabilisation is bounded (it trades unbounded-but-unbacked fiat intervention for bounded-but-backed monetisation); and it is a design whose pricing primitive’s empirical reality rests on the authors’ companion papers rather than on new or independently replicated evidence here. What it offers is a candidate way to close a long-standing gap. The companion papers argue that interest is a removable parameter and that κ prices the instruments of an Islamic financial system on real data; this paper proposes that the same κ can serve as the clearing object for the monetary system above them, and the companion general-equilibrium paper shows the proposal survives a fully micro-founded model. The same κ thus supplies the riba-free market-clearing price that the Islamic and Chicago-Plan full-reserve programmes both required and neither had.

Acknowledgements

This research received no specific grant from any funding agency in the public, commercial, or not-for-profit sectors. All errors remain our own.

References

- Akerer, D., Hugonnier, J., & Jermann, U. (2025). Perpetual futures pricing. *Mathematical Finance*, 35, DOI: 10.1111/mafi.70018.
- Adrian, T., & Mancini-Griffoli, T. (2019). The rise of digital money. *IMF Fintech Note* 19/01.
- Ahmed, S., & Bhuyan, R. (2026a). From perpetual contracts to Islamic credit: the random stopping-time equivalence and a riba-free foundation for sukuk, takaful, and credit protection. *SSRN Working Paper* 6322858.
- Ahmed, S., & Bhuyan, R. (2026b). The κ -yield curve: construction, estimation methodology, and real-data evidence from sovereign sukuk. *SSRN Working Paper* 6322938.
- Ahmed, S., & Bhuyan, R. (2026c). An Islamic credit default swap on smart-contract infrastructure. *SSRN Working Paper* 6323519.
- Ahmed, S., & Bhuyan, R. (2026d). The κ -framework: a riba-free foundation for Islamic monetary systems. *SSRN Working Paper* 6845098.
- Ahmed, S., & Bhuyan, R. (2026e). Interest-free monetary equilibrium: A general-equilibrium foundation for the κ -Standard. *SSRN Working Paper* 6890259.
- Al-Jarhi, M. A. (1981/1983). A monetary and financial structure for an interest-free economy: institutions, mechanism and policy. In Ahmad, Iqbal & Khan (eds), *Money and Banking in Islam*. ICRIE/Institute of Policy Studies.
- Askari, H., Iqbal, Z., Krichene, N., & Mirakhor, A. (2012). *Risk Sharing in Finance: The Islamic Finance Alternative*. Wiley.
- Auer, R., Cornelli, G., & Frost, J. (2023). Rise of the central bank digital currencies. *International Journal of Central Banking*, 19(4), 185–214.

- Bell, V., Joyce, M., Liu, Z., & Young, C. (2012). The distributional effects of asset purchases. *Bank of England Quarterly Bulletin*, 52(3), 254–266.
- Bank for International Settlements (2023). Blueprint for the future monetary system: improving the old, enabling the new. *BIS Annual Economic Report 2023*, Chapter III.
- Benes, J., & Kumhof, M. (2012). The Chicago Plan revisited. *IMF Working Paper* 12/202.
- Borio, C. (2014). The financial cycle and macroeconomics: what have we learnt? *Journal of Banking & Finance*, 45, 182–198.
- Brunnermeier, M. K., & Niepelt, D. (2019). On the equivalence of private and public money. *Journal of Monetary Economics*, 106, 27–41.
- Buiter, W. H. (2007). Seigniorage. *Economics: The Open-Access, Open-Assessment E-Journal*, 1, 2007–10.
- Cantillon, R. (1755). *Essai sur la Nature du Commerce en Général* [Essay on the Nature of Trade in General]. Reprinted Liberty Fund, 2010.
- Chapra, M. U. (1985). *Towards a Just Monetary System*. The Islamic Foundation.
- Cochrane, J. H. (2014). Toward a run-free financial system. In *Across the Great Divide: New Perspectives on the Financial Crisis*. Hoover Institution Press.
- Cochrane, J. H. (2018). Michelson–Morley, Fisher, and Occam: the radical implications of stable quiet inflation at the zero bound. *NBER Macroeconomics Annual*, 32, 113–226.
- Cochrane, J. H. (2023). *The Fiscal Theory of the Price Level*. Princeton University Press.
- Colciago, A., Samarina, A., & de Haan, J. (2019). Central bank policies and income and wealth inequality: A survey. *Journal of Economic Surveys*, 33(4), 1199–1231.
- Domanski, D., Scatigna, M., & Zabai, A. (2016). Wealth inequality and monetary policy. *BIS Quarterly Review*, March, 45–64.
- Duffie, D., & Singleton, K. J. (1999). Modeling term structures of defaultable bonds. *Review of Financial Studies*, 12(4), 687–720.
- Dyson, B., Hodgson, G., & van Lerven, F. (2016). *Sovereign Money: An Introduction*. Positive Money.
- El-Gamal, M. A. (2006). *Islamic Finance: Law, Economics, and Practice*. Cambridge University Press.
- Fisher, I. (1936). *100% Money* (2nd ed.; first ed. 1935). Adelphi.
- Friedman, M. (1960). *A Program for Monetary Stability*. Fordham University Press.
- Goodhart, C. A. E. (1998). The two concepts of money: implications for the analysis of optimal currency areas. *European Journal of Political Economy*, 14(3), 407–432.
- Hall, R. E., & Reis, R. (2015). Maintaining central-bank financial stability under new-style central banking. *NBER Working Paper* 21173.
- El Hamiani Khatat, M. (2016). Monetary policy in the presence of Islamic banking. *IMF Working Paper* 16/72.

- Iqbal, Z., & Mirakhor, A. (2011). *An Introduction to Islamic Finance: Theory and Practice* (2nd ed.). Wiley.
- Jakab, Z., & Kumhof, M. (2015). Banks are not intermediaries of loanable funds—and why this matters. *Bank of England Working Paper* 529.
- Khan, M. S., & Mirakhor, A. (1987). *Theoretical Studies in Islamic Banking and Finance*. IRIS Books.
- Leeper, E. M. (1991). Equilibria under ‘active’ and ‘passive’ monetary and fiscal policies. *Journal of Monetary Economics*, 27(1), 129–147.
- McLeay, M., Radia, A., & Thomas, R. (2014). Money creation in the modern economy. *Bank of England Quarterly Bulletin*, Q1.
- Pennacchi, G. (2012). Narrow banking. *Annual Review of Financial Economics*, 4, 141–159.
- Saiki, A., & Frost, J. (2014). Does unconventional monetary policy affect inequality? Evidence from Japan. *Applied Economics*, 46(36), 4445–4454.
- Sargent, T. J., & Wallace, N. (1982). The real-bills doctrine versus the quantity theory: a reconsideration. *Journal of Political Economy*, 90(6), 1212–1236.
- Selgin, G., & White, L. H. (1994). How would the invisible hand handle money? *Journal of Economic Literature*, 32(4), 1718–1749.
- Simons, H. C. (1936). Rules versus authorities in monetary policy. *Journal of Political Economy*, 44(1), 1–30.
- Sims, C. A. (1994). A simple model for study of the determination of the price level and the interaction of monetary and fiscal policy. *Economic Theory*, 4(3), 381–399.
- Sims, C. A. (2013). Paper money. *American Economic Review*, 103(2), 563–584.
- Werner, R. A. (2014). Can banks individually create money out of nothing? The theories and the empirical evidence. *International Review of Financial Analysis*, 36, 1–19.
- Wicksell, K. (1898/1936). *Interest and Prices* (Geldzins und Güterpreise). Macmillan.
- Woodford, M. (2001). Fiscal requirements for price stability. *Journal of Money, Credit and Banking*, 33(3), 669–728.
- Woodford, M. (2003). *Interest and Prices: Foundations of a Theory of Monetary Policy*. Princeton University Press.
- Ahmed, S., & Bhuyan, R. (2026a). The κ -Framework Meets the Data: An Out-of-Sample Empirical Validation Across Seventeen Data Providers. Working paper.
- Ahmed, S., & Bhuyan, R. (2026b). The Separation Theorem: Interest as a Removable Parameter in No-Arbitrage Pricing. Working paper.

EDITORIAL NOTE TO THIS PRINTING

Paper 9 · Interest-Free Monetary Equilibrium

Working paper, revised July 2026 SSRN 6890259

Role in the programme. The general-equilibrium check on Paper 8. **What it establishes.** A determinate DSGE without an interest-bearing asset, Dynare-validated and disciplined against real sovereign data; time preference stays in preferences while the interest contract leaves the asset space. **Corrected or extended later.** The first-claim was hedged to “to our knowledge” in June 2026. **Not established here.** Long-horizon behaviour under heterogeneity; the paper holds it open and prices what it can (the welfare cost lands in Paper 14’s range).

Replication for every number in this paper: DOI 10.5281/zenodo.21290557 · github.com/shehzadahmed-xx/kappa-replication

Interest-Free Monetary Equilibrium: A General-Equilibrium Foundation for the κ -Standard

Shehzad Ahmed*

Rafiqul Bhuyan†

June 2026

Abstract

We build what is, to our knowledge, the first dynamic stochastic general-equilibrium model of a monetary economy that operates *without an interest rate*. Money is sovereign equity priced by a riba-free convergence intensity κ (the κ -Standard of the companion paper), credit is intermediated through profit-and-loss-sharing pools, and the central bank conducts policy by quantity rather than by a policy rate. The model retains time preference β in preferences while removing the interest *contract* from the asset space, so the intertemporal margin is an asset-pricing condition on state-contingent equity returns and money demand is the Fisherian wedge in *shadow* form. We establish price-level determinacy. Because the backing is real equity and policy is a quantity operation, the economy escapes real-bills indeterminacy. And because the equity-backing identity $M/P = \mathcal{V}$ pins money's value to the backing—so the quantity rule and the fiscal valuation are the same equation—the fiscal-theory regime is the *only* determinacy regime available to it (it is the equity-backing, not merely the absence of an interest rate, that forces this). The nominal determinacy condition is closed-form, $\gamma_s < e^{\tilde{\kappa}} - 1$, with the convergence intensity playing the role the interest rate plays in conventional fiscal theory; the partial-equilibrium determinacy result of the companion paper is its exogenous-surplus corner. Endogenous credit intensity introduces a real-block determinacy threshold that a backing-bounded countercyclical open-market rule restores, uniting the stabilisation and determinacy results of the companion paper. Optimal policy separates into a Friedman rule for liquidity and a maximal, backing-constrained credit-stabilisation rule. A structural Bernanke–Gertler–Gilchrist micro-foundation of κ shows the reduced-form determinacy hazard softens once credit conditions move through a net-worth state. The model is disciplined throughout by live sukuk-market data: the determinacy ceiling is mapped across the real Cbonds cross-section of 12 sovereign issuers (and a parallel corporate universe). We estimate the credit-shock process from the κ time series — its persistence and its size matter comparably, with the buffer *diverging* as persistence approaches a random walk — and derive a buffer-sizing rule that lets a central bank read its adequate Sovereign Stabilisation Fund off its own credit dynamics. Finally, the same equity-backing identity that forces the fiscal closure also zeroes the Cantillon redistribution of money creation: because money's value is pinned to real backing, issuance tethered to output transfers no purchasing power to its first recipients, so distributional neutrality is a *corollary* of determinacy rather than an added egalitarian assumption—and absent a positive-inflation mandate the efficiency dividend of productivity growth accrues to every money holder rather than to owners of appreciating assets. All theoretical results are cross-validated in Dynare.

JEL: E12, E43, E58, E63, D31, G12. **Keywords:** interest-free monetary policy, fiscal theory of the price level, determinacy, Islamic monetary economics, convergence intensity, financial accelerator, Cantillon effect.

*Independent University, Bangladesh.

†Alabama A&M University.

1 Introduction

Modern monetary models are organised around a nominal interest rate: it is the policy instrument, it appears in the consumption Euler equation, and a Taylor-type rule on it delivers a determinate price level. Remove the interest rate and three things break at once—the instrument, the intertemporal optimality condition, and the standard route to determinacy. Whether a coherent monetary general equilibrium can be built without any of them is not merely a theological curiosity: it is the unanswered question behind a century of proposals for an interest-free or fully-reserved monetary order. The Islamic full-reserve tradition [Al-Jarhi, 1981, Chapra, 1985] and, from a different premise, the Western Chicago-Plan tradition [Fisher, 1936, Benes & Kumhof, 2012] both wanted a debt-free, state-issued money base separated from interest-bearing credit; neither supplied a no-arbitrage price for the assets backing that base, because the natural candidate—an interest rate—is exactly what they sought to remove.

The companion paper [the “ κ -Standard”; Ahmed & Bhuyan, 2026a] closes that gap at the level of pricing. It proposes a monetary base built from sovereign equity priced by a *riba*-free convergence intensity κ —a stopping-time intensity that prices credit at zero interest—and proves price-level determinacy under a fiscal closure with *exogenous* surpluses and *exogenous* κ . That result is a partial-equilibrium statement: it pins the price level given a predetermined nominal stock and a real backing, but it does not specify where the surpluses, the credit intensity, or money demand come from. The natural and pressing question—whether the determinacy survives in a fully micro-founded general equilibrium with endogenous surpluses, endogenous credit, and a micro-founded money demand—is the subject of this paper. Answering it requires confronting head-on the fact that the usual scaffolding of monetary theory is unavailable, and rebuilding the intertemporal margin, money demand, and the determinacy argument without an interest rate.

Contribution. We make five contributions. *First*, we show that an economy with time preference but no interest *instrument* is internally consistent: β lives in preferences and the stochastic discount factor, the intertemporal margin is carried by state-contingent equity returns, and the nominal interest rate survives only as an untraded shadow wedge (Section 2). *Second*, we establish price-level determinacy in closed form: the nominal condition is $\gamma_s < e^{\bar{\kappa}} - 1$, with $\bar{\kappa}$ occupying the position the interest rate occupies in conventional fiscal theory [Leeper, 1991, Sims, 1994, Woodford, 2001, Cochrane, 2023], and the companion paper’s result is its exogenous-surplus corner (Section 4). *Third*, endogenous credit intensity creates a real-block determinacy threshold that a backing-bounded countercyclical κ -OMO restores, and optimal policy separates into a Friedman rule and a maximal credit-stabilisation rule (Sections 4–5). *Fourth*, a structural Bernanke–Gertler–Gilchrist micro-foundation disciplines the reduced-form hazard, showing it softens once credit moves through a net-worth state (Section 6). *Fifth*, we discipline the model with live sukuk-market data: we map the determinacy ceiling across the real Cbonds cross-section of sovereign and corporate issuers, estimate the credit-shock process from the κ time series—finding that persistence and size drive the required backing comparably, the buffer diverging as persistence nears a random walk—and derive a buffer-sizing rule for the Sovereign Stabilisation Fund (Sections 4, 7, 8). Every determinacy and welfare result is cross-validated in Dynare.

Related literature. The paper sits at the intersection of four literatures. The first is the *fiscal theory of the price level* and the determinacy literature [Leeper, 1991, Sims, 1994, Woodford, 2001, 2003, Leeper & Leith, 2016, Cochrane, 2023], from which we borrow the insight that a determinate price level can be delivered by fiscal backing rather than by an active interest-rate rule. Our contribution is to show that the fiscal theory is not merely one

admissible closure but the *forced* one for a monetary system with no interest instrument, and to derive the active/passive boundary in closed form with the convergence intensity in place of the interest rate. This perspective also turns the absence of a policy rate into a determinacy *advantage*: an economy that cannot run an active interest-rate rule is immune by construction to the self-fulfilling deflationary equilibria that such rules admit at the zero bound [Benhabib et al., 2001, 2002].

The second is *Islamic and risk-sharing monetary economics*. The aspiration to an interest-free monetary architecture is long-standing [Al-Jarhi, 1981, Chapra, 1985], and the recognition that an equity/risk-sharing system transmits policy directly to the real sector and may be more stable goes back to Khan & Mirakhor [1989] and is developed in the risk-sharing program of Askari et al. [2012]. Existing analyses of monetary policy without a conventional rate, however, are descriptive or partial-equilibrium—most directly the operational treatment of El Hamiani Khatat [2016]. We advance this literature to a micro-founded general equilibrium with a closed-form determinacy theorem and a globally-solved, welfare-ranked optimal policy; the Western full-reserve / Chicago-Plan tradition [Fisher, 1936, Benes & Kumhof, 2012] reaches a structurally similar debt-free base from different premises and inherits the same determinacy question.

The third is the *financial-frictions* literature. We micro-found the endogenous credit intensity with the external-finance premium of Bernanke et al. [1999], Carlstrom & Fuerst [1997], whose asset-price feedback descends from the collateral-constraint mechanism of Kiyotaki & Moore [1997]; our backing-bounded κ -OMO rule is the interest-free analogue of the intermediary-balance-sheet credit policy of Gertler & Karadi [2011]; the need for a global rather than local solution of the fire-sale region follows Brunnermeier & Sannikov [2014]; and the welfare rationale for sizing the Sovereign Stabilisation Fund is the pecuniary-externality / macroprudential argument of Bianchi [2011] and Jeanne & Korinek [2019], with the occasionally-binding constraint solved as in Guerrieri & Iacoviello [2015], Mendoza [2010], Kocherlakota [2000]. The fourth is *optimal monetary policy*: the classical-dichotomy structure follows the money-in-utility tradition of Sidrauski [1967], and the optimal liquidity prescription is the Friedman rule [Friedman, 1969, Chari et al., 1996, Lucas, 2000] expressed through a shadow rate no agent can contract on.

2 The model

Time is discrete; there is a representative household, a representative firm, a sovereign, and an Islamic Central Bank (ICB). There is no interest-bearing asset of any maturity.

2.1 The interest-free household

The household has money-in-utility preferences $\mathbb{E}_0 \sum_t \beta^t [u(C_t) + v(m_t)]$, $m_t \equiv M_t/P_t$, $\beta \in (0, 1)$. Time preference is retained: Islam prohibits a predetermined return on a loan, not the discounting of future utility. The asset menu is two equity claims: base money (a nominal store of value yielding liquidity services and real return $R_{t+1}^m = 1/\pi_{t+1}$) and shares in a *mudarabah* profit-and-loss-sharing (PLS) pool (real ex-dividend price q_t , state-contingent dividend d_{t+1} , return $R_{t+1}^a = (q_{t+1} + d_{t+1})/q_t$). With SDF $\mathcal{M}_{t+1} = \beta u'(C_{t+1})/u'(C_t)$, the first-order conditions are

$$1 = \mathbb{E}_t[\mathcal{M}_{t+1} R_{t+1}^a], \quad \frac{v'(m_t)}{u'(C_t)} = 1 - \mathbb{E}_t[\mathcal{M}_{t+1} R_{t+1}^m]. \quad (1)$$

The intertemporal margin is an asset-pricing condition on a *state-contingent* equity return, not a consumption Euler in a policy rate.

Proposition 1 (Interest-free money demand is well-posed). *Define shadow rates $1/R_t^f \equiv \mathbb{E}_t \mathcal{M}_{t+1}$ and $1/(1+i_t^{sh}) \equiv \mathbb{E}_t[\mathcal{M}_{t+1}/\pi_{t+1}]$. Then the opportunity cost of money in (1) equals $i_t^{sh}/(1+i_t^{sh})$ —the Fisherian wedge in shadow form. Money is held in positive quantity iff $i_t^{sh} > 0$, with satiation (the Friedman rule) at $i_t^{sh} = 0$. The shadow rate i_t^{sh} is a function of the SDF and inflation only; no instrument pays it. Removing the interest rate removes a price, not a margin.*

Proof. Substituting the definition of i_t^{sh} into the money-demand condition in (1), $\frac{v'(m_t)}{u'(C_t)} = 1 - \mathbb{E}_t[\mathcal{M}_{t+1}R_{t+1}^m] = 1 - \mathbb{E}_t[\mathcal{M}_{t+1}/\pi_{t+1}] = 1 - 1/(1+i_t^{sh}) = i_t^{sh}/(1+i_t^{sh})$, since $R_{t+1}^m = 1/\pi_{t+1}$. With $u'(C_t) > 0$ the right-hand side is strictly positive iff $i_t^{sh} > 0$; because v' is strictly decreasing with $v'(\infty) = 0$, there is a unique finite m_t solving it, and at $i_t^{sh} = 0$ the right-hand side is zero, so $v'(\bar{m}) = 0$ (satiation). Finally, i_t^{sh} is defined as a moment of $\mathcal{M}_{t+1} = \beta u'(C_{t+1})/u'(C_t)$ and π_{t+1} and enters no budget-constraint payoff; hence it is a shadow price, not a traded instrument. $\square$

This is the macroeconomic form of the program’s Separation Theorem: time preference β generates a positive shadow real return $R^f = 1/\mathbb{E}[\mathcal{M}]$ and, with inflation, a positive shadow nominal rate, but the only *traded* intertemporal price is the state-contingent equity return R^a . The interest rate is not assumed away; it is demoted from a policy instrument to an untraded wedge. It is precisely this demotion that confines the economy to passive money—there is no i to make active—and so, as Section 4 shows, to the fiscal-theory determinacy regime.

2.2 Firms and the κ cost-of-capital wedge

Output is $Y_t = A_t K_{t-1}^\alpha \bar{L}^{1-\alpha}$, capital $K_t = (1 - \delta_K)K_{t-1} + I_t$, investment funded only by the PLS pool. Pool claims resolve at the convergence intensity κ_t with recovery δ , priced at $\iota = 0$ as in the companion papers; the household-required return R^a obtains only if the firm’s net marginal product covers both time preference and the expected credit loss:

$$\mathbb{E}_t[A_{t+1}\alpha K_t^{\alpha-1} - \delta_K] = \underbrace{(R_t^f - 1)}_{\text{productivity / time value}} + \underbrace{\kappa_t(1 - \delta)}_{\text{expected credit loss}}. \quad (2)$$

Neither term is *riba*: the first is the marginal product of real capital (the return to a productive asset, which Islam permits), and the second is actuarial compensation for an expected loss, not a predetermined excess over principal. The convergence intensity κ_t is the wedge by which credit conditions raise the cost of capital: a rise in κ_t raises the required marginal product, lowers capital, and contracts output. This is the cost-of-capital channel through which the credit cycle reaches the real economy. In the steady state this wedge is the entire difference between the gross return on capital and the *riba*-free required return $1/\beta$.

2.3 Money as sovereign equity: the fiscal block

The sovereign runs a real primary surplus s_t . The consolidated sovereign-sukuk book—the ICB’s only asset and the backing of base money—is a claim on $\{s_{t+j}\}$ priced at $\iota = 0$: $\mathcal{V}_t = \mathbb{E}_t^\mathbb{Q}[e^{-\kappa t}(s_{t+1} + \mathcal{V}_{t+1})]$. Money’s real value equals the real backing,

$$\frac{M_t}{P_t} = \mathcal{V}_t, \quad (3)$$

the equilibrium upgrade of the companion paper’s accounting identity. The fiscal rule $s_t = \bar{s} + \gamma_s(\mathcal{V}_{t-1} - \bar{\mathcal{V}}) + e_t^s$ is the Leeper knob; the ICB sets the money *quantity* (no interest feedback), so the system is passive-money by construction.

2.4 Endogenous credit intensity

In the baseline, $\kappa_t = \bar{\kappa} + \psi_\kappa(-\hat{y}_t) + \xi_t$ (reduced form), where $\psi_\kappa \geq 0$ is the procyclical credit feedback and ξ_t an exogenous credit shock; Section 6 replaces this with a structural Bernanke–Gertler–Gilchrist net-worth block. The ICB’s countercyclical κ -OMO offsets the feedback, $\psi^{\text{eff}} \equiv \psi_\kappa - \phi_o$, with $\phi_o \leq \phi_o^{\text{max}}$ bounded by the Sovereign Stabilisation Fund (SSF).

3 Steady state and the interest-free cost of capital

With the four blocks specified, we evaluate the model at its deterministic steady state, where the interest-free structure of the cost of capital is most transparent. At that steady state $\bar{M} = \beta$, $R^a = 1/\beta$, and the dividend–price ratio is $\bar{d}/\bar{q} = (1 - \beta)/\beta$. The cost of capital (2) is

$$\text{MPK} - \delta_K = \left(\frac{1}{\beta} - 1\right) + \bar{\kappa}(1 - \delta) > 0, \quad (4)$$

positive and riba-free: a time-preference/productivity component plus an actuarial credit-loss component, with no interest term. Money demand satisfies $v'(\bar{m}) = u'(\bar{C})(1 - \beta/\pi)$, and the valuation block gives $\bar{m} = \bar{V} = \bar{s}/(e^{\bar{\kappa}} - 1)$, tying the money base to the surplus and $\bar{\kappa}$. The shadow nominal rate $i^{sh} = \pi/\beta - 1$ exists but is untraded.

4 Price-level determinacy

4.1 The nominal block: a closed-form region

A system that issues money against the value of an asset book is, structurally, a *real-bills* regime, and real-bills regimes are the textbook case of price-level indeterminacy [Sargent & Wallace, 1982]. The κ -Standard escapes this for two reasons already in the design: the backing is real equity (a price-level-independent real value), and policy is a quantity operation with no fixed-price convertibility.

To see this, hold the real block at its stationary values—the endowment core standard in establishing fiscal-theory determinacy—and linearise the backing recursion. With κ -discount factor $\beta_\kappa \equiv e^{-\bar{\kappa}}$ and the steady-state identity $\bar{V}(e^{\bar{\kappa}} - 1) = \bar{s}$, the linearised backing equation is $\hat{V}_t = -\hat{\kappa}_t + \beta_\kappa \mathbb{E}_t \hat{s}_{t+1} + \beta_\kappa \mathbb{E}_t \hat{V}_{t+1}$ (deviations scaled by $\bar{V}$). Substituting the fiscal rule $\hat{s}_t = \gamma_s \hat{V}_{t-1} + \hat{e}_t^s$, so that $\mathbb{E}_t \hat{s}_{t+1} = \gamma_s \hat{V}_t$, collapses the monetary–fiscal core to a single forward-looking valuation equation in one jump variable,

$$\hat{V}_t = \frac{\beta_\kappa}{1 - \beta_\kappa \gamma_s} \mathbb{E}_t \hat{V}_{t+1} - \frac{1}{1 - \beta_\kappa \gamma_s} \hat{\kappa}_t, \quad (5)$$

with the price level following residually from (3), $\hat{p}_t = \hat{M}_t - \hat{V}_t$.

Theorem 1 (Riba-free price-level determinacy). *The price level is locally determinate iff $\gamma_s < e^{\bar{\kappa}} - 1 \approx \bar{\kappa}$. The Blanchard–Kahn eigenvalue of the valuation block is $\lambda = e^{\bar{\kappa}} - \gamma_s$, and $\gamma_s < e^{\bar{\kappa}} - 1 \Leftrightarrow |\lambda| > 1$.*

Proof. Equation (5) is a scalar linear expectational difference equation in the single non-predetermined variable $\hat{V}_t$ with autoregressive-in-expectation coefficient $a \equiv \beta_\kappa/(1 - \beta_\kappa \gamma_s)$ and a stationary forcing process $\hat{\kappa}_t$. By Blanchard–Kahn, with one jump variable and no predetermined endogenous state, determinacy requires exactly one generalised eigenvalue outside the unit circle. Writing (5) as $\mathbb{E}_t \hat{V}_{t+1} = \lambda \hat{V}_t + a^{-1}(1 - \beta_\kappa \gamma_s)^{-1} \hat{\kappa}_t$ with $\lambda = 1/a = (1 - \beta_\kappa \gamma_s)/\beta_\kappa = e^{\bar{\kappa}} - \gamma_s$, the unique eigenvalue is λ . Determinacy is $|\lambda| > 1 \Leftrightarrow e^{\bar{\kappa}} - \gamma_s > 1 \Leftrightarrow \gamma_s < e^{\bar{\kappa}} - 1$, which also implies the regularity $1 - \beta_\kappa \gamma_s > 0$. When it holds the unique bounded




Figure 1: Nominal determinacy region $\gamma_s < e^{\bar{\kappa}} - 1$ (left) and the eigenvalue $\lambda = e^{\bar{\kappa}} - \gamma_s$ crossing unity (right), with real sovereign $\bar{\kappa}$ values from the companion yield-curve paper. The line $\gamma_s = 0$ is the companion paper’s Theorem.

solution is the forward sum $\hat{\mathcal{V}}_t = -(1 - \beta_\kappa \gamma_s)^{-1} \sum_{j \geq 0} a^j \mathbb{E}_t \hat{\kappa}_{t+j}$, pinning $\hat{\mathcal{V}}_t$ by fundamentals, and $\hat{p}_t = \hat{M}_t - \hat{\mathcal{V}}_t$ pins the price level given the predetermined nominal stock $\hat{M}_t$. $\square$ $\square$

A real-bills regime is the limiting case $\gamma_s \rightarrow \infty$ with fixed-price convertibility, in which $\hat{M}_t$ becomes endogenous and only the ratio $\hat{M}_t - \hat{p}_t = \hat{\mathcal{V}}_t$ is pinned: this is exactly the [Sargent & Wallace \[1982\]](#) indeterminacy that determinacy here is earned against.

The convergence intensity $\bar{\kappa}$ plays the role the interest rate plays in conventional fiscal theory; passive money is automatic (the eigenvalue is independent of the money-rule coefficient). This is a consequence of the equity-backing identity (3): because money’s real value is pinned to the backing, $M_t/P_t = \mathcal{V}_t$, an independent money-growth (active-money) anchor is not separately available—the quantity rule and the fiscal valuation are the *same* equation—so by [Leeper \[1991\]](#) the economy is confined to the passive-money/active-fiscal regime, and the fiscal closure is the *only* determinacy theory available to it. We stress the scope: it is the equity-backing of money, not the mere absence of an interest rate, that forces the fiscal closure. Absent the identity (3), a quantity-theoretic money-growth anchor would be the standard alternative; under it, the quantity rule and the price level are no longer the same object. We note the obverse of this elegance plainly: the fiscal theory is itself among the most contested closures in monetary economics, and an economy whose *only* anchor is fiscal-backing credibility has no second line of defence if that credibility fails — the determinacy result is conditional on the fiscal theory’s logic, and inherits its controversies.

Corollary 1 (Nesting the companion paper). *At $\gamma_s = 0$ (exogenous surpluses) $\lambda = e^{\bar{\kappa}} > 1$ for all $\bar{\kappa} > 0$, recovering $P_t = M_t/\mathcal{V}_t$ —the companion paper’s determinacy theorem—as the $\gamma_s = 0$ corner of the general-equilibrium region $\gamma_s \in [0, e^{\bar{\kappa}} - 1]$.*

The determinacy ceiling across the real sovereign-sukuk universe. The threshold $\gamma_s^* = e^{\bar{\kappa}} - 1$ is not a free parameter but a number every actual sovereign issuer pins down through its own credit spread. Computing $\bar{\kappa}$ for the full cross-section of sovereign US-dollar sukuk on Cbonds—78 bonds from 12 issuers, of which 60 with a positive quoted G-spread and maturity beyond 0.3 years enter the medians, with $\bar{\kappa}$ taken as the median G-spread over loss-given-default ($\delta = 0.60$)—gives an active-fiscal ceiling that is strictly positive for every sovereign and monotone in credit quality, from 0.72% for Qatar (AA) and 0.92% for




Figure 2: The active-fiscal determinacy ceiling $\gamma_s^* = e^{\bar{\kappa}} - 1$ for the full Cbonds cross-section of sovereign US-dollar sukuk (12 issuers, 78 bonds, of which 60 quoted enter the medians; $\delta = 0.60$). Each issuer lies on the determinacy boundary at its empirical $\bar{\kappa}$; the shaded region below is determinate (active fiscal). The ceiling is strictly positive and ordered by credit quality across the entire universe, from Qatar (AA) to Pakistan (CCC). Issuers with a single outstanding US-dollar sukuk (Qatar, Hong Kong, Philippines, Benin, Pakistan) are single-instrument and indicative; Saudi Arabia (12 bonds), Indonesia (20), and Bahrain (10) are densely covered.

Malaysia (A) through 1.8–2.0% for Saudi Arabia and Indonesia (investment grade) to 5.4% for Bahrain (B), 6.4% for Benin, and 9.1% for Pakistan (CCC); see Figure 2. Weaker-credit sovereigns admit a *wider* active-fiscal region, because their steeper κ -discount makes the real value of money less sensitive to the fiscal feedback. Two features of the estimator are worth noting. The cross-sectional $\bar{\kappa}$ is a current, mixed-tenor snapshot (median G-spread across an issuer’s outstanding sukuk), whereas Figure 1 uses the companion paper’s through-cycle constant-maturity series; the determinacy ceiling therefore moves with credit conditions. Türkiye illustrates this: its spreads have compressed substantially since the 2022–2024 stress, so it sits at a lower current $\bar{\kappa}$ (4.9%) than its crisis-inclusive long-run mean (7.4%). That time-variation is itself the point: the active-fiscal ceiling of Theorem 1 is an empirically meaningful, issuer-specific and state-dependent constraint, not an abstraction, and it is strictly positive and credit-ordered across the entire real sovereign-sukuk universe.

Sovereign versus corporate: the ceiling decomposes by credit. The determinacy ceiling inherits the additive structure of the credit spread. Computing $\bar{\kappa}$ for the Cbonds corporate US-dollar sukuk universe (213 bonds from the broader Cbonds corporate screener, median G-spread 105 bp, $\bar{\kappa} = 2.6\%$) and comparing within country, the corporate active-fiscal ceiling lies *above* the sovereign’s in every case, by the corporate credit premium—a median of +31 bp of spread, and as much as +115 bp for Türkiye (Figure 3). A monetary authority whose backing portfolio included corporate as well as sovereign κ -priced claims would therefore operate with a wider active-fiscal determinacy region, the extra width being exactly the corporate risk premium it bears. The determinacy condition thus decomposes as cleanly as the spread it is built from: ceiling $\approx$ sovereign credit + corporate credit.




Figure 3: Active-fiscal determinacy ceiling $\gamma_s^* = e^{\hat{\kappa}} - 1$, sovereign versus corporate US-dollar sukuk, for the six countries with both (Cbonds; $\delta = 0.60$). The corporate ceiling exceeds the sovereign's by the corporate credit premium in every country.

4.2 The real block: a credit-feedback threshold

With separable money-in-utility the model dichotomises: the real allocation is determined independently of the price level, and the nominal block of Theorem 1 takes $\hat{\kappa}_t$ as forcing. Endogenous credit enters the equity Euler through the cost-of-capital wedge, making the real block a 2×2 system in predetermined $\hat{k}_t$ and jump $\hat{c}_t$ whose transition $T(\psi^{\text{eff}})$ has a trace independent of ψ^{eff} and a determinant increasing in it.

Proposition 2 (Real determinacy threshold). *There is a threshold $\psi^* > 0$ such that the real block is determinate iff $\psi^{\text{eff}} = \psi_{\kappa} - \phi_o < \psi^*$. Above it both eigenvalues leave the unit circle and no bounded equilibrium exists—the procyclical credit loop is explosive. At the program calibration $\psi^* \approx 0.62$.*

Proof. The homogeneous real dynamics are a 2×2 system $(\hat{k}_{t+1}, \mathbb{E}_t \hat{c}_{t+1})' = T(\psi^{\text{eff}})(\hat{k}_t, \hat{c}_t)'$ with $\hat{k}_t$ predetermined and $\hat{c}_t$ a jump, where the credit feedback ψ^{eff} enters only the $(2, 1)$ entry of T through the cost-of-capital wedge $-(1 - \delta)\hat{\kappa}_t$ with $\hat{\kappa}_t = -\psi^{\text{eff}}\alpha\hat{k}_t$. Direct computation shows $\text{tr } T$ is independent of ψ^{eff} while $\det T$ is strictly increasing in ψ^{eff} (the relevant coefficient is $-a_{12}(\beta/\sigma)(1 - \delta)\alpha > 0$, since $a_{12} = -\bar{C}/\bar{K} < 0$). With the trace fixed, the smaller eigenvalue $\lambda_- = \frac{1}{2}(\text{tr } T - \sqrt{\text{tr } T^2 - 4\det T})$ increases monotonically in $\det T$ and crosses unity at the unique $\psi^{\text{eff}} = \psi^*$ solving $\det T = \text{tr } T - 1$; for $\psi^{\text{eff}} > \psi^*$, $\lambda_-, \lambda_+ > 1$ and the saddle path is lost. Evaluating T at the calibration of the table gives $\psi^* \approx 0.62$. $\square$

The procyclical loop is the general-equilibrium form of the auto-contraction that Borio [2014] identifies as the central hazard of asset-linked monetary aggregates: a credit shock raises κ , raises the cost of capital, contracts investment and output, and—if the feedback is strong enough—raises κ further.

4.3 Backing-constrained determinacy: uniting stabilisation and determinacy

The ICB cannot move the real allocation through the price level (dichotomy) but can lean on the market-clearing κ through κ -OMO, $\psi^{\text{eff}} = \psi_{\kappa} - \phi_o$, bounded by the SSF.




Figure 4: Backing-constrained real determinacy in (ψ_κ, ϕ_o) (left) and the amplified output response to a credit shock as $\psi^{\text{eff}} \rightarrow \psi^*$ (right). The companion paper’s backing rule extends the determinate region from ψ^* to $\psi^* + \phi_o^{\max}$.

Theorem 2 (Backing-constrained real determinacy). *Determinacy holds with no policy iff $\psi_\kappa < \psi^*$; with the countercyclical rule it is restorable iff $\psi_\kappa < \psi^* + \phi_o^{\max}$; beyond that frontier the SSF is too small. Full determinacy requires this and the nominal condition $\gamma_s < e^{\bar{\kappa}} - 1$.*

This is the general-equilibrium form of the companion paper’s “bounded-but-backed” stabilisation: monetary policy keeps the credit loop inside the determinate region, with reach bounded by the backing. The companion paper’s stabilisation result (its §6) and determinacy result (its §7) become one statement.

4.4 Dynare validation

The full nonlinear model is solved in Dynare. With four leads it has four forward-looking variables, so determinacy requires four explosive eigenvalues. The tool reproduces both thresholds in the correct direction: the baseline ($\psi^{\text{eff}} = 0.30$, $\gamma_s = 0.015$) returns four explosive roots (determinate); $\psi^{\text{eff}} = 0.70 > \psi^*$ returns five (“no stable equilibrium”—an extra real root); and $\gamma_s = 0.05 > e^{\bar{\kappa}} - 1$ returns three (“indeterminacy”—the nominal root falls below unity). Crossing the real threshold adds an explosive root; crossing the nominal threshold removes one.

5 Optimal policy

Because the model dichotomises, optimal policy separates into two non-interacting instruments.

Proposition 3 (Friedman rule). *Welfare in the liquidity margin is maximised at $i^{\text{sh}} = 0$, i.e. gross inflation $\pi^* = \beta$ (satiation in real balances). The ICB attains it by money growth $\mu = \beta$, despite the absence of an interest instrument.*

Proof. In the separable MIU block the household’s real-balance first-order condition equates the marginal utility of liquidity to the shadow opportunity cost i_t^{sh} ; since $u_m \geq 0$ with $u_m \rightarrow 0$ as real balances rise, welfare in this margin is non-decreasing in real balances and maximised at satiation $u_m = 0$, attained iff $i^{\text{sh}} = 0$ [Sidrauski, 1967, Friedman, 1969]. By the Fisher relation in shadow form $i^{\text{sh}} = \beta^{-1}\pi - 1$, $i^{\text{sh}} = 0 \iff \pi^* = \beta$; and from the money-market




Figure 5: Welfare loss falls as the credit loop is leaned against (left); the optimal rule $\phi_o^* = \phi_o^{\text{max}}$ holds ψ^{eff} below ψ^* up to the backing-constrained frontier (right).

equilibrium $\pi = \mu$ in the stationary equilibrium, so $\mu = \beta$ implements it. The dichotomy (Proposition 4 block) ensures this liquidity instrument does not interact with the real/credit margin, so the prescription is exact and instrument-independent. $\square$ $\square$

Proposition 4 (Credit stabilisation is a backing-bounded corner). *Stationary consumption volatility is strictly increasing in ψ^{eff} and diverges at ψ^* ; welfare loss is therefore strictly decreasing in ϕ_o , and the optimum is the corner $\phi_o^* = \phi_o^{\text{max}}$ —lean as hard against credit deterioration as the SSF backs. There is no interior optimum; the binding constraint is the backing.*

Proof. To second order, the consumption-equivalent welfare loss from fluctuations is $\frac{1}{2}\sigma \text{Var}(\hat{c})$. The stationary variance $\text{Var}(\hat{c})$, obtained from the discrete Lyapunov equation of the real-block saddle solution, is increasing in ψ^{eff} because a larger feedback raises the stable root toward unity (it diverges as $\psi^{\text{eff}} \rightarrow \psi^*$). Since $\psi^{\text{eff}} = \psi_{\kappa} - \phi_o$ is decreasing in ϕ_o , welfare loss is strictly decreasing in ϕ_o over the feasible range; with no convex cost on ϕ_o other than the backing constraint $\phi_o \leq \phi_o^{\text{max}}$, the maximiser is the corner $\phi_o^* = \phi_o^{\text{max}}$. Numerically the loss falls from 0.0257% of consumption at $\psi^{\text{eff}} = 0.57$ to 0.0094% at $\psi^{\text{eff}} = 0$, and Dynare order-2 theoretical moments reproduce $\text{std}(\hat{c})$ to three to four significant figures. $\square$ $\square$

The optimal rule is thus $\mu^* = \pi^* = \beta$ and $\phi_o^* = \phi_o^{\text{max}}$, acting on disjoint margins. Monetary policy in the κ -Standard is not price-level control (that is fiscal) but credit stabilisation under a backing budget. Welfare costs of departures are conventional: a money-growth deviation to 3%/5%/8% shadow rates costs 0.07/0.17/0.39% of consumption. The credit-stabilisation losses are a different and smaller margin—basis points of consumption—but the value of the rule at $\psi_{\kappa} > \psi^*$ is discrete, since there it restores a bounded equilibrium where none otherwise exists.

5.1 Distributional incidence: the inflation tax and the Cantillon channel

The backing identity that delivers determinacy also fixes the model's *distributional* incidence. Because $M_t/P_t = \mathcal{V}_t$ (3), every holder's real balances equal the real backing, and log-differencing gives the equilibrium inflation $\hat{\pi}_t = \hat{M}_t - \hat{\mathcal{V}}_t$. New money ΔM_t accompanied by new backing $\Delta \mathcal{V}_t$ therefore transfers purchasing power from existing holders to its first recipient by exactly the induced inflation,

$$R_t = \mathcal{V}_t(\hat{M}_t - \hat{\mathcal{V}}_t) = \mathcal{V}_t \hat{\pi}_t, \quad (6)$$

the general-equilibrium counterpart of the accounting Cantillon transfer of the companion κ -Standard paper, now with $\hat{\pi}_t$ the model's *endogenous* inflation rather than an exogenous slippage. Equation (6) is the channel through which the gains from money creation—and from the productivity growth that monetary policy accommodates—are allocated across holders.

Proposition 5 (Zero Cantillon transfer under the equity-backing closure). *In the passive-money/active-fiscal equilibrium (Theorem 1), base money is the value of the κ -priced book, so money created by monetising a new asset enters with matching backing: along the issuance margin $\hat{M} = \hat{V}$, and by (6) the issuance channel contributes $\hat{\pi} = 0$ and a zero transfer. The only residual inflation is the revaluation of the standing book under the mean-reverting credit intensity, which is zero-mean in the stationary distribution, so $\mathbb{E}[R_t] = 0$. Under the Friedman rule (Proposition 3, $\mu = \beta$) the marginal inflation-tax wedge is additionally zero. The regime thus attains the liquidity optimum and aggregate distributional neutrality simultaneously, up to κ -mispricing.*

Proof. The equity-backing identity (3) makes M_t the nominal value of the backing $P_t \mathcal{V}_t$; an issuance against a new asset of κ -priced value $\Delta \mathcal{V}_t$ raises both sides equally, so its contribution to $\hat{M}_t - \hat{V}_t$ vanishes and $R = 0$ on that margin by (6). The remaining variation in $\hat{M}_t - \hat{V}_t$ is the book revaluation driven by κ_t , whose innovation is zero-mean and mean-reverting, so $\mathbb{E}_t \hat{\pi}_{t+1} \rightarrow 0$ and the stationary mean transfer is zero. The Friedman-rule claim is Proposition 3: at the shadow satiation rate $i^{sh} = 0$ the inflation-tax marginal distortion u_m is zero. $\square$ $\square$

Why the active-money counterfactual cannot match it. Determinacy confined the architecture to passive money (Section 4): with no interest instrument, the quantity rule and the fiscal valuation are the *same* equation, $M/P = \mathcal{V}$. A conventional active-money regime instead sets money growth μ by a policy rule *independent* of the backing, so $\hat{M} - \hat{V} \neq 0$ systematically; with a positive-inflation mandate $\hat{\pi} > 0$ and $R > 0$ accrues one-directionally to the first recipients of new credit and reserves—the asset-price Cantillon channel that the empirical literature on post-2008 unconventional monetary policy identifies as a driver of wealth concentration. The equity-backing closure breaks this not as an add-on but as a corollary of (3): the equation that pins the price level is the equation that zeroes the Cantillon transfer. The single remaining lever is the inflation objective—under a quantity target with no positive-inflation mandate, productivity growth ($\hat{g}_V > \hat{g}_M$) appears as $\hat{\pi} < 0$ and $R < 0$, distributing the efficiency dividend to *every* money holder rather than to owners of appreciating assets. That this neutrality is a property of the same valuation equation that forces the fiscal closure—rather than an independent egalitarian assumption—is the paper's distributional result.

6 A structural credit intensity

So far the credit feedback ψ_κ has been a reduced-form parameter. We now micro-found it, and use the structural model to ask whether the determinacy threshold of Proposition 2 survives.

6.1 The external-finance premium and the net-worth state

Following the costly-state-verification logic of Carlstrom & Fuerst [1997] and Bernanke et al. [1999], the premium an entrepreneur pays on external funds is increasing in leverage. In our *riba*-free setting that premium *is* the credit intensity, so $\hat{\kappa}_t = \chi(\hat{\kappa}_t - \hat{n}_t)$, with net worth $\hat{n}_t$ a predetermined state accumulating retained returns. This delivers a micro-foundation:

the reduced-form feedback is recovered as the leverage projection $\psi_\kappa \approx \chi \cdot \text{lev}$, so the free parameter of Section 4 is the product of a credit-market primitive (the premium elasticity $\chi \approx 0.05$, [Bernanke et al., 1999](#)) and steady-state leverage.

6.2 Micro-foundation softens the determinacy hazard

What the structural exercise really does is move the Phase-3 threshold. That threshold required credit and output to move *within* the period. Routing κ through the *predetermined* net-worth state breaks the contemporaneous loop: in Dynare the structural model is determinate across every configuration tested—premium elasticity up to 5, leverage up to 8, survival down to 0.3—and a net-worth shock propagates as a hump (peaking near quarter 17) rather than amplifying contemporaneously (Figure 6). The Phase-3 explosive loop is thus disciplined as a reduced-form artifact.

6.3 Tobin’s Q and the revaluation channel

The one channel still absent is balance-sheet *revaluation*: with no capital adjustment costs Tobin’s $Q \equiv 1$, so asset prices cannot move net worth. We add a convex adjustment cost, making the price of installed capital a forward-looking jump $q_t = 1 + \phi(I_t/K_{t-1} - \delta_K)$; the return to capital then carries a capital-gains term, and net worth becomes a *leveraged* claim on the market value $q_t K_t$ in the standard [Bernanke et al. \[1999\]](#) form, $N_t = \theta_n[R_t^k q_{t-1} K_{t-1} - \beta^{-1}(q_{t-1} K_{t-1} - N_{t-1})]$, with the premium responding to market-value leverage $\hat{\kappa}_t = \chi(\hat{q}_t + \hat{k}_t - \hat{n}_t)$. The accelerator is now complete—a credit shock lowers q , which erodes the leveraged net worth faster than it erodes assets, raising leverage and κ further—and in the solved model q moves ($\text{std}(\hat{q}) \approx 0.67\%$), confirming the channel is active.

Proposition 6 (Determinacy is robust to the revaluation channel). *With Tobin’s Q and the leveraged net-worth law, the interest-free general equilibrium is locally determinate at every configuration tested, including leverage up to 10, premium elasticity up to 1.5, and adjustment cost as low as $\phi = 0.02$ (highly volatile q). The credit-shock output response is of the same order with and without Q and declines in the premium elasticity, the premium acting as a procyclical tax on capital.*

Proof (analytical structure, numerically verified). Augment the real block with Tobin’s Q and the leveraged net-worth law and linearise. The forward-looking states are consumption and q ; the predetermined states are capital k and net worth n . The Blanchard–Kahn count—number of unstable roots equal to the number of forward variables—is invariant to the premium elasticity and adjustment cost because those parameters enter only the predetermined-state transition, not the count of forward variables, while q remains pinned by the consumption Euler. Hence determinacy is a structural property of the dichotomy, not of the calibration; we confirm it numerically by computing the root count across the stated grid (leverage to 10, premium elasticity to 1.5, ϕ to 0.02), at every node of which exactly one unstable root obtains. $\square$

The result is intuitive: even with asset-price revaluation, the accelerator runs through the predetermined states k and n , while the forward-looking price q is pinned by the consumption Euler, so the Blanchard–Kahn root count is robust. The modest amplification is consistent with [Kocherlakota \[2000\]](#), who shows that the quantitative bite of standard net-worth models is small absent occasionally-binding constraints or fire-sale externalities—a separate, nonlinear agenda outside the local-determinacy scope of this paper.




Figure 6: A net-worth shock propagates as a hump (structural) rather than a contemporaneous impulse (reduced form) (left); the reduced-form feedback is micro-founded as $\psi_\kappa = \chi \cdot \text{lev}$ (right).




Figure 7: With Tobin's Q and the leveraged net-worth law, the interest-free economy is determinate across the entire (leverage, χ) space, including the region the contemporaneous reduced form labelled explosive (red dashed frontier) (left); the credit-shock output response with revaluation versus without is modest and declines in χ (right).

6.4 Synthesis

Taken together, Sections 4–6 give a clean account of credit and determinacy: a contemporaneous reduced-form feedback has a determinacy threshold ψ^* ; micro-founding the credit intensity through a net-worth state removes the within-period loop and yields robust determinacy with hump-shaped propagation; and the asset-price revaluation channel leaves determinacy intact, adding only modest amplification. The countercyclical κ -OMO rule is accordingly best read as the welfare/volatility instrument of Section 5, not as determinacy insurance.

7 The nonlinear frontier: occasionally-binding backing

The robustness of Sections 4–6 is a *linear*, normal-times statement. The one regime it cannot describe is the one the Sovereign Stabilisation Fund exists to prevent: the κ -OMO rule supports credit only down to the remaining backing buffer $S_t \geq 0$, so a large enough

credit shock *exhausts* the fund, the support cap binds, and the procyclical feedback runs unchecked. We close the model’s account of credit with a tractable nonlinear model of this occasionally-binding constraint, in two states—the output gap $\hat{y}_t$ and the buffer S_t :

$$a_t = \min(\phi_o \max(\kappa_t^{\text{des}}, 0), S_t), \quad \hat{\kappa}_t = \kappa_t^{\text{des}} - a_t, \quad S_{t+1} = \min(\bar{S}, S_t - a_t + \rho_S(\bar{S} - S_t)), \quad (7)$$

with $\kappa_t^{\text{des}} = [\psi_\kappa(-\hat{y}_{t-1}) + \xi_t]$ clipped to $[-\bar{\kappa}, \kappa_{\max}]$ (the intensity is non-negative and credit deterioration saturates once defaults occur), output $\hat{y}_t = \rho_y \hat{y}_{t-1} - b \hat{\kappa}_t$, and the normal-times feedback calibrated stable ($\rho_y + b\psi_\kappa = 0.94 < 1$) so that all nonlinearity comes from the binding constraint.

Proposition 7 (State-dependent fire sales and backing adequacy). *The response to a credit shock is state-dependent: the same 4% shock troughs at -0.80% when the SSF is full but -2.88% when it is depleted—a $3.6\times$ fire-sale amplification—and the ergodic output distribution is left-skewed (skewness ≈ -0.25). Crisis frequency and tail risk fall monotonically in the buffer: crises (output below -5%) occur with frequency $0.63\%, 0.17\%, 0.04\%, 0.00\%$ at buffers $\bar{S} = 1, 2, 3, 5\%$, with the 5th-percentile output gap rising from -2.98% to $+0.06\%$. Under transitory (i.i.d.) shocks a buffer of 3–5% of annual κ -support capacity drives the crisis frequency below 0.1% ; the shock scale itself is empirically grounded (the CIR stationary standard deviation of κ implied by the real sovereign-sukuk estimates of Ahmed & Bhuyan [2026b] has a cross-issuer median of 0.97% , close to the 1% used here). Once the shock’s persistence is estimated, the required buffer is substantially larger, as the next paragraph shows; the 3–5% figure is the transitory-shock benchmark, not the realistic one.*

This is the nonlinear realisation of “bounded-but-backed”: in normal times the constraint is slack and the economy is the linear one of Sections 4–6; only a large shock against an exhausted fund produces a fire sale, and the model maps the shock distribution into the backing needed to absorb it—a quantitative answer to the companion paper’s principal open question (Figure 8). The model is a tractable reduced form of the mechanism, not a global solution of the full DSGE; a global rational-expectations solve is provided below, and a piecewise-linear [Guerrieri & Iacoviello, 2015] or projection solve on the $(\hat{y}, S)$ state space is the natural quantitative deepening.

Empirical credit-shock persistence sharply raises the required buffer. The 3–5% figure assumes *transitory* (i.i.d.) credit shocks. The data say otherwise. Fitting an AR(1) to the real constant-maturity κ series of Ahmed & Bhuyan [2026b] gives a monthly persistence of $\rho_\xi \approx 0.77$ (a half-life of 2.7 months, within their reported 1.4–4.3) and an *innovation* standard deviation of 0.73% (the implied *stationary* standard deviation is $0.73\%/\sqrt{1 - \rho_\xi^2} \approx 1.1\%$, consistent with the 0.97% median quoted above). Persistence is the dominant driver of the tail: replacing the i.i.d. benchmark (1.0% s.d.) with the estimated AR(1) (0.73% innovation, $\approx 1.1\%$ stationary s.d.) raises the crisis frequency at a 3% buffer from 0.04% to 13.4% , because credit deterioration now *clusters* and drains the fund over successive months rather than in a single period (the effect is, if anything, larger at a common innovation size). The buffer required to hold the crisis frequency below 1–2% rises several-fold, from roughly 1–5% under i.i.d. shocks to 11–13% under the empirical persistence. This does not overturn the mechanism—larger backing still monotonically reduces crises—but it sharpens the boundedness lesson: a Sovereign Stabilisation Fund must be sized for *persistent* credit stress, and a realistically-sized fund reduces the frequency and depth of fire sales without eliminating them. The honest reading of the result is therefore more cautionary once the credit process is taken from the data.




Figure 8: *Left*: the output response to a credit shock is absorbed when the SSF is full and a deep fire sale when it is depleted ($3.6\times$ amplification). *Right*: crisis frequency and the 5th-percentile output gap versus SSF buffer capacity: larger backing monotonically lowers crisis frequency and tail risk, with an adequate buffer near 3–5% *under transitory shocks* (the persistence-corrected requirement is 11–13%; see text).

A buffer-sizing rule, and why the sovereign book is the binding case. The dependence on the credit process can be made operational. Mapping the model over a grid of persistence and volatility and recording, at each point, the buffer that holds crisis frequency below 1% yields a buffer-sizing function. A log–log fit over the simulated grid gives

$$\bar{S}^*(\rho_\xi, \sigma_\xi) \approx \sigma_\xi^{2.3} (1 - \rho_\xi)^{-2.2} \quad (\sigma_\xi, \bar{S}^* \text{ in percent}), \quad (8)$$

which reproduces the grid to within a few percentage points (it is an empirical fit over the relevant (ρ_ξ, σ_ξ) range, not a closed-form law; the fitted exponents drift with the window). Two features matter. Volatility and persistence enter with *comparable* elasticities ($\sigma_\xi^{2.3}$ versus $(1 - \rho_\xi)^{-2.2}$), so a given proportional change in shock size moves the required buffer about as much as the same proportional change in $1 - \rho_\xi$ — size is not second-order. What is distinctive about persistence is that the requirement *diverges* as $\rho_\xi \rightarrow 1$ (the near-random-walk limit), whereas volatility enters as a regular power (Figure 9). A central bank can read an indicative adequate fund off its own estimated credit dynamics.

Two empirical points anchor the rule. First, the sovereign sukuk process ($\rho_\xi = 0.77$, $\sigma_\xi = 0.73\%$) gives $\bar{S}^* \approx 12\%$. Second, we re-estimate the corporate process on a *much larger* sample than was previously available — constant-maturity κ series for **26 corporate issuers** (183 sukuk; Cbonds daily G -spread histories, 2018–2026), against the seven issuers of the earlier draft. It is *not* negligibly persistent: the per-issuer AR(1) coefficient has median $\rho_\xi = 0.84$ (mean 0.71, IQR 0.56–0.89) with $\sigma_\xi = 0.74\%$. Re-estimating the *sovereign* process by the *identical* 5-year constant-maturity method (six issuers) gives median $\rho_\xi = 0.87$ ($\sigma_\xi = 0.67\%$), statistically indistinguishable from the corporate 0.84: on a common footing the two credit processes are equally persistent.¹

¹The headline $\rho_\xi = 0.77$ used for the 12% sovereign buffer is the earlier five-issuer, headline-tenor estimate, retained for continuity; the like-for-like figure (0.87) is marginally higher, which would only raise the required buffer. A third, independent check on a freshly-pulled 2026 Cbonds panel—the four sovereign issuers with sufficient five-year constant-maturity history (Indonesia, Saudi Arabia, Sharjah, Bahrain), monthly AR(1)—gives a median persistence of $\rho_\xi \approx 0.71$ (half-life ≈ 2 months), *below* the headline and so implying, on that cut, a buffer no larger than the one adopted. Across all three samples persistence lies in the 0.71–0.87 range, bracketing the 0.77 used here, so the 12% figure is not an artifact of the original small sample.




Figure 9: The SSF buffer-sizing function $\bar{S}^*(\rho_\xi, \sigma_\xi)$ —the backing needed to hold crisis frequency below 1%—over credit-shock persistence and volatility. The required buffer rises steeply with persistence ($\propto (1 - \rho_\xi)^{-2.2}$). The empirically-estimated sovereign sukuk process (blue) requires $\approx 12\%$; the corporate process (green), re-estimated on 26 issuers ($\rho_\xi = 0.84$), is comparably persistent and requires a buffer of the *same order* (≈ 8 – 30% as the central ρ_ξ runs from the sample mean to the median).

Corporate persistence therefore implies a buffer of the *same order* as the sovereign’s (roughly 8–30% as the central ρ_ξ runs from the mean to the median), not the sub-one-percent that the small seven-issuer sample suggested: corporate and sovereign κ -processes are *comparably* persistent, so a κ -Standard backed by either needs a fund of the same order, and the buffer must be calibrated to each issuer’s own dynamics rather than assumed small for corporates. The estimate is, however, fragile in exactly the way the rule predicts: because $\bar{S}^*$ scales as $(1 - \rho_\xi)^{-2.2}$, near-random-walk persistence ($\rho_\xi \rightarrow 1$) makes the required buffer both large and sensitive to the persistence estimate.

A global solution validates the mechanism and the rule. The model above uses the mechanical backing rule of the companion paper. To confirm that the fire-sale mechanism and the backing-adequacy conclusion survive rational expectations—and to ask whether a forward-looking ICB would manage the buffer differently—we solve the ICB’s full dynamic program by value-function iteration on the $(\hat{y}, S)$ state space, with the planner choosing support to minimise a discounted crisis-averse loss subject to the occasionally-binding constraint $a_t \leq S_t$. The constrained-optimal policy lowers the planner’s loss by about 22% relative to the mechanical rule, almost entirely by removing a small mean bias; the two policies deliver *similar* crisis statistics (crisis frequency 0.047% versus 0.045%, fifth-percentile gap -1.19% versus -1.12% at $\bar{S} = 3\%$). The simple backing-bounded rule of Section 5 thus captures most of the crisis-prevention value, and the global solution confirms the fire-sale amplification and the backing-adequacy trade-off under rational expectations rather than overturning them (Figure 10). This global solve uses *transitory* (i.i.d.) credit shocks — its crisis frequencies ($\approx 0.04\%$ at $\bar{S} = 3\%$) are the i.i.d. benchmark above, so it confirms the mechanism and the mechanical-versus-optimal comparison in that regime. The persistence-driven buffer levels of equation (8) (11–13%) rest on the reduced-form AR(1) analysis and would require a global




Figure 10: Global rational-expectations solution. *Left*: the optimal support policy $a^*(\hat{y}, S)$ is state-dependent, rising with both the depth of the stress and the available buffer. *Right*: the ergodic output distribution under the optimal policy nearly coincides with that under the simple backing-bounded rule—the rule is near-optimal.

solve with a persistent shock state to confirm directly — the natural next quantitative step.

Idiosyncratic versus systematic persistence. The 11–13% figure prices the shock process of a *single issuer's* constant-maturity series ($\rho_\xi = 0.77$ monthly, half-life 2.7 months). An economy-wide stabilisation fund, however, absorbs the *systematic* factor: the cross-sectional median of $\hat{\kappa}$ across the full traded universe. Constructing that index from the complete Cbonds history panel (273 bonds aggregated to 37 issuer-month series, 127 months, 2015–2026) and estimating its dynamics directly gives $\rho = 0.921$ (half-life 8.4 months) with innovation $\sigma = 0.66\%$ — the median smooths idiosyncratic noise, and what remains is the slow-moving common credit factor. By the fitted law of equation (8), $S^* \propto \sigma^{2.3}(1 - \rho)^{-2.2}$, the systematic calibration scales the requirement by a factor of ≈ 8 : **a fund that must absorb the systematic factor needs on the order of 80–100% of annual support capacity**, not 11–13%. The distinction sharpens rather than overturns the paper's central result — it is the persistence-divergence of S^* evaluated at the persistence that actually obtains for the aggregate — and it implies a portfolio reading: the 11–13% benchmark applies to funds whose exposures diversify across issuers' idiosyncratic shocks, while the systematic component must be either pre-funded at near-annual scale or hedged outside the credit domain altogether (the cross-domain construction taken up in the companion safe-asset analysis).

8 Calibration and empirical discipline

The model has three kinds of numbers, and it is worth separating them. The *structural* parameters (Table 1) are standard preference and technology constants or Bernanke–Gertler–Gilchrist conventions. The *empirical* inputs (Table 2) are estimated from traded sukuk on Cbonds and are what every headline result rests on. The *reduced-form* parameters of the occasionally-binding block of Section 7 ($\psi_\kappa, \phi_o, b, \rho_y, \rho_S$) are stylised: they are not separately estimated but chosen so that normal-times dynamics are stable and the mechanism is transparent, and the conclusions drawn from that block are differences (state-dependence, the persistence comparison) rather than levels.

The empirical discipline is the distinctive feature: every quantity that moves a headline result is read from live sukuk markets rather than assumed. The convergence intensity, the

Structural parameter	Value	Source
β (discount)	0.97	$\sim 3\%$ real, program-consistent
σ (CRRA)	2.0	standard
α (capital share)	0.33	standard
δ_K (depreciation)	0.08	standard
G/Y	0.20	standard
χ (premium elasticity)	0.05	Bernanke et al. [1999]
leverage $\bar{K}/\bar{N}$	2.0	Bernanke et al. [1999]
θ_n (survival)	0.95	Bernanke et al. [1999]

Table 1: Structural calibration—standard or BGG conventions.

recovery, the credit-shock process and the asset-class scale all come from Cbonds, and they were re-pulled and re-estimated for this paper rather than inherited.

Empirical input	Value (Cbonds)	Enters
Asset-class scale	USD 761.3 bn outstanding, 4,701 issues, USD 277.3 bn international (Cbonds market statistics, 31 May 2026)	establishes the base is of monetary scale
Recovery δ	0.60 (senior sovereign)	spread $\rightarrow \kappa$ map
Baseline $\bar{\kappa}$	2.8% (Indonesia, 7y)	structural steady state
Sovereign $\bar{\kappa}$ cross-section	0.7–9.1% across 12 issuers (AA–CCC), 78 bonds (60 quoted)	determinacy ceilings, Fig. 2
$\bar{\kappa}$ constant-maturity dynamics	half-life 1.4–4.3 mo (7-year tenor, four issuers), Feller satisfied in 12/12 series	CIR- κ validity
Sovereign credit-shock AR(1)	$\rho_\xi = 0.77$, $\sigma_\xi = 0.73\%$ (half-life 2.7 mo)	buffer adequacy, Eq. (8)
Systematic κ -index (panel median, 127 mo)	$\rho = 0.921$ (half-life 8.4 mo), $\sigma = 0.66\%$	systematic-factor buffer ($\approx 8\times$)
Corporate credit-shock AR(1)	$\rho_\xi = 0.84$ median (0.71 mean), $\sigma_\xi = 0.74\%$ (26 issuers, 183 sukuk)	comparable to sovereign
Corporate $\bar{\kappa}$ premium	+31 bp median (213 corporate bonds, broader screener)	corporate ceiling, Fig. 3

Table 2: Empirical inputs, all estimated from traded sukuk on Cbonds. The determinacy region, the shock process, and the buffer-sizing rule are read from the data, not posited.

Three threads run through these inputs and reinforce one another. The *determinacy region* is populated by the real sovereign cross-section: the active-fiscal ceiling $e^{\bar{\kappa}} - 1$ is strictly positive and credit-ordered for every issuer, and widens for corporates by their credit premium. The *credit-shock process* that the Sovereign Stabilisation Fund must absorb is estimated directly from the κ time series, and its persistence—not just its size—is what determines the required backing. The *buffer-sizing rule* then closes the loop, mapping the estimated (ρ_ξ, σ_ξ) into the fund a central bank would actually need. Because all three are read from the same source, the empirical architecture is internally consistent: the sovereign sukuk that populate the determinacy map are the same instruments whose dynamics set the shock process and the backing requirement. The model is calibrated where convention is appropriate and estimated wherever a number carries economic weight.

9 Conclusion

An interest-free monetary economy is internally coherent and, under an explicit fiscal closure, determinate. Three findings organise the paper. First, the removal of interest is a statement about the *contract space*, not about preferences: time preference β remains, the intertemporal margin is carried by state-contingent equity returns, and the nominal interest rate survives only as an untraded shadow wedge that money demand reads as its opportunity cost. This is what allows a fully specified monetary general equilibrium to exist with no policy rate anywhere in it.

Second, price-level determinacy holds and is closed-form. Because the backing is real equity and policy is a quantity operation, the economy escapes the real-bills indeterminacy that an asset-backed money base otherwise courts, and the determinacy condition $\gamma_s < e^{\bar{\kappa}} - 1$ places the convergence intensity exactly where the interest rate sits in conventional fiscal theory. Crucially, given the equity-backing identity $M/P = \mathcal{V}$ —which makes the quantity rule and the fiscal valuation the same equation, leaving no separate active-money anchor—the fiscal-theory regime is not one closure chosen among alternatives but the only one available, and removing the interest rate *selects* the determinacy theory rather than leaving determinacy open. (It is the equity-backing of money, not the absence of an interest rate alone, that forces this; absent the identity a quantity-theoretic anchor would be the standard alternative.) The partial-equilibrium result of the companion paper is the exogenous-surplus corner of this region.

Third, the credit channel is determinate-friendly. A contemporaneous reduced-form feedback admits a determinacy threshold, but micro-founding the credit intensity as a Bernanke–Gertler–Gilchrist external-finance premium—first with a net-worth state, then with the full Tobin’s- Q revaluation channel—leaves the economy robustly determinate, with only modest amplification consistent with Kocherlakota [2000]. The price-level determinacy of the κ -Standard thus holds in partial equilibrium, in general equilibrium with endogenous surpluses, and across every financial-friction micro-foundation we examine. Optimal policy separates into a Friedman rule for liquidity and a maximal, backing-constrained credit-stabilisation rule, the latter understood as a welfare instrument rather than as determinacy insurance.

Finally, the linear robustness is a normal-times statement. The one regime outside it—a large credit shock against an exhausted Sovereign Stabilisation Fund—is characterised by an occasionally-binding model (Section 7) that produces state-dependent fire-sale amplification ($3.6\times$), a left-skewed ergodic distribution, and a quantitative backing-adequacy rule. Under transitory shocks a buffer of 3–5% of annual κ -support capacity suffices; but once the credit-shock process is estimated from the real sukuk data—which is persistent, with an AR(1) half-life near three months—the required buffer rises several-fold, to 11–13%, because credit stress clusters and drains the fund over successive months — and to near-annual scale (80–100%) if the fund must absorb the *systematic* factor, whose measured persistence ($\rho = 0.921$ on the full-universe index) far exceeds any single issuer’s. This answers the companion paper’s principal open empirical question, and answers it more cautiously than a transitory-shock calibration would: the interest-free economy is determinate and well-behaved in normal times, the one regime in which it is not is exactly the one the backing fund exists to mitigate, and the fund must be sized for persistent rather than one-off credit stress.

Beyond determinacy and stabilisation, the same backing identity carries a distributional corollary (Section 5.1). Because $M/P = \mathcal{V}$ pins money’s value to the real backing, money issued against output transfers no purchasing power to its first recipients: the Cantillon redistribution that an untethered active-money regime routes to asset-holders is zero in the κ -Standard, up to κ -mispricing, and—absent a positive-inflation mandate—the efficiency dividend of productivity growth accrues to every money holder rather than to owners of

appreciating assets. Distributional neutrality is therefore not an egalitarian assumption layered onto the model but a consequence of the very equation that forces the fiscal closure: the equation that pins the price level is the equation that zeroes the Cantillon transfer.

The model turns the companion paper’s monetary architecture into a general-equilibrium foundation. The remaining step is quantitative rather than conceptual: a full global solution of the occasionally-binding model [Guerrieri & Iacoviello, 2015], which would sharpen the backing-adequacy numbers without changing their direction. That is the natural next step for a quantitative interest-free business-cycle model.

References

- Ahmed, S., & Bhuyan, R. (2026a). The κ -Standard: debt-free money creation under a riba-free convergence intensity—uniting the Chicago Plan and Islamic full-reserve monetary theory. Working paper (companion paper).
- Ahmed, S., & Bhuyan, R. (2026b). The κ -yield curve: construction, estimation methodology, and real-data evidence from sovereign sukuk. *SSRN Working Paper* 6322938.
- Al-Jarhi, M. A. (1981). A monetary and financial structure for an interest-free economy. In *Money and Banking in Islam*. ICRIE.
- Askari, H., Iqbal, Z., Krichene, N., & Mirakhor, A. (2012). *Risk Sharing in Finance: The Islamic Finance Alternative*. Wiley.
- Benes, J., & Kumhof, M. (2012). The Chicago Plan revisited. *IMF Working Paper* 12/202.
- Benhabib, J., Schmitt-Grohé, S., & Uribe, M. (2001). The perils of Taylor rules. *Journal of Economic Theory*, 96(1–2), 40–69.
- Benhabib, J., Schmitt-Grohé, S., & Uribe, M. (2002). Avoiding liquidity traps. *Journal of Political Economy*, 110(3), 535–563.
- Bernanke, B., Gertler, M., & Gilchrist, S. (1999). The financial accelerator in a quantitative business cycle framework. *Handbook of Macroeconomics*, 1C, 1341–1393.
- Bianchi, J. (2011). Overborrowing and systemic externalities in the business cycle. *American Economic Review*, 101(7), 3400–3426.
- Borio, C. (2014). The financial cycle and macroeconomics: what have we learnt? *Journal of Banking & Finance*, 45, 182–198.
- Brunnermeier, M. K., & Sannikov, Y. (2014). A macroeconomic model with a financial sector. *American Economic Review*, 104(2), 379–421.
- Carlstrom, C. T., & Fuerst, T. S. (1997). Agency costs, net worth, and business fluctuations. *American Economic Review*, 87(5), 893–910.
- Chari, V. V., Christiano, L. J., & Kehoe, P. J. (1996). Optimality of the Friedman rule in economies with distorting taxes. *Journal of Monetary Economics*, 37(2), 203–223.
- Chapra, M. U. (1985). *Towards a Just Monetary System*. The Islamic Foundation.
- Cochrane, J. H. (2023). *The Fiscal Theory of the Price Level*. Princeton University Press.
- El Hamiani Khatat, M. (2016). Monetary policy in the presence of Islamic banking. *IMF Working Paper* WP/16/72.

- Fisher, I. (1936). 100% Money and the public debt. *Economic Forum*, 406–420.
- Friedman, M. (1969). The optimum quantity of money. In *The Optimum Quantity of Money and Other Essays*. Aldine.
- Gertler, M., & Karadi, P. (2011). A model of unconventional monetary policy. *Journal of Monetary Economics*, 58(1), 17–34.
- Guerrieri, L., & Iacoviello, M. (2015). OccBin: a toolkit for solving dynamic models with occasionally binding constraints easily. *Journal of Monetary Economics*, 70, 22–38.
- Jeanne, O., & Korinek, A. (2019). Managing credit booms and busts: a Pigouvian taxation approach. *Journal of Monetary Economics*, 107, 2–17.
- Khan, M. S., & Mirakhor, A. (1989). The financial system and monetary policy in an Islamic economy. *JKAU: Islamic Economics*, 1, 39–57.
- Kiyotaki, N., & Moore, J. (1997). Credit cycles. *Journal of Political Economy*, 105(2), 211–248.
- Kocherlakota, N. R. (2000). Creating business cycles through credit constraints. *Federal Reserve Bank of Minneapolis Quarterly Review*, 24(3), 2–10.
- Leeper, E. M. (1991). Equilibria under ‘active’ and ‘passive’ monetary and fiscal policies. *Journal of Monetary Economics*, 27(1), 129–147.
- Leeper, E. M., & Leith, C. (2016). Understanding inflation as a joint monetary–fiscal phenomenon. In *Handbook of Macroeconomics*, Vol. 2B, 2305–2415. Elsevier.
- Lucas, R. E. (2000). Inflation and welfare. *Econometrica*, 68(2), 247–274.
- Mendoza, E. G. (2010). Sudden stops, financial crises, and leverage. *American Economic Review*, 100(5), 1941–1966.
- Sargent, T. J., & Wallace, N. (1982). The real-bills doctrine versus the quantity theory. *Journal of Political Economy*, 90(6), 1212–1236.
- Sidrauski, M. (1967). Rational choice and patterns of growth in a monetary economy. *American Economic Review*, 57(2), 534–544.
- Sims, C. A. (1994). A simple model for study of the determination of the price level. *Economic Theory*, 4(3), 381–399.
- Woodford, M. (2001). Fiscal requirements for price stability. *Journal of Money, Credit and Banking*, 33(3), 669–728.
- Woodford, M. (2003). *Interest and Prices: Foundations of a Theory of Monetary Policy*. Princeton University Press.

EDITORIAL NOTE TO THIS PRINTING

Paper 10 · Money as a Mirror

Working paper, revised July 2026 SSRN 6906422

Role in the programme. The intellectual lineage: the argument did not begin with a blockchain. **What it establishes.** Continuity from al-Ghazali and Ibn Taymiyyah on money and the charge for time, through to the Western full-reserve tradition converging on debt-free money from the other direction. **Corrected or extended later.** No substantive corrections to date; it is the corpus's historical anchor. **Not established here.** Jurisprudential rulings; history motivates the question, it does not answer it.

Replication for every number in this paper: DOI 10.5281/zenodo.21290557 · github.com/shehzadahmed-xx/kappa-replication

Money as a Mirror: Classical Islamic Monetary Ontology and the Foundations of Interest-Free Pricing*

Shehzad Ahmed[†]

Rafiqul Bhuyan[‡]

June 2026

Abstract

Contemporary Shariah-compliant financial engineering defends the exclusion of interest almost entirely through twentieth- and twenty-first-century institutional economists, framing it as a regulatory adjustment to fractional-reserve banking. This paper recovers a deeper and older foundation. Drawing on al-Ghazali’s teleology of money as a “mirror,” Ibn Taymiyyah’s conception of money as an immutable measure (*mi’yar al-amwal*), Ibn Khaldun’s labour theory of value, and al-Maqrizi’s account of currency competition, it argues that the prohibition of *riba* follows by construction from a classical ontology of *what money is*, and that a *riba*-free pricing primitive—the convergence intensity κ , which a purpose-built blockchain substrate is designed to enforce at the consensus level—is the mechanical instantiation of that ontology rather than a modern compromise with it. The contribution is conceptual: it supplies the κ -program with its missing foundation and grounds, in named classical authority, claims its companion papers currently assert without one. Throughout, the substrate is said to *operationalize* classical principles; we do not claim the classical authors endorsed distributed consensus.

Keywords: al-Ghazali; money as a mirror; *riba*; Islamic monetary theory; interest-free pricing; convergence intensity (κ); Ibn Taymiyyah; Ibn Khaldun; history of economic thought; blockchain; Shariah.

JEL classification: B11; E42; Z12; P48; O33.

1 Introduction

Modern Islamic finance carries a quiet embarrassment at its foundation. Its instruments—the *murabaha* mark-up, the *ijara* lease, the *sukuk* certificate—are routinely defended as compliant *workarounds*: structures that reproduce the cash-flows of conventional finance while avoiding the contractual form of interest. The literature that justifies them is overwhelmingly contemporary, and overwhelmingly institutional: El-Gamal (2006), Chapra (1985), Al-Jarhi (1981), Usmani (2002), and the standards of AAOIFI. That literature is indispensable for modern practice, but it tends to present the exclusion of interest as a *regulatory adjustment*—a constraint imposed on an otherwise conventional monetary system. On this telling, *riba*-freedom is something Islamic finance must engineer around, at cost.

*This paper provides the classical conceptual foundation of the κ -rate research program; companion papers develop the formal no-arbitrage pricing, empirical, and general-equilibrium results referenced herein.

[†]Department of Finance, Independent University Bangladesh (IUB), Dhaka. 2120922@iub.edu.bd.

[‡]Corresponding author. Department of Accounting and Finance, Alabama A&M University, Normal, AL, USA. rafiquel.bhuyan@aamu.edu.

This paper argues that the framing is backwards. The classical Islamic monetary canon—the work of jurists and historians writing between the eleventh and fifteenth centuries—does not treat the prohibition of interest as a constraint bolted onto a theory of money. It derives the prohibition *from* a theory of money: an account of what money is for, from which the impermissibility of earning money from money follows as a matter of definition. Recover that ontology, and interest-freedom ceases to be a compliance cost and becomes the natural state of a correctly-specified monetary system.

The argument matters now because, for the first time, that classical ontology can be *enforced* rather than merely asserted. The companion papers of this program construct a riba-free pricing primitive—the convergence intensity κ , a stopping-time hazard that plays the role the risk-free rate r plays in conventional finance, but with no interest input—and embed it in a blockchain substrate¹ whose consensus rules are designed to make a positive interest parameter structurally impossible. The present paper supplies what that construction has lacked: its conceptual foundation. It shows that each core invariant of the substrate is the algorithmic form of a specific classical principle.

A discipline governs the whole. We claim that the substrate *operationalizes* or *instantiates* classical principles; we never claim that al-Ghazali or Ibn Taymiyyah *endorsed* Byzantine-fault-tolerant consensus, which would be a plain anachronism. The classical authors supply the normative ontology of money; the protocol supplies a mechanism that realizes it. The verb is “operationalizes,” never “validates.”

2 The modern compromise and its silent gap

The intellectual scaffolding of contemporary Islamic monetary economics is a twentieth-century construction. Al-Jarhi (1981) and Chapra (1985) set out the case for a non-interest monetary system; Khan & Mirakhor (1987), Iqbal & Mirakhor (2011), and Askari et al. (2012) developed its banking and macro-prudential implications; El-Gamal (2006) supplied the now-standard operational analysis of riba and *gharar*. This body of work is rigorous and necessary, and the present paper depends on it. But it shares a characteristic feature: it argues at the level of *institutions and contracts*—reserve regimes, profit-and-loss-sharing ratios, the legal form of permissible instruments—and rarely descends to the ontological question of what money *is*. Where it reaches for first principles, it reaches for the Qur’anic prohibition (Q 2:275) and moves directly to modern implementation, leaving a gap of some thirteen centuries between scripture and the 1980s—an instance, within Islamic finance itself, of the lacuna that Ghazanfar & Islahi (1990) and Ghazanfar (2003) term the Schumpeterian “Great Gap”: the neglected centuries of medieval Islamic economic thought between the Greeks and the Latin Scholastics (Schumpeter, 1954).

That gap is not empty in the historical record; it is empty only in the program literature and in the working canon of Islamic finance. A smaller scholarly tradition has excavated the classical monetary writings—Islahi (2001) on al-Ghazali and Islahi (2012) on Ibn al-Qayyim, with Islahi (2014) surveying the tradition to the fifteenth century, Ghassan (2019) on the equation of exchange in Ibn Taymiyyah and Ibn Khaldun, and Figuera (2018) on medieval Islamic monetary thought as a whole, with Uddin (2019) surveying Islamic monetary economics. The present

¹For the reader coming from the history of economic thought: the substrate is the design of a proof-of-stake blockchain (the κ -Chain) whose consensus rules are specified to (i) reject any transaction implying a positive interest parameter ι and slash a validator who includes one; (ii) compute κ each block as a stake-weighted median of validators’ market observations; and (iii) pay validators only from real transaction fees, with no inflationary issuance. These three invariants— $\iota = 0$ enforcement, the κ oracle, and the fee-only reward—plus the execution-isolation choice discussed with al-Maqrīzī below, are what the classical principles are mapped onto in the Synthesis. At the time of writing the κ -Chain consensus layer is a specification; the EVM reference implementation of the instruments (Baraka Protocol) runs on a public testnet, while the consensus substrate described here is a design.

paper continues their recovery and extends it in a direction those recoveries did not pursue: it connects the classical ontology to a *working mechanism*. The modern compromise asks how to avoid interest within a monetary system designed around it. The classical canon asks what money is such that interest was never its proper use.

A second body of work bounds the contribution from the other side, and the space between them must be stated precisely. A growing literature already reads the classical scholars *forward* onto digital money, asking whether cryptocurrencies satisfy al-Ghazali’s and Ibn Taymiyyah’s criteria for a valid medium of exchange—and reaching mixed verdicts, from qualified acceptance (Alzubaidi & Abdullah, 2017; Haq & Ali, 2018) to scepticism on grounds of volatility and speculation (Siswanto et al., 2020). That literature renders permissibility verdicts on instruments that already exist—is a given coin money in the Shariah sense?—and takes the artefact as given. The present paper inverts the direction: it treats the classical ontology as a *design specification* and exhibits a substrate *design* engineered to satisfy it by construction, in which the prohibition of interest is enforced as a consensus invariant rather than adjudicated after the fact, and in which the *riba*-free measure κ is given a formal, no-arbitrage definition rather than asserted. The novelty is therefore neither the classical recovery (Islahi, 2001; Figuera, 2018) nor the act of relating the canon to digital money (Alzubaidi & Abdullah, 2017; Siswanto et al., 2020), but the demonstration that a specific, fully-specified mechanism instantiates the ontology as an enforced and priced invariant.

The neglect this paper addresses is not incidental; it has a name in the history of economic thought. Schumpeter’s *History of Economic Analysis* (Schumpeter, 1954) posited a “Great Gap”—“over 500 years” that, on his account, could safely be leapt over on the way from the Greeks to St Thomas Aquinas—in which nothing of analytical significance was written; the leapt-over centuries are precisely those of medieval Islamic scholarship. Ghazanfar & Islahi (1990) and the contributors to Ghazanfar (2003) have shown the claim untenable: the gap is filled, densely, by the very authors on which this paper draws. Yet the recovery has remained largely internal to the history-of-thought literature; it has not been carried into the working canon of Islamic finance, which continues to reach past the medieval monetary writers to the Qur’anic text and then forward to the twentieth century. The wager here is that this canon is not merely of antiquarian interest but *operationally* consequential: it specifies, with a precision the modern compromise lacks, what a monetary system must do to be interest-free by construction rather than by exception. The *maqasid* (higher-objectives) tradition frames the prohibition of *riba* not as an isolated rule but as protection of a wider good—the integrity of exchange and the circulation of wealth—which is the register in which the classical monetary authors wrote.

3 Al-Ghazali (1058–1111): the teleology of money and the mirror

Abu Hamid al-Ghazali’s account of money, set out in the *Ihya’ Ulum al-Din*, is not a list of money’s functions but a teleology—an argument about the *purpose* for which money exists, from which its permitted and forbidden uses follow. Money, for al-Ghazali, is a remedy for the structural failures of barter: the absence of a common denominator and the impossibility of relating dissimilar goods. Its defining property is that it has no value of its own. “Dirhams and dinars,” he writes, “are not created for any particular persons, they are useless by themselves; they are just like stones. They are created to circulate from hand to hand, to govern and to facilitate transactions” (al-Ghazali, *Ihya’*; trans. Islahi, 2001).

From this absence of intrinsic value al-Ghazali draws his central image. A medium can measure all things precisely *because* it is itself featureless: “a thing [such as money] can be exactly linked to other things if it has no particular form or feature of its own, for example, a mirror which has no colour but can reflect all colours. Same is the case with money—it has no

purpose of its own but it serves as a medium for the purpose of exchanging goods” (al-Ghazali, *Ihya'* (al-Ghazali, c.1105); trans. Islahi, 2001). Money is a mirror: it carries no colour, and exists only to reflect the colours—the values—of real goods.

The prohibitions follow by construction. To hoard money (*kanz*; cf. Q 9:34)² is to withdraw the mirror from circulation, defeating the purpose for which it was created. And to trade money for money at a premium—*riba*—is to make the mirror an end in itself. In al-Ghazali’s own framing, when “a person is permitted to sell [or exchange] money with money, then such transactions will become his goal” (al-Ghazali, *Ihya'*; trans. Islahi, 2001)—the medium is forced to generate its own colour, a return arising from the mirror rather than from the goods it should reflect. Interest is therefore not a contingent abuse to be regulated but a *category error*: a demand that a colourless thing produce colour.

What the κ -substrate operationalizes. The protocol’s substrate invariant—the interest parameter ι in the funding-engine state is asserted to equal zero, with the consensus layer rejecting at inclusion any state transition that implies $\iota > 0$ and slashing a proposer who admits one—is a mechanical instantiation of the mirror principle. The native unit is structurally prevented from accruing a time-value return; it can only reflect realized commercial risk, which the program measures as the convergence intensity κ . In an extension of al-Ghazali’s image, κ is the colour of the goods—an intensity read from real trade—and the substrate guarantees the mirror never manufactures colour of its own.

Two further elements of al-Ghazali’s account deserve emphasis, because each has a direct structural analogue. The first is his critique of the *store-of-value* function. Al-Ghazali does not deny that money can be held, but he warns that when the store-of-value motive predominates—when money is desired for its own accumulation rather than for the exchange it enables—the medium is diverted from its purpose; hoarding (*kanz*), he writes, is like imprisoning a just ruler, depriving the community of the circulation through which value is realized. The second is his treatment of debasement and counterfeiting, which he condemns not merely as fraud but as a corruption of the measure itself: to adulterate the coinage is to falsify the yardstick against which all goods are priced, an injustice to every party to every subsequent exchange. Both concerns—that money not become an end, and that the measure not be corrupted—reappear, transposed into mechanism, in the substrate’s refusal of inflationary issuance (debasement by another name) and its protection of the published measure from manipulation. The two ways al-Ghazali identifies by which a monetary system betrays its purpose—letting the medium become the goal, and corrupting the measure—are precisely the two failures the architecture is built to foreclose.

4 Ibn Taymiyyah (1263–1328): money as an immutable metric

Where al-Ghazali supplies the ontology, the Hanbali jurist Ibn Taymiyyah supplies the operational mandate. Money, he argues, is fundamentally a *measure*: “*Athman* [prices, that which is paid as price] are meant to be a measurement of objects of value (*mi'yar al-amwal*), through which the quantities of the objects of value (*maqadir al-amwal*) are known; and they are never meant to be used for themselves” (Ibn Taymiyyah, *Majmu' al-Fatawa* 29; trans. Islahi, 2001). Money is the yardstick of every contract, and like al-Ghazali’s mirror it is corrupted the moment it becomes the object rather than the instrument of exchange.

From the metric conception Ibn Taymiyyah derives two claims that bear directly on protocol design. The first is normative: because money measures the fairness of all obligations, the authority’s tampering with it—debasement by alloy or by reducing weight—is a systemic

²We use *kanz*, the Qur’anic register for the hoarding of money, rather than *ihtikar*, which in fiqh usage denotes the monopolistic hoarding of commodities.

injustice (*zulm*), since it warps the standard against which every party’s labour and goods are valued. The second is analytical: the quantity of money in circulation should be proportioned to the volume of transactions, so that the metric neither inflates nor deflates the values it records—an early articulation of a quantity-theoretic discipline (Ghassan, 2019).

An honest caveat against overreading. Ibn Taymiyyah is sometimes conscripted for a strict commodity-money (gold-and-silver-only) position. The texts do not support this. He explicitly accepted *fulus*—copper token money—provided it is minted “according to the just value of the people’s transactions, without injustice to them” (Ibn Taymiyyah, c.1315). His objection is to *debasement and distortion of the metric*, not to non-commodity money as such, and the present account does not enlist him against token or fiat money per se.

What the κ -substrate operationalizes. Two design elements descend from Ibn Taymiyyah, with different strengths. First, the oracle module operationalizes the demand that the value-metric be insulated from manipulation: index observations enter consensus through ABCI++ vote extensions and are aggregated by a stake-weighted median, so that no single proposer—no “authority”—can debase the published κ . Second, his quantity-tracks-transactions principle maps most directly onto the fee-only, non-inflationary reward of Section 5: the medium expands only with realized activity, never by fiat issuance.

It is worth dwelling on the originality of Ibn Taymiyyah’s second claim, because it is easily underestimated. To hold that the quantity of money should be proportioned to the volume of transactions is to grasp, in the fourteenth century, the intuition that later crystallized as the quantity theory of money: that the value of the monetary unit is not intrinsic but depends on the relation between the stock of money and the goods it must measure. Ghassan (2019) reconstructs this as a “Shariah rationale money” in which the velocity and volume of transactions are endogenous and the stock is the variable the authority must discipline. The normative and the analytical claims are, for Ibn Taymiyyah, one: the authority’s duty to keep the measure stable just *is* the requirement that issuance track real activity rather than fiscal convenience. This is the seam along which the substrate’s fee-only, zero-emission reward is cut—issuance, in effect, is bound to realized transactions because there is no issuance at all beyond the fees those transactions generate.

5 Ibn Khaldun (1332–1406): the labour theory of value and velocity

Ibn Khaldun’s *Muqaddimah* completes the picture by tying the value money measures to its ultimate source. Value, for Ibn Khaldun, originates in human effort: “profit is the value realized from human labour”, and “when there is more labour, the value realized from it increases among the people” (Ibn Khaldun, 1377/2015, 273ff.). Money itself is not wealth but a vehicle through which wealth, produced by labour, is acquired and measured. Prosperity, in his account, rises with the *velocity* of money through real commercial activity, and falls when money is sequestered in passive accumulation rather than circulated through production and trade.

The implication is a sharp criterion of legitimacy for return. A gain that tracks productive labour, or genuine exposure to commercial risk, is value realized; a gain that accrues to capital without corresponding effort or risk is not value created but value redistributed—a dilution of the real measure. This is the conceptual ancestor of the program’s distinction between κ (an intensity earned from real market activity) and r (a return to the mere passage of time).

What the κ -substrate operationalizes. The fee-only wakala reward of the distribution module instantiates the Khaldunian criterion at the level of network economics. Block producers

receive no inflationary emission whatsoever; their entire compensation is drawn from the transaction fees generated by real trades clearing on the order book. Validators are therefore paid for the actual labour of producing blocks—computation, latency, security, storage—and for nothing else, with no value minted for idle capital. Network income rises and falls with transaction velocity, precisely as Ibn Khaldun holds that legitimate prosperity must.

Ibn Khaldun’s monetary claims sit inside a larger theory of *umran* (civilisation) and its rise and decline, and the embedding matters. For Ibn Khaldun the wealth of a society is a function of the division of labour and the density of cooperative production—*asabiyya* furnishing the social cohesion that makes it possible—and money is the instrument through which that produced value is measured and exchanged, its circulation tracking the vitality of the underlying production. Decline sets in, on his account, when the state’s extraction (through confiscatory taxation or monetary manipulation) outpaces the productive labour that sustains it, so that the measure is loaded with charges unrelated to real value. The lesson the architecture draws is narrow but exact: a network’s compensation should be a function of the real economic activity it processes, not of a claim levied on the mere passage of time or the mere holding of the token. The fee-only reward is the Khaldunian criterion enforced at the level of protocol incentives—validators prosper precisely to the extent that real exchange occurs, and not otherwise.

6 Al-Maqrizi (1364–1442): currency competition and systemic integrity

The last of the four, Ibn Khaldun’s student Ahmad al-Maqrizi, contributes a systemic warning rather than a definition. In *Ighathat al-Umma bi-Kashf al-Ghumma*, his analysis of the Mamluk monetary crisis, al-Maqrizi documents how the over-issuance of low-quality copper *fulus* alongside the gold and silver standard drove the sound coin out of circulation and produced ruinous inflation—an early statement of the principle later named Gresham’s Law, that “bad money drives out good” (al-Maqrizi, *Ighatha*; trans. al-Maqrizi, c.1405; cf. Figuera, 2018). His prescription was the restoration of a sound standard, with debased token coin confined to petty transactions.

What the κ -substrate operationalizes—and an explicit limit. The protocol’s choice to keep execution within a sovereign, isolated virtual machine rather than bridging natively to an external interest-bearing ($r > 0$) settlement layer can be *framed* in al-Maqrizi’s terms: an interest-bearing wrapped asset admitted into a zero-interest substrate is the digital analogue of debased coin admitted to a sound system. But this is an *analogy, not a derivation*, and it is the weakest of the four mappings for two reasons. First, al-Maqrizi’s subject is the debasement of *coinage*, not the interoperability of execution environments. Second, Ibn Taymiyyah’s acceptance of *fulus* shows the classical tradition is not monolithically hostile to “lesser” money. The execution-isolation decision therefore stands on its independent technical grounds—determinism, light-client compatibility, the avoidance of a consensus fork, and the refusal to depend on infrastructure the protocol does not control—with al-Maqrizi supplying historical resonance rather than load-bearing proof.

The historical specificity of al-Maqrizi’s analysis is part of its force. Writing as an eyewitness-historian of the Mamluk crisis of 806–808/1403–1405, he traces a causal chain—fiscal mismanagement, the over-minting of copper *fulus*, the flight of gold and silver from circulation, and the collapse of the price system—with a rigour that anticipates not only Gresham but the broader recognition that a debased medium does not merely lose value but *disorders* the measurement of all other values. The point for the present argument is deliberately modest: al-Maqrizi establishes that the integrity of a monetary system can be destroyed from within by the admission of an inferior, over-issued instrument alongside a sound one; whether that

observation transfers to the admission of an interest-bearing wrapped asset into an interest-free execution environment is a question of analogy, not derivation, and the analogy is offered as illustration rather than proof.

7 Supporting voices: a tradition, not four figures

The four principal figures are the prominent peaks of a continuous tradition, and three further voices confirm that the ontology is a school rather than a selection. **Ibn Rushd** (Averroes, 1126–1198), in the *Bidayat al-Mujtahid*, gives the prohibition of riba its most systematic juristic analysis, distinguishing riba al-fadl (excess in the exchange of fungibles) from riba al-nasi’ah (excess through deferment) and—characteristically—not merely stating the law but reconstructing the *rationale* the schools offer for it, locating the wisdom of the prohibition in the preservation of equivalence and the prevention of exploitation in exchange (Ibn Rushd, c.1180). His move from rule to underlying reason is the juristic counterpart of al-Ghazali’s move from function to teleology.

Ibn al-Qayyim (1292–1350), in *I’lam al-Muwaqqi’in*, supplies the structural logic that connects the two categories of riba. He classifies riba al-nasi’ah as the manifest usury (*jali*) and riba al-fadl as the hidden (*khafi*), the latter forbidden not in itself but as *sadd al-dhari’ah*—the blocking of a means to the former, “prohibited to stop it from becoming” open usury (Islahi, 2012; Ibn al-Qayyim, c.1345). This anticipatory logic is precisely the reasoning the companion perpetual-futures paper (Ahmed & Bhuyan, 2026d) invokes in its treatment of *bay’ al-kali bil kali* (debt-for-debt), and it grounds the substrate’s posture of foreclosing interest at the level of structure rather than policing it after the fact.

Al-Dimashqi (Abu al-Fadl Ja’far ibn Ali al-Dimashqi, eleventh–twelfth century), in his merchant manual *Kitab al-Ishara ila Mahasin al-Tijara*, descends from jurisprudence to commercial practice, treating money’s role in valuation and prudent exchange and the ethics of honest dealing (al-Dimashqi, 11th–12th c.). His exact dates are debated—the sole internal evidence is a coin he names—but his authorship is established, and his presence shows the ontology of money as neutral measure was not the preserve of theologians but the working assumption of the merchant class itself.

Taken together, these three voices close the distance between the high monetary theory of al-Ghazali and Ibn Taymiyyah and the daily practice of the market. Ibn Rushd supplies the jurist’s reasoned account of *why* the prohibition holds, refusing to leave it as bare command; Ibn al-Qayyim supplies the structural logic by which a lesser, hidden form of riba is foreclosed precisely to protect against the greater, manifest form—the *sadd al-dhari’ah* reasoning the architecture mirrors in foreclosing interest *ex ante* rather than policing it *ex post*; and al-Dimashqi shows that the conception of money as a neutral measure, far from being a theologian’s abstraction, was the operating assumption of the merchant himself. The cumulative effect is to establish that the ontology this paper recovers is not a doctrine assembled from convenient fragments but a tradition—continuous across four centuries and across the schools—Shafi’i, Maliki, and Hanbali alike—and the divide between the scholar’s study and the trader’s ledger.

8 Synthesis: a classical ontology of interest-free pricing

Read together, the canon yields four load-bearing principles. From al-Ghazali and Ibn Taymiyyah: money is a **neutral measure**—a mirror, a *mi’yar*—with no value of its own, corrupted the moment it becomes the object rather than the instrument of exchange. From Ibn Khaldun: legitimate value originates in **labour and real activity**, so that return must track effort or genuine risk, never the mere passage of time over idle capital. From Ibn Taymiyyah again and from al-Maqrizi: the **integrity of the measure** must be defended. And from Ibn

Rushd and Ibn al-Qayyim: the prohibition is to be secured **structurally**, by reason and by the blocking of means.

These are one ontology seen from four angles, and the κ -substrate instantiates each in a distinct invariant:

Component	Rule	Classical anchor	Operationalizes
<code>x/perp</code>	$\iota = 0$ at inclusion; slashing	al-Ghazali (+ Ibn Taymiyyah)	neutral mirror; no time-value r
<code>x/oracle</code>	ABCI++ median κ	Ibn Taymiyyah	integrity of the value-metric
<code>x/distribution</code>	0% inflation; fee-only	Ibn Khaldun (+ quantity discipline)	return tracks labour/velocity
<code>x/evm</code>	sovereign isolated execution	al-Maqrizi (<i>illustrative</i>)	resists “bad-money” pollution
posture	foreclose $\iota > 0$ ex ante	Ibn Rushd; Ibn al-Qayyim	prohibition by reason / <i>sadd a</i>

Table 1: Classical principles and the substrate invariants that operationalize them. Module names refer to the κ -Chain design: `x/perp` (the funding engine), `x/oracle` (the κ oracle), `x/distribution` (the validator-reward module), `x/evm` (the isolated execution environment).

The mapping is bounded, and its strength varies by figure: tight for al-Ghazali’s mirror and the $\iota = 0$ invariant, strong for Ibn Taymiyyah and Ibn Khaldun, explicitly illustrative—the weakest—for al-Maqrizi, and structural for the Ibn Rushd/Ibn al-Qayyim posture.

9 From ontology to mechanism: the κ -substrate

The bridge from this ontology to a working system is the convergence intensity κ . [Ackerer et al. \(2025\)](#) establish that, when the interest parameter of a perpetual contract is set to zero, a unique no-arbitrage price still exists; the companion papers show that this interest-free price is isomorphic to the reduced-form pricing of credit, with κ playing the role of the default intensity ([Ahmed & Bhuyan, 2026a](#)). κ is therefore a measure read entirely from realized market activity, with no interest input—the formal counterpart of al-Ghazali’s “colour of the goods.”

At the level of pricing, the program’s Separation Theorem renders this mirror an identity rather than a figure of speech. A claim’s value embeds a charge for the passage of *time* (the interest parameter ι) and a charge for the intensity of real *convergence or risk* (the hazard κ), and the no-arbitrage price remains uniquely defined when the first is set to zero ([Ackerer et al., 2025](#); [Ahmed & Bhuyan, 2026a](#)):

$$\text{value} = f\left(\underbrace{\iota}_{\text{time-value}}, \underbrace{\kappa}_{\text{real risk}}\right), \quad \iota = 0 \implies \text{value} = f(0, \kappa).$$

Money may reflect the colour of the goods (κ) but may not manufacture colour of its own (ι); setting $\iota = 0$ does not break pricing, it isolates the reflection. And the mirror is not merely a metaphor: a companion empirical study measures κ directly on real data and finds it behaves as a probability intensity must—strictly non-negative across a cross-section of 195 actively-quoted USD sukuk, and Feller-consistent in every estimated issuer series of a separate constant-maturity dynamics panel ([Ahmed & Bhuyan, 2026b](#))—so the program’s κ behaves, in modern data, as the classical ontology requires of a measure of real risk—an instantiated property rather than an aspiration. The substrate is designed to do three things the classical authors could only prescribe. It *enforces* the neutrality of money by making $\iota > 0$ unrepresentable in valid state. It *protects the measure* by computing κ from a stake-weighted median of validator observations, so that no single authority can debase it. And it *ties reward to labour* by paying validators only from realized transaction fees. What al-Ghazali, Ibn Taymiyyah and Ibn Khaldun could establish only as principle, the substrate is specified to establish as invariant—to be checked every block, by every honest node.

10 Conclusion and honest boundaries

This paper establishes a genuine classical lineage for interest-free pricing, and in doing so closes the thirteen-century citation gap in the program literature — and shows how the working canon of Islamic finance, which still reaches past the medieval monetary writers, can close its own. The exclusion of interest is not, on this account, a modern regulatory patch; it is the recovery of an ontology in which money is a neutral measure whose proper use was never to breed itself.

The boundaries are as important as the claim. The paper does *not* establish that the classical authors endorsed blockchain, consensus, or any modern mechanism—the relation is operationalization, not endorsement. The al-Maqrizi mapping is explicitly the weakest, illustrative rather than probative. And the canon settles only the *riba* question; the contract-level questions of *gharar* (excessive uncertainty), *maysir* (speculation), and *qabd* (constructive possession) on cash-settled instruments are not resolved here and remain matters for a qualified Shariah board. What the classical ontology supplies is not a fatwa but a foundation: the demonstration that a correctly-specified monetary substrate is interest-free not by constraint but by construction.

It bears restating where the argument stops. The classical ontology recovered here speaks to *riba*, and the substrate enforces what that ontology requires. It does not, and cannot, settle the adjacent prohibitions—*gharar* (excessive uncertainty), *maysir* (speculation), and *qabd* (possession)—which turn on asset ontology, contractual intent, and delivery rather than on the time-value of money, and which therefore remain matters for qualified jurisprudential adjudication, particularly for cash-settled synthetic instruments. A companion note sets out that division explicitly and proposes the product-layer conditions under which asset-backed instruments built on the *riba*-free foundation can satisfy the remaining pillars (Ahmed & Bhuyan, 2026c). The claim of the present paper is correspondingly bounded: that the foundation is sound, classical, and enforceable—not that the foundation alone renders every instrument built upon it compliant.

References

- Ackerer, D., Hugonnier, J., & Jermann, U. (2025). Perpetual Futures Pricing. *Mathematical Finance*. DOI 10.1111/mafi.70018.
- Ahmed, S., & Bhuyan, R. (2026a). From Perpetual Contracts to Islamic Credit: The Random Stopping-Time Equivalence and a Riba-Free Foundation for Sukuk, Takaful, and Credit Protection. SSRN Working Paper 6322858.
- Ahmed, S., & Bhuyan, R. (2026b). The κ -Yield Curve: Construction, Estimation Methodology, and Real-Data Evidence from Sovereign Sukuk. SSRN Working Paper 6322938.
- Ahmed, S., & Bhuyan, R. (2026c). The Jurisprudential Boundaries of an Interest-Free Substrate: Riba, Gharar, Maysir, and Qabd in Cash-Settled Synthetic Instruments. Research note, κ -rate program.
- Ahmed, S., & Bhuyan, R. (2026d). The Interest Parameter in Perpetual Futures: Shariah Analysis and Empirical Evidence from Centralized and Decentralized Exchanges. SSRN Working Paper 6322778.
- Al-Jarhi, M. A. (1981). *A Monetary and Financial Structure for an Interest-Free Economy*. International Centre for Research in Islamic Economics, Jeddah.
- Alzubaidi, I. B., & Abdullah, A. (2017). Developing a Digital Currency from an Islamic Perspective: Case of Blockchain Technology. *International Business Research* 10(11), 79–87.
- Askari, H., Iqbal, Z., Krichene, N., & Mirakhor, A. (2012). *Risk Sharing in Finance: The Islamic Finance Alternative*. Wiley.
- Chapra, M. U. (1985). *Towards a Just Monetary System*. The Islamic Foundation, Leicester.

- al-Dimashqi, Abu al-Fadl Ja'far ibn Ali. *Kitab al-Ishara ila Mahasin al-Tijara*. Ed. H. Ritter, "Ein arabisches Handbuch der Handelswissenschaft," *Der Islam* 7 (1917), 1–91.
- El-Gamal, M. A. (2006). *Islamic Finance: Law, Economics, and Practice*. Cambridge University Press.
- Figuera, S. (2018). Considerations on Islamic Economic Thought Regarding Monetary Matters in the Middle Ages. *OEconomia* 8(1), 1–28.
- Ghassan, H. B. (2019). Re-examining the Equation of Exchange according to Shari'ah Rationale Money. Munich Personal RePEc Archive, No. 91666.
- al-Ghazali, Abu Hamid. *Ihya' Ulum al-Din*, Book of Shukr. (As translated in Islahi, 2001.)
- Ghazanfar, S. M., & Islahi, A. A. (1990). Economic Thought of an Arab Scholastic: Abu Hamid al-Ghazali. *History of Political Economy* 22(2), 381–403.
- Ghazanfar, S. M. (Ed.) (2003). *Medieval Islamic Economic Thought: Filling the Great Gap in European Economics*. RoutledgeCurzon, London.
- Haq, H. B., & Ali, S. T. (2018). Navigating the Cryptocurrency Landscape: An Islamic Perspective. arXiv:1811.05935.
- Ibn Khaldun, A. R. ([1377] 2015). *The Muqaddimah: An Introduction to History*. Trans. F. Rosenthal. Princeton University Press.
- Ibn al-Qayyim al-Jawziyya. *I'lam al-Muwaqqi'in 'an Rabb al-'Alamin*.
- Ibn Rushd (Averroes). *Bidayat al-Mujtahid wa Nihayat al-Muqtasid*. Trans. I. A. K. Nyazee (1994–96), *The Distinguished Jurist's Primer*. Garnet, Reading.
- Ibn Taymiyyah, A. *Majmu' al-Fatawa*, vol. 29.
- Iqbal, Z., & Mirakhor, A. (2011). *An Introduction to Islamic Finance: Theory and Practice*, 2nd ed. Wiley.
- Islahi, A. A. (2001). *An Analytical Study of al-Ghazali's Thought on Money and Interest*. Munich Personal RePEc Archive, No. 41438.
- Islahi, A. A. (2012). Economic Thought of Ibn al-Qayyim (1292–1350). Munich Personal RePEc Archive, No. 41369.
- Islahi, A. A. (2014). *Contributions of Muslim Scholars to the History of Economic Thought and Analysis up to the 15th Century*. Munich Personal RePEc Archive, No. 53462.
- Khan, M. S., & Mirakhor, A. (1987). *Theoretical Studies in Islamic Banking and Finance*. The Institute for Research and Islamic Studies, Houston.
- al-Maqrizi, T. A. *Ighathat al-Umma bi-Kashf al-Ghumma*. Trans. A. Allouche (1994), *Mamluk Economics: A Study and Translation of al-Maqrizi's Ighathah*. University of Utah Press.
- Schumpeter, J. A. (1954). *History of Economic Analysis*. Oxford University Press, New York.
- Siswantoro, D., Handika, R., & Mita, A. F. (2020). The Requirements of Cryptocurrency for Money, an Islamic View. *Heliyon* 6(1), e03235.
- Uddin, M. A. (2019). Islamic Monetary Economics: Insights from the Literature. In *Islamic Monetary Economics and Institutions*, Springer, 39–53.
- Usmani, M. T. (2002). *An Introduction to Islamic Finance*. Kluwer Law International.

EDITORIAL NOTE TO THIS PRINTING

Paper 11 · Jurisprudential Boundaries

Working paper, revised July 2026 · SSRN 6906458

Role in the programme. The fiqh frame: where the mathematics hands over. **What it establishes.** The four product-layer conditions for the asset-backed class, and the precise statement of what $\iota = 0$ does and does not settle. **Corrected or extended later.** The 2026 addendum absorbs Paper 16's risk-premium result and narrows the reserved question to its residuals (tokenised qabd, the wakala fee form). **Not established here.** Whether the conditions satisfy the pillars; that determination belongs to a board of scholars, and the paper says so.

Replication for every number in this paper: DOI 10.5281/zenodo.21290557 · github.com/shehzadahmed-xx/kappa-replication

The Jurisprudential Boundaries of an Interest-Free Substrate: Riba, Gharar, Maysir, and Qabd in Cash-Settled Synthetic Instruments

A Research Note*

Shehzad Ahmed[†]

Rafiqul Bhuyan[‡]

June 2026

Abstract

Efforts to engineer Shariah-compliant financial infrastructure frequently conflate the elimination of interest (*riba*) with the holistic resolution of all contractual defects (*fasad*). A decentralised ledger can structurally eliminate the time-value of money by coding monetary neutrality ($\iota = 0$) into consensus; but *fiqh al-muamalat* treats *riba*, *gharar* (excessive uncertainty), *maysir* (gambling), and *qabd* (possession) as distinct pillars requiring independent remedies. Synthesising classical and contemporary jurisprudence on cash-settled synthetic derivatives, this paper states precisely where algorithmic compliance ends and qualitative judgment must begin, and proposes four product-layer rules under which asset-backed instruments built on an interest-free substrate address the remaining pillars by construction, with the intent residual reserved to Shariah-board judgment. The claim is bounded: the substrate resolves *riba al-nasiah*; it does not, on its own, settle *gharar*, *maysir*, or *qabd* for cash-settled synthetics.

Keywords: *riba*; *gharar*; *maysir*; *qabd*; Islamic derivatives; cash-settled synthetics; Islamic CDS; *tabarru*; *takaful*; AAOIFI; Shariah compliance.

JEL classification: G13; K12; P48; Z12; G23.

1 Introduction

The promise of a Shariah-compliant decentralised ledger is easily overstated. A recurring pattern in the literature on Islamic cryptocurrency and digital finance is to establish that some interest-bearing feature has been removed and to infer that the resulting instrument is therefore permissible—an inference that collapses the several distinct prohibitions of *fiqh al-muamalat* into the single question of interest. This note resists that collapse. It takes as given that an interest-free consensus substrate can be built—that *riba* can be foreclosed at the level of the protocol’s own rules—and asks the harder question that remains: what such a substrate does *not* settle, and under what conditions instruments built upon it can satisfy the prohibitions it leaves open.

*A concept note companion to the κ -rate research program. It states precisely what an interest-free consensus substrate resolves and what remains for jurisprudential adjudication, and is intended as a basis for engagement with qualified Shariah scholarship.

[†]Department of Finance, Independent University Bangladesh (IUB), Dhaka. 2120922@iub.edu.bd.

[‡]Corresponding author. Department of Accounting and Finance, Alabama A&M University, Normal, AL, USA. rafiquel.bhuyan@aamu.edu.

The contribution is twofold. First, it draws the boundary precisely, distinguishing *riba*—a defect of monetary extraction, resolvable mechanically—from *gharar*, *maysir*, and *qabd*, which are defects of contractual ontology, intent, and delivery, and are not. Second, it proposes four product-layer rules—each an established principle of Islamic commercial law—under which asset-backed instruments built on the *riba*-free foundation address the remaining pillars by construction, and it identifies the residual that no architecture can resolve and that properly belongs to jurisprudential judgment. The register throughout is deliberately conservative: the aim is to state what is and is not claimed, so that the line between engineering and jurisprudence is drawn rather than blurred. The remainder treats the *riba*/ontology boundary; then *qabd*, *gharar*, and *maysir* in turn; then the four structuring rules and a comparative taxonomy of what is resolved and what remains.

2 The epistemic boundary: *riba* vs. contractual ontology

The κ -Chain’s consensus layer mechanically eliminates *riba al-nasiah* by enforcing a strict $\iota = 0$ invariant: a transaction whose decoded state transition implies a positive interest parameter is rejected at inclusion, and a proposer who admits one is slashed.¹² That a zero-interest environment is *necessary but not sufficient* is the settled position of the modern critical literature. El-Gamal (2006) and Usmani (2002) criticise the industry’s “Shariah arbitrage”—replicating conventional cash-flows without interest while ignoring the economic substance of the underlying contract. Code can guarantee that money does not breed money; it cannot adjudicate the *ontological reality* of the asset traded. Resolving *riba* therefore *isolates* the remaining defects; it does not cure them. The classical monetary ontology under which $\iota = 0$ is the operationalization of money as a neutral measure—rather than a regulatory expedient—is developed in the companion paper (Ahmed & Bhuyan, 2026); the present note begins where that argument ends.

3 Qabd (possession): the hardest pillar for cash-settled synthetics

A valid sale classically requires transfer of ownership and the assumption of ownership risk (*daman*) before profit is taken. AAOIFI Shariah Standard No. 20 (Sale of Commodities in Organized Markets) and Usmani (2002) require genuine delivery for forward and futures structures; SS-20 in fact declares conventional futures and options impermissible as formed and traded (AAOIFI, 2015), and Kamali (1999) surveys the prospects for an Islamic derivatives market within precisely this constrained space. al-Amine (2008) notes that cash-settled derivatives, trading abstract claims rather than assets in hand, risk violating *bay’ al-kali’ bi-l-kali’* (the prohibited sale of debt for debt; cf. AAOIFI SS-10 on *Salam*).

A blockchain guarantees *qabd hukmi* (constructive possession) of native margin collateral via cryptographic control. It cannot manufacture possession of an external real asset merely *referenced* by an oracle. Cash-settled synthetic settlement is, on this point, a legal fiction.

The classical fiqh of possession is more textured than a single rule. Jurists distinguish *qabd haqiqi* (physical possession—taking the commodity in hand) from *qabd hukmi* (constructive

¹Throughout, the *riba* foreclosed at consensus is *riba al-nasiah*—the deferment-based excess that interest instantiates and that a pricing layer can host. *Riba al-fadl* in spot exchanges of ribawi commodities is a matter of product-layer exchange (*sarf*) rules and falls under the structuring rules of Section 5, not the consensus invariant.

²For the reader outside distributed systems: the substrate is a proof-of-stake blockchain whose consensus rules are specified to make a positive interest parameter unrepresentable in valid state; $\mathbf{x}/\text{shariah}$ denotes its transaction-admission gate, and ABCI++ vote extensions are the mechanism by which validators jointly determine an external price. At the time of writing the consensus layer is a specification; the instruments’ reference implementation runs on a public testnet.

possession—legal control without physical transfer, as when goods are placed at the buyer’s disposal or represented by a recognised document of title). That distinction is exactly what makes a digital analogue conceivable: a cryptographically controlled, registry-backed token over a real, identified asset can plausibly constitute *qabd hukmi*, in the way a warehouse receipt or a registered title does. What it cannot do is constitute possession of an asset that does not exist within, or is merely referenced by, the system—the situation of a cash-settled synthetic whose “underlying” is an external price index. The *salam* contract is the instructive exception that proves the rule: classical law permits the sale of a described, deferred-delivery commodity precisely *because* delivery of a real, specified good is contracted for and the price is paid in full at contract—conditions the standards on organised markets preserve and which cash settlement removes. The requirement throughout is *daman*: that the party claiming profit has first borne the risk of ownership.

4 Gharar (excessive uncertainty): the non-existent underlying

al-Dharir (1997) defines *gharar* as arising chiefly where the object of contract is non-existent, undeliverable, or unspecified. al-Amine (2008) and Usmani (2002) argue synthetic derivatives carry a fundamental *gharar* because the subject matter is a mathematical abstraction (a price index), not a tangible commodity — *la tabi’ ma laysa ’indak* (“do not sell what you do not possess”); Dusuki & Abozaid (2007) situate this within the broader *maqasid* critique of form-over-substance engineering.

An ABCI++ BFT-median oracle eliminates *measurement* *gharar*—the price is exact, transparent, and manipulation-resistant. It cannot eliminate *ontological* *gharar*: that the parties trade a feed, not a deliverable.

Not all uncertainty voids a contract; the jurisprudence is graduated. al-Dharir (1997) distinguishes *gharar fahish* (excessive, contract-voiding uncertainty) from *gharar yasir* (minor, tolerated uncertainty inseparable from ordinary commerce), and locates the prohibited kind chiefly in three conditions: that the object of contract be non-existent, that it be undeliverable, or that its essential attributes be unspecified. A cash-settled synthetic can fail on the first two at once: there is no object to deliver, because the “object” is a number derived from a price feed. This is what separates the *gharar* a protocol can cure from the *gharar* it cannot. A deterministic oracle removes uncertainty about *measurement*—the price is exact, public, and manipulation-resistant, so the specification defect (*jahala*) of an ambiguous reference price is eliminated—but no oracle can supply an object where there is none. The asset-backing rule below is the direct response: by requiring that the object of contract be a real, existing, deliverable asset, it converts the prohibited *gharar fahish* of the non-existent underlying back into the tolerated *gharar yasir* of ordinary commercial risk.

5 Maysir (gambling) vs. tahawwut (legitimate hedging)

Islamic economics distinguishes *tijara* (legitimate commercial risk) from *maysir/qimar* (zero-sum gambling). Bacha (1999) and Obaidullah (1998) observe that in cash-settled futures one party’s gain exactly equals the counterparty’s loss, driven by price movement rather than real value creation—structurally close to *qimar*. A *minority* view (Kamali, 2000) holds that a derivative used strictly for *tahawwut* (hedging) serves a valid *maslaha*; but distinguishing hedger from speculator requires assessing intent (*niyya*) and off-chain exposure. A central limit order book is blind to *niyya*: it cannot tell a farmer hedging a real harvest from a trader gambling at leverage. The execution layer cannot filter *maysir*.

The line between *maysir* and legitimate risk is among the most contested in the field, and the contest is substantive rather than terminological. *Qimar* (gambling) is classically defined as

an arrangement in which each party’s gain is contingent on the other’s loss, with no creation of real value and no productive risk borne—a description that fits a naked, cash-settled position uncomfortably well. The defence of derivatives in Islamic law turns, accordingly, not on denying this structure but on the *purpose* the instrument serves: Kamali (2000), in the most sustained argument for permissibility, holds that where a derivative is used for *tahawwut*—genuine hedging of a real economic exposure—it advances a recognised *maslaha* (public interest) and should not be assimilated to gambling. The difficulty, which Kamali concedes, is evidential: hedging and speculation are distinguished by the trader’s purpose and off-contract position, neither of which is legible to a matching engine. This is why the structural response below does not attempt to read intent but to remove the *occasion* for pure speculation—by requiring, through the insurable-interest rule, that a protection buyer hold the exposure being protected, so that the instrument is necessarily a hedge of something owned rather than a wager on something not.

The aggregation objection. A sophisticated objection runs: for a *diversified* holder of many κ -instruments, the law of large numbers renders the portfolio’s return near-deterministic — is the result not functionally a fixed return, *riba* by another route? We submit the question to the reader’s judgment with one observation from the classical materials: the prohibition attaches to the *contract* — a predetermined excess owed by a specific counterparty irrespective of outcome — and not to the statistical stability that *emerges* when many genuinely risk-bearing positions are aggregated. A *mudarabah* fund spread across a thousand ventures, a *takaful* operator’s pooled book, a broadly diversified equity portfolio: each yields stable aggregate returns, and none has been held to constitute *riba* on that account, because every constituent position remains individually loss-bearing and no party owes the stability by obligation. Whether this distinction is dispositive for κ -portfolios is, like everything in this note, for qualified scholarship to rule upon; we record the precedent so the question arrives with its strongest available answer attached.

6 Structuring for compliance: four product-layer rules

The constructive answer is not to claim code resolves *gharar*, *maysir*, and *qabd*—it cannot—but to make the *compliant instrument class* satisfy them structurally, enforced at the same consensus gate (**x/shariah**) that the design already specifies to enforce $\iota = 0$, leverage caps, and full-notional (anti-naked) collateral. Four rules complete it. The rules themselves are not novel: asset-backing, insurable interest, the prohibition of naked positions, and *tabarru* risk-sharing are established compliance principles (Kamali, 2000; AAOIFI, 2015; El-Gamal, 2006). The contribution here is their enforcement as *consensus-level invariants* on a *riba*-free substrate, and the explicit boundary that separates them from the *riba* the substrate has already foreclosed.

1. **R1 — Asset-backing whitelist** (*gharar*). A tradable instrument must reference an on-chain token representing a real, registered, deliverable asset (sukuk certificate, warehouse receipt, commodity or property title), not a bare price feed. The oracle resolves measurement *gharar*; asset-backing resolves existential *gharar* (the subject of contract exists and can be delivered).
2. **R2 — Possession-gating** (*qabd*). Enforce *qabd hukmi*: a party must hold the asset-token (cryptographic control plus a registry record, e.g. the Qatar Financial Centre’s (QFC) Digital Receipt System) before reselling forward or taking profit. The gate rejects any forward or sell order from a party not in possession—the on-chain form of *la tabi’ ma laysa ’indak*—so ownership risk (*daman*) rests on the holder before resale.

3. **R3 — Insurable-interest / no-naked rule** (*maysir* and *qabd*) for credit instruments. For an Islamic CDS, require the protection *buyer* to hold the reference obligation (insurable interest), converting a naked wager into a *kafala* (guarantee) over an owned, real exposure; the protection seller posts full notional (anti-naked, part of the gate’s specification and enforced in the reference implementation). With neither side naked, the contract is risk transfer on a real owned exposure, not a zero-sum bet — the structural (naked, zero-sum) surface of *maysir* removed by construction, with the intent residual remaining for the Board rather than being policed by code.
4. **R4 — Expiry / delivery and tabarru wrappers** (*gharar* and *maysir*). Prefer expiring, deliverable structures (*salam*, *murabaha*, *ijara*, expiring options on owned assets) over perpetual cash settlement, and use *tabarru* (mutual-donation) takaful pools for risk products — risk-sharing tied to real loss, not a price bet.

The honest residual. The intent (*niyya*) component of *maysir* is not code-governable. R1–R4 remove the naked, zero-sum surface; the remainder is for the Board through access classes (hedging-only or institutional tracks vs. open retail) and position-purpose attestation. Code minimises the surface; the Board rules on the residual.

Two-track architecture. *Track A*—the perpetual-futures DEX—is *riba*-free ($\iota = 0$) but cash-settled and synthetic, and therefore remains subject to an open fatwa on *gharar/maysir*. *Track B*—the asset-backed compliant class (*sukuk*, commodity *murabaha*, tokenised real-world assets, iCDS-as-*kafala*-on-owned-exposure, takaful)—addresses *gharar* and *qabd* by construction and removes the structural (naked, zero-sum) component of *maysir* via R1–R4, on the same *riba*-free substrate, with the intent residual reserved to the Board and final adjudication of all three pillars resting with qualified scholarship. The two tracks also order the launch sequence: asset-linked instruments—e.g. protection on owned *sukuk* exposures, vault-secured tokenised commodities, and a constant-maturity *sukuk*-yield benchmark for institutional liquidity management—precede any cash-settled synthetic, which remains the hard case pending the open fatwa.

Two worked examples make the rules concrete. Consider, first, a tokenised sovereign *sukuk* traded forward on the compliant venue. Asset-backing is satisfied because the token represents a registered certificate over a real, income-producing asset, not a price index (R1); possession-gating requires that a seller hold the token—and so bear ownership risk—before transferring it forward, the transfer recorded as *qabd hukmi* (R2); and the instrument is an expiring, deliverable claim rather than a perpetual cash settlement (R4). Consider, second, an Islamic credit-protection contract over that *sukuk*. The protection seller posts the full notional as collateral, so no uncovered (naked) protection can be written (the specified anti-naked rule); the insurable-interest rule additionally requires the protection *buyer* to hold the reference *sukuk*, converting the contract from a wager on a third party’s default into a *kafala*—a guarantee of an exposure the buyer actually carries (R3). In each case the instrument that clears the gate is not a synthetic abstraction but a claim tied, at both ends, to a real and owned asset. The rules do not make speculation impossible; they make the compliant class structurally a class of hedges and real-asset transactions rather than bets.

7 Comparative taxonomy

8 The *riba* question, measured (addendum, 2026)

The boundary this paper draws can now be located empirically. An out-of-sample validation of the κ -corpus (Ahmed & Bhuyan, 2026, P15) confirms that ι —the predetermined, guaranteed component that defines *riba al-nasiah*—is removable in the data (the conventional-venue funding floor vanishes at $\iota = 0$, the systematic transfer goes to zero); the removability is given a rigorous

Pillar	Core defect	Resolved by	Residual risk
Riba	yield on idle capital / time	substrate ($\iota = 0$, consensus; <i>riba al-nasiah</i>)	<i>riba al-fadl</i> in spot exchange: product layer (<i>sarf</i> rules)
Gharar	asset non-existence / non-delivery	product (asset-backing, R1); oracle for measurement	synthetics: ontological gharar
Qabd	ownership/risk before profit	product (tokenised RWA + receipts, R2)	synthetics: possession a fiction
Maysir	zero-sum gambling	product (<i>tabarru</i> + real-exposure, R3–R4)	intent (<i>niyya</i>) not code-governable

Table 1: What the substrate resolves and what remains: pillar, defect, locus of resolution, residual.

footing in the Separation Theorem, where the pricing operator is analytic at $\iota = 0$ for any positive convergence intensity (Ahmed & Bhuyan, 2026, P17). What remains is a return of the permitted kind, and a companion paper (Ahmed & Bhuyan, 2026, P16) measures how much of it is realised hazard versus risk premium. Decomposing the convergence intensity into the physically realised intensity κ_P and the priced intensity $\kappa_Q = s/(1 - \delta)$, the risk-premium share $1 - \kappa_P/\kappa_Q$ varies across domains by almost two orders of magnitude: a measured $\sim 18\%$ for takaful (premiums equal expected loss plus a margin), and the great majority for sovereign credit (bounded below near 45% by a short default record, plausibly $\sim 85\%$), highest for the safest credit and largest in crises.

This sharpens, and does not dissolve, the reservation this paper makes. It shows that the open question is not “is κ interest?”—the predetermined charge is gone—but the narrower and genuinely contested one: *is a fully contingent, loss-bearing risk premium, stripped of every predetermined charge, riba in substance, or the permitted form of return under al-ghunm bil-ghurm?* Under *al-ghunm bil-ghurm*—entitlement to gain tied to liability for loss, the basis of *mudarabah* and *musharakah*—a contingent, loss-bearing, non-predetermined return is the permitted form of profit; rarity of realised loss does not convert risk-bearing into a stipulated increment. The stricter view holds that a systematic, priced return is the discount rate in substance, and that pricing default risk on debt raises *bay‘ al-dayn* and *gharar* independently. The decomposition confirms both concerns are real and cannot adjudicate between them. Its practical corollary is sharp: the realised-hazard products (takaful) are sound on form *and* substance and need no resolution of the open question, while the credit and monetary layer stands on the contested ground and, by the safest-credit result, on its most premium-laden part—a frontier gated explicitly on the Board’s ruling, R3/R4 (§6) being the structural pre-conditions and this the substantive question reserved.

A pricing resolution that narrows the reservation. The structural rules R1–R4 already recast credit protection as a *kafala* over an owned, possessed exposure inside a *tabarru* pool—form-sound by construction. The decomposition adds a resolution on *substance* available at the product layer *without* a prior ruling on the open question. Because κ separates into a realised part κ_P and a priced part κ_Q , the compliant cooperative-guarantee form can price its contribution at the *realised hazard* $\kappa_P(1 - \delta)$ —the actuarially fair expected loss, exactly as a *takaful* contribution is set—rather than at the market-priced $\kappa_Q(1 - \delta)$ a tradeable spread quotes. The contested risk-premium wedge $1 - \kappa_P/\kappa_Q$ is then not *sold* to a protection writer for profit but retained by the mutual pool and returned as surplus, with a disclosed *wakala* fee for administration. Priced this way, credit-loss cover is structurally the takaful case this paper already holds riba-free in substance: a mutual indemnity, gated on insurable interest and *qabd hukmi* over the reference obligation, funded at expected loss. The reservation therefore narrows sharply. The open question is no longer the broad “is a priced risk premium riba?” but the

two residuals that survive the κ_P -priced mutual structure— whether tokenised title constitutes valid *qabd* for the reference obligation, and whether the *wakala* loading is genuinely a fee for service rather than the risk premium reintroduced under another name—both narrower, both answerable, both reserved to the Board. We do not represent this as settling the matter; we record that the product layer can meet the boundary far closer than a cash-settled spread priced at κ_Q could, and that where it cannot reach is now precisely located.

9 Conclusion and request for judgment

The $\iota = 0$ substrate is a definitive, structural resolution of *riba al-nasiah*—a clean, interest-free ledger and a market-determined, non-negative pricing measure. It is a necessary foundation, not a sufficient condition, for holistic compliance. *Gharar*, *maysir*, and *qabd* are matters of asset ontology, contractual intent, and delivery—the domain of jurisprudential interpretation. It is on these precise boundaries—how to construct, restrict, collateralise, or asset-back synthetic pricing environments built on a *riba*-free substrate—that qualified Shariah scholarship must adjudicate; this note sets them out so that the line between what an architecture can settle and what it cannot is drawn explicitly rather than elided.

There is a more general point in this division of labour. The *maqasid* tradition holds that the prohibitions of commercial law are not arbitrary constraints but protections of identifiable goods—the integrity of exchange, the linkage of return to real risk, the circulation rather than the hoarding of wealth. Read in that light, an interest-free substrate and the four structuring rules are not a checklist of avoidances but a positive design: a system whose default state is the exchange of real, owned, deliverable value, in which interest is absent because the architecture has no place to put it and speculation is curtailed because the instruments are tied to ownership. What such a system cannot do is legislate the intention of those who use it, and it should not pretend to. The honest position—and, we have argued, the strongest one before a jurisprudential audience—is to build the foundation as soundly as the engineering permits, to state plainly the boundary at which engineering ends, and to refer what lies beyond it to those whose competence it is.

References

- AAOIFI (2015). *Shariah Standards*, SS-20 (Sale of Commodities in Organized Markets) and SS-10 (Salam). Accounting and Auditing Organization for Islamic Financial Institutions, Bahrain.
- Ahmed, S., & Bhuyan, R. (2026). Money as a Mirror: Classical Islamic Monetary Ontology and the Foundations of Interest-Free Pricing. Working paper, κ -rate program.
- al-Amine, M. al-Bashir M. (2008). *Risk Management in Islamic Finance: An Analysis of Derivatives Instruments in Commodity Markets*. Brill (Arab and Islamic Laws Series 1).
- al-Dharir, S. M. al-Amin (1997). *Al-Gharar in Contracts and Its Effect on Contemporary Transactions*. Islamic Research and Training Institute (IRTI), Eminent Scholars' Lecture Series 16.
- Bacha, O. I. (1999). Derivative Instruments and Islamic Finance: Some Thoughts for a Reconsideration. *International Journal of Islamic Financial Services* 1(1).
- Dusuki, A. W., & Abozaid, A. (2007). A Critical Appraisal on the Challenges of Realizing Maqasid al-Shariah in Islamic Banking and Finance. *IIUM Journal of Economics and Management* 15(2), 143–165.
- El-Gamal, M. A. (2006). *Islamic Finance: Law, Economics, and Practice*. Cambridge University Press.
- Kamali, M. H. (2000). *Islamic Commercial Law: An Analysis of Futures and Options*. Islamic Texts Society, Cambridge.

- Kamali, M. H. (1999). Prospects for an Islamic Derivatives Market in Malaysia. *Thunderbird International Business Review* 41(4–5), 523–540.
- Obaidullah, M. (1998). Financial Engineering with Islamic Options. *Islamic Economic Studies* 6(1), 73–103.
- Usmani, M. T. (2002). *An Introduction to Islamic Finance*. Kluwer Law International.
- Ahmed, S., & Bhuyan, R. (2026). The κ -Framework Meets the Data: An Out-of-Sample Empirical Validation Across Seventeen Data Providers.
- Ahmed, S., & Bhuyan, R. (2026). The Riba Question, Reduced: A Risk-Premium Decomposition of the Convergence Intensity.
- Ahmed, S., & Bhuyan, R. (2026). The Separation Theorem: Interest as a Removable Parameter in No-Arbitrage Pricing.

EDITORIAL NOTE TO THIS PRINTING

Paper 12 · Liveness under Macro Winters

Working paper, revised July 2026 · SSRN 6924459

Role in the programme. The stress question: does a measured primitive survive when markets go quiet or hostile. **What it establishes.** Liveness behaviour of the κ estimator and the chain's economics under funding droughts and macro stress scenarios. **Corrected or extended later.** No substantive corrections to date. **Not established here.** Behaviour under conditions outside the simulated scenario set; mainnet economics.

Replication for every number in this paper: DOI 10.5281/zenodo.21290557 · github.com/shehzadahmed-xx/kappa-replication

Liveness under Macro Winters: Mechanism Design for Fee-Only Blockchains

Shehzad Ahmed*

June 2026

Abstract

Every blockchain whose consensus layer is financed purely by transaction fees — by design (zero-emission chains) or by destiny (Bitcoin as its block subsidy decays) — faces a liveness hazard that emission-funded chains conceal: in the aftermath of a persistent macroeconomic shock, transaction volume can remain below validator operating cost for months. We show the binding risk is the *duration* of such “winters,” not their depth, and that the intuitive remedy — counter-cyclical fee multipliers — is both pro-cyclical for users and manipulable by the validators who produce the volatility measurements. We propose instead a pre-funded mutual support pool (a *liveness sleeve*) with three design commitments: a *conjunction trigger* (volume and a stress index, each sustained, so that the one signal validators control never suffices alone); a *shortfall payout* capped at a published reference operating cost; and a slow payout decay whose role, we show, is not to discourage idleness but to make a validator’s receipts strictly increasing in its own honest fee income — without decay, censoring one’s own fee base is payoff-neutral inside the subsidised state. Under the shortfall structure, censorship aimed at forcing the trigger is strictly dominated whenever the support floor lies below normal income (Proposition 1). Monte-Carlo experiments on six-year paths with heterogeneous validators locate the feasible design window — replacement ratio $\beta \in [1.2, 1.6]$ of reference cost with decay at most 4% per week — and a reinforcement-learning adversary (a 35%-stake cartel free to censor up to 75% of transactions) converges to near-zero censorship, never beating honesty on matched evaluation episodes. Survival of a thirty-validator set through a black-swan winter (volume -85% , half-life twelve months) rises from 56% to 77% with the sleeve in place. The mechanism requires no emission, no fee increases on distressed users, and no discretionary bailout authority.

Keywords: blockchain security budget, validator economics, mechanism design, fee markets, liveness, mutual insurance, zero-emission consensus. **JEL:** D82, G22, L86.

1 Introduction

The long-run security of a proof-of-stake or proof-of-work blockchain rests on a simple flow constraint: the consensus layer’s income must cover its operating cost. Most chains today satisfy the constraint with predetermined issuance — inflationary block rewards that pay validators for the passage of time regardless of network activity. A growing class of designs rejects issuance, on monetary grounds (zero-emission “sound money” designs; interest-free substrates in which a predetermined time-based return is prohibited by construction (Ahmed, 2026b)) or has it taken away by schedule: Bitcoin’s block subsidy halves toward zero, and

*Independent University, Bangladesh. The mechanism analysed here originates in the validator-economics design of an interest-free Layer-1 specification (Ahmed, 2026b); the liveness problem and the mechanism are general to any chain whose consensus rewards are funded by transaction fees alone.

the question of whether fees alone can fund its security — the *security budget problem* — is well posed in the literature (Carlsten et al., 2016; Budish, 2018; Easley et al., 2019).

That literature has concentrated on two failure modes: strategic deviations in block production when fee income is volatile at the timescale of blocks (Carlsten et al., 2016), and the adequacy of steady-state fee income against attack costs (Budish, 2018). This paper analyses a third, slower failure mode that emission-funded chains never exhibit and the fee-market literature has largely passed over: the **macro winter**. A severe macroeconomic shock produces, first, a burst of activity — liquidations, de-risking, repricing — during which fee income *rises*; and then a long aftermath in which volume settles far below its prior level and recovers only slowly. For a fee-only consensus layer, the acute phase is self-financing; the aftermath is the threat. If the winter outlasts validators’ working capital, the chain loses operators precisely when the remaining operators’ incentives are weakest, compounding censorship risk, liveness risk, and stake concentration.

Two intuitive remedies fail. *Counter-cyclical fee multipliers* — raising protocol fee floors during stress to keep validators whole — are pro-cyclical for users (they tax distressed transactors to preserve infrastructure margins) and create a perverse measurement problem: the validators who would benefit from the multiplier are the same agents who produce, order, and can distort the on-chain measurements of stress. *Discretionary treasuries* reintroduce the trusted bailout authority that the architecture was designed to exclude. We propose and analyse a third route: a rule-based, pre-funded mutual support pool — a *liveness sleeve* — filled counter-cyclically from fees in good times and disbursed under a conjunction trigger with a shortfall payout in certified winters.

Contributions. *First*, we formalise the winter-liveness problem and show its binding dimension is shock persistence, not depth: the support pool a chain needs diverges as the volume process’s autocorrelation approaches unity, mirroring (and inheriting the mathematics of) buffer-sizing results for macroeconomic stabilisation funds (Ahmed & Bhuyan, 2026). *Second*, we give the mechanism three commitments whose interaction is the design’s substance: (i) a conjunction trigger over volume and an independent stress index, each sustained over a hysteresis window, with *fixed* governance constants rather than trailing quantiles (trailing windows are gameable by slow grinding); (ii) a shortfall payout $\max(0, \beta \cdot \text{OpEx}_{\text{ref}} - y_i)$ against a *published* reference cost, performance-weighted by consensus-measurable work; (iii) a slow payout decay. *Third*, we prove the central incentive result (Proposition 1): under the shortfall structure, volume censorship aimed at forcing or exploiting the trigger is strictly dominated whenever the support floor lies below normal income — and we show the decay parameter’s true role is to break the payoff-neutrality of self-censorship *inside* the subsidised state, not to discourage idleness. *Fourth*, we quantify the design window by simulation (six-year paths, heterogeneous validators, joint volume-and-stress crises) and stress the incentive claim with a learned adversary: a PPO-trained 35%-stake censor cartel converges to near-zero censorship and never beats honesty on matched episodes. *Fifth*, we derive sizing guidance: the pool’s cap should be set by the tail winter’s *duration* (twelve months of reference cost in our calibration), and the replacement ratio $\beta > 1$ should be read as a dispersion allowance over the audited validator cost distribution, not a profit margin.

Related work. The security-budget line begins with Nakamoto (2008)’s subsidy schedule and its consequences: Carlsten et al. (2016) show fee-only mining is unstable at the block timescale; Budish (2018) bounds the steady-state security a fee flow can buy; Easley et al. (2019) document the economics of the fee market’s emergence. The fee-mechanism line — EIP-1559 and its analyses (Buterin et al., 2019; Roughgarden, 2021), and equilibrium fee formation (Huberman et al., 2021) — treats congestion pricing within a period, not

income insurance across periods; our mechanism is complementary and operates at a different timescale. On the funding side, [Saleh \(2021\)](#) analyses proof-of-stake without wasteful issuance. The sleeve’s mutual, donation-based structure has a long institutional pedigree in mutual insurance; its specific instantiation here originates in a takaful (mutual protection) pool specified for an interest-free Layer-1 ([Ahmed, 2026b](#)), where a predetermined validator return would violate the substrate’s own monetary constraint — a case in which the incentive requirement and the monetary-design requirement coincide (Section 6).

2 The winter-liveness problem

Consider a chain whose consensus layer earns only transaction fees. Normalise aggregate fee income in the normal state to 1 per period, of which a fixed wakala-style share w (here $w = 0.6$) accrues to validators; let aggregate reference operating cost be OpEx_{ref} (here 0.36, a 40% normal-state margin). Validator i bears cost c_i with $\sum_i c_i = \text{OpEx}_{\text{ref}}$, finances deficits from working capital up to a credit line (two months of own cost in our calibration), and exits when the line is exhausted. Volume follows a log-AR(1) with rare crisis jumps; a stress index κ_t (in our calibration, the chain’s published credit-intensity state, with monthly autocorrelation 0.774 estimated from real instrument panels ([Ahmed & Bhuyan, 2026](#))) jumps jointly with crises.

Two features structure the problem. *First*, the acute phase of a crash is self-financing: liquidations and de-risking raise fee income exactly when volatility peaks. The threat is the aftermath — volume 60–85% below normal with a recovery half-life of months. *Second*, the required support diverges in persistence: for a winter with depth D and recovery half-life h , the cumulative income shortfall scales like $D \times h$, so a pool adequate for moderate winters fails in persistent ones at any depth. It is always the length of the winter, not its severity, that exhausts a finite reserve. Both features are inherited by any insurance mechanism built on top, and both militate against financing the gap out of *contemporaneous* user fees.

3 The mechanism

Definition 1 (Liveness sleeve). *A liveness sleeve is a protocol-owned pool S_t with the following rules. **Fill:** in any period with volume above its long-run median and the sleeve inactive, a skim s (here 10%) of fees accrues to S_t , stopping at cap $\bar{S} = C$ months of OpEx_{ref} ; a genesis allocation seeds part of the cap. **Trigger:** the sleeve activates only if measured volume is below $\underline{V}$ for K consecutive periods and the stress index exceeds a fixed constant $\bar{\kappa}_{99}$ for K consecutive periods. **Payout:** while active, validator i with fee income $y_{i,t}$ receives*

$$\pi_{i,t} = \omega_{i,t} \max(0, \beta c_i^{\text{ref}} - y_{i,t}) (1 - d)^{\tau_t},$$

where c_i^{ref} is the published reference cost, $\omega_{i,t} \in [0, 1]$ a performance weight built from consensus-measurable work (blocks signed, vote participation, oracle submissions within tolerance of the stake-weighted median, uptime), d a weekly decay, and τ_t weeks since activation. **Exit:** deactivation requires volume above $1.3 \underline{V}$ sustained one week, then a one-month cooldown. **Governance:** $(\underline{V}, \bar{\kappa}_{99}, \beta, d, C)$ are fixed constants amendable by ordinary governance within board-bounded invariants; the reference cost is set by published methodology, not validator self-report.

Three commitments deserve comment before the analysis. The *conjunction* exists because measured volume is the one trigger input validators control: a cartel can censor transactions to suppress it. An independent, manipulation- expensive stress leg (here, a stake-weighted-median oracle state whose falsification requires majority collusion and is slashable ex post)

must therefore co-sign every activation. The *fixed constants* exist because trailing-quantile thresholds — the adaptive alternative — can be ground down slowly by sustained, mild suppression. The *shortfall* form of the payout is the heart of the mechanism, as the next section shows.

4 Censorship is strictly dominated

Fix a cartel controlling stake share λ and consider any strategy that suppresses a fraction $\chi_t \in [0, 1]$ of network transactions in period t . Censored transactions do not execute, so the cartel’s fee income falls one-for-one with the fee base it destroys. Let $y_t(\chi)$ denote aggregate validator fee income under suppression χ , strictly decreasing in χ .

Proposition 1 (Shortfall dominance). *Suppose the support floor lies below normal income, $\beta \text{OpEx}_{\text{ref}} < y^{\text{normal}}$, and the decay satisfies $d > 0$. Then for every period and state: (i) outside the active state, any $\chi_t > 0$ strictly reduces the cartel’s payoff; (ii) inside the active state, the cartel’s receipts $y_t(\chi) + \max(0, \beta \text{OpEx}_{\text{ref}} - y_t(\chi))(1 - d)^\tau$ are strictly increasing in $y_t(\chi)$, hence strictly decreasing in χ_t ; and (iii) any strategy that incurs censorship costs to cause activation is strictly dominated by honesty, since the payout in the activated state is capped below the income the cartel destroys to reach it.*

Proof. (i) Outside activation there is no payout; receipts equal fee income, which is strictly decreasing in χ . (ii) Inside activation, receipts as a function of own income y are $R(y) = y + (\beta \text{OpEx} - y)\delta$ on $y < \beta \text{OpEx}$ with $\delta = (1 - d)^\tau < 1$, so $R'(y) = 1 - \delta > 0$; and $R(y) = y$ above the floor. Receipts are therefore strictly increasing in income everywhere, and censorship, which only destroys income, strictly reduces them. (iii) To move the system from inactive to active, the cartel must hold measured volume below $\underline{V}$ for K periods, destroying at least $\lambda w (1 - \underline{V}) y^{\text{normal}} K$ of its own income with zero payout during the approach (the sleeve is inactive), to reach a state whose per-period receipts are bounded by $\beta \text{OpEx}_{\text{ref}} < y^{\text{normal}}$ and declining in τ . Every such path is pointwise dominated by the honest path. $\square$

Remark 1 (The decay’s true role). *At $d = 0$, part (ii) fails at the margin: $R'(y) = 0$ below the floor, so a cartel already inside the active state is indifferent to destroying its own fee base — censorship becomes costless at the margin, and the mechanism loses its strict incentive to keep producing blocks that carry fees. The decay’s purpose is precisely to make receipts strictly increasing in honest income inside the subsidised state. Its anti-idleness framing is secondary: in our simulations the trigger’s false-positive rate against ordinary downturns is zero, so the “lazy validator in a minor dip” never reaches the sleeve at all. This inverts the Bagehot intuition with which we began the design: penal pricing of the support is neither necessary (the conjunction gate and the performance weights carry the selection burden) nor sufficient (aggressive decay guts protection in long winters); what is essential is that support never exceed the floor and that receipts always rise with honest work.*

5 Quantitative design

5.1 Reference model and calibration

We simulate six-year paths at hourly resolution (and replicate at daily resolution inside a cadCAD engine (Ahmed, 2026a)) with thirty validators of lognormally dispersed cost, the volume and stress processes of Section 2, and Poisson crises: minor (volume -35% , half-life one month, roughly annual) and major — the winters — in three severities: moderate (-65% , half-life six months), severe (-75% , nine months), black swan (-85% , twelve months).




Figure 1: A matched severe winter. Top to bottom: volume against the trigger floor $\underline{V}$; the stress index against its $p99$ constant; the pool in months of reference cost; payouts during the activation episode.

The stress index is calibrated to the full-universe panel-median intensity series (273 bonds aggregated to a 127-month index, 2015–2026): measured monthly AR(1) $\rho = 0.921$ — an 8.4 -month half-life, substantially more persistent than single-issuer estimates, confirming that winters are the relevant regime — with the $p95/p99$ governance constants taken from the real distribution. The real decade contains two stress episodes: the 2020 episode (peak jump +305bp of spread, four months above $p95$, measured decay half-life three months from peak) and a 2016 episode (+200bp, slower decay). The measured 2020 episode is, in intensity units, roughly $2.5\times$ larger than the scenario severe winter (+762bp vs. +300bp of κ) but decays much faster (three months vs. nine); both are run below. Volume remains a scenario input, as no fee-only chain history yet exists. Figure 1 shows one matched winter: the joint shock, the pool’s drawdown, and the payout episode.

5.2 Trigger discrimination and survival

Across all simulated paths without a true winter, the conjunction trigger fired **zero** times: ordinary downturns never reach the sleeve, so no minor-downturn moral hazard arises at the trigger stage. History gives the same answer: on the real 127-month index, the κ gate was open (sustained above $p99$) in just **two months of the decade** (1.6%; seven months at $p95$) — outside those windows no volume manipulation, however aggressive, could have fired the sleeve. Through forced winters at the validated parameters ($\beta = 1.3$, $d = 2\%$ /week, cap twelve months, genesis seed four months), survival of the validator set improves monotonically in severity — the sleeve is a tail instrument. The table reports the daily cadCAD implementation; the hourly reference model yields the same ordering with survival several points higher in every cell (daily aggregation deepens effective drawdowns against a fixed credit line), and all qualitative conclusions — the monotone severity gradient, the near-exhaustion of the pool only in the black-swan case, and the location of the feasible window — are common to both:

Winter	without sleeve	with sleeve	Δ
moderate (−65%, hl 6mo)	85.1%	90.2%	+5.1pp
severe (−75%, hl 9mo)	72.6%	84.4%	+11.8pp
black swan (−85%, hl 12mo)	56.3%	76.5%	+20.2pp

The pool approaches exhaustion only in the black-swan case, which is what sizes the cap: a six-month cap — adequate for every moderate winter — fails exactly in the persistent tail, the quantitative form of the persistence-divergence point of Section 2. The *measured* 2020 episode makes the same point with real data: replayed through the model (+762bp measured jump, measured three-month decay), it is nearly harmless — 96% of validators survive *without any sleeve* — despite a stress jump $2.5\times$ the scenario severe winter’s. A huge fast shock is survivable; the dangerous winter is the slow one. Persistence, not magnitude, is the binding risk — now an observation, not only a simulation.

5.3 The feasible window

Sweeping (β, d) on severe winters under the joint criterion *survival* $\geq 90\%$ and *subsidised income* $\leq 85\%$ of normal locates the feasible window at $\beta \in [1.2, 1.6]$, $d \leq 4\%$ /week (Figure 2). Two readings matter. First, our own initial calibration — $\beta = 0.8$ with 10% /week decay, the Bagehot-flavoured “penalty rate” — lies outside the window: it is incentive-safe but protects almost nothing in long winters, improving severe-winter survival by under two percentage points. Second, $\beta > 1$ is not profit: the reference cost is an aggregate, and the upper tail of a lognormal cost distribution needs $1.2\text{--}1.4\times$ the reference to reach its own floor. The binding constraint throughout the window is protection, not incentives — subsidised income runs at 71–77% of normal at $\beta = 1.3$, so the preference inversion (no validator prefers the subsidised state) holds with room to spare.

5.4 A learned adversary

Proposition 1 covers censorship strategies; simulations can stress its premises against strategies we did not enumerate. We train a PPO agent (Schulman et al., 2017) playing a 35%-stake cartel that chooses, daily, a censorship intensity in $\{0, 25\%, 50\%, 75\%\}$, observing measured volume, the stress ratio, the sleeve state, and the pool level, and rewarded with its stake share of total validator income (fees plus payouts). Episodes run three years; half contain a severe winter, so exploitable stress states, if any, appear in training. After 150,000 steps the learned policy converges to mean suppression 1.4% — effectively honesty — and on thirty matched evaluation episodes it beats the zero-censorship policy **never** (mean per-episode




Figure 2: The feasible window on severe winters: survival (left) and subsidised-income ratio (right) over (β, d) ; the contour bounds the jointly feasible region; the star marks the initial Bagehot-flavoured calibration, which fails the protection criterion.




Figure 3: The learned censor-cartel versus honesty on matched seeds: cumulative cartel income (left) and the distribution of the learned policy’s advantage (right) — negative in every episode.

disadvantage -0.39 income units; Figure 3). The agent’s verdict agrees with the proposition: the trigger can sometimes be forced, but forcing it never pays.

5.5 Reproducibility

The hourly reference model, the daily cadCAD implementation (partial state update blocks), the adversary environment and PPO training script, the trained policy, the parameter grids behind Figure 2, and all figure generators are included in the replication package accompanying this paper; every number in this section is reproducible from a single script invocation per experiment.

6 When incentive design and monetary design coincide

Run the mechanism’s failure mode in reverse. A badly parameterised sleeve — a generous floor, no decay, stake-weighted rather than performance-weighted disbursal — pays validators a predetermined amount for the passage of time, regardless of work done. That is, structurally,

an inflationary block reward funded from a pool rather than from issuance: exactly the predetermined time-based return that zero-emission designs exist to exclude, and that interest-free monetary substrates prohibit as a matter of construction (Ahmed, 2026b). The conditions that keep the sleeve incentive-compatible — payments that are outcome-contingent (shortfall only), earned (performance-weighted), bounded (capped at audited cost, decaying), and exceptional (conjunction-triggered) — are the same conditions that keep it from being a disguised fixed return. We do not claim this coincidence is deep in general; we observe that in this mechanism the incentive-compatible point and the no-predetermined-return point are the same point, which suggests that “returns must track contemporaneous real work” is doing the load-bearing work in both arguments.

7 Limitations

These results are design-stage. The volume process is a scenario calibration — no fee-only chain has yet lived through the winters we simulate — though the stress side is now fully measured: the index dynamics, the governance constants, and the major-winter magnitude are taken from the 127-month full-universe record rather than assumed. Validator cost dispersion is assumed lognormal; the reference-cost audit process is modelled as exact. The adversary’s strategy class is censorship intensity; selective censorship, timing games against the cooldown and hysteresis windows, oracle-leg manipulation by majority coalitions, and collusion with off-chain positions are future work, as is an implementation in a production consensus engine. Proposition 1 is robust to these extensions in one direction only: any strategy whose effect on the trigger passes through destroying the cartel’s own fee base inherits the dominance argument; strategies that move the stress leg without touching fee income do not, which is why the stress leg must remain majority-collusion-expensive and slashable.

8 Conclusion

Fee-only consensus layers — whether by zero-emission design or by subsidy decay — face a winter-liveness problem whose binding dimension is persistence. A pre-funded liveness sleeve with a conjunction trigger, a shortfall payout against a published reference cost, and a slow decay protects the validator set through tail winters (black-swan survival +20 percentage points in our experiments) while making censorship strictly dominated and teaching a learned adversary to be honest. The design inverts the Bagehot instinct: the support should be *adequate and strictly increasing in honest work*, not punitive; selection belongs to the trigger, not the price. The cap belongs to the tail winter’s duration, and the replacement ratio to the cost distribution’s dispersion. None of it requires issuance, discretion, or taxing distressed users — which is what a consensus layer built to exclude predetermined returns requires, and, we suspect, what Bitcoin’s fee-only future will eventually need as well.

References

- Ahmed, S. (2026a). An integrated simulation framework for decentralised finance protocols: cadCAD, reinforcement learning, game theory, and mechanism design. *SSRN Working Paper* 6323618.
- Ahmed, S. (2026b). The κ -Chain: an interest-free ($\iota = 0$) Layer-1 monetary substrate. Technical whitepaper, v0.7.6 (design specification).

- Ahmed, S., & Bhuyan, R. (2026). Interest-free monetary equilibrium: a general-equilibrium foundation for the κ -Standard. Working paper.
- Budish, E. (2018). The economic limits of Bitcoin and the blockchain. *NBER Working Paper* 24717.
- Buterin, V., Conner, E., Dudley, R., Slipper, M., Norden, I., & Bakhta, A. (2019). EIP-1559: Fee market change for ETH 1.0 chain. Ethereum Improvement Proposals.
- Carlsten, M., Kalodner, H., Weinberg, S. M., & Narayanan, A. (2016). On the instability of Bitcoin without the block reward. In *Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security* (pp. 154–167).
- Easley, D., O’Hara, M., & Basu, S. (2019). From mining to markets: The evolution of bitcoin transaction fees. *Journal of Financial Economics*, 134(1), 91–109.
- Huberman, G., Leshno, J. D., & Moallemi, C. (2021). Monopoly without a monopolist: An economic analysis of the Bitcoin payment system. *Review of Economic Studies*, 88(6), 3011–3040.
- Nakamoto, S. (2008). Bitcoin: A peer-to-peer electronic cash system. White paper.
- Roughgarden, T. (2021). Transaction fee mechanism design. In *Proceedings of the 22nd ACM Conference on Economics and Computation* (p. 792).
- Saleh, F. (2021). Blockchain without waste: Proof-of-stake. *Review of Financial Studies*, 34(3), 1156–1190.
- Schulman, J., Wolski, F., Dhariwal, P., Radford, A., & Klimov, O. (2017). Proximal policy optimization algorithms. *arXiv:1707.06347*.

EDITORIAL NOTE TO THIS PRINTING

Paper 13 · Warping the Glass

Working paper, revised July 2026 · SSRN 6925040

Role in the programme. Oracle integrity: when the adversary attacks the estimator instead of the contract. **What it establishes.** Manipulation signatures, a change-point monitor for the oracle, and the depth rule that caps how much value a manipulated print can reprice relative to bonded stake. **Corrected or extended later.** No substantive corrections to date; the chain implements these defences in consensus. **Not established here.** Detection guarantees against attack classes outside the studied signatures.

Replication for every number in this paper: DOI 10.5281/zenodo.21290557 · github.com/shehzadahmed-xx/kappa-replication

Warping the Glass: Measurement Integrity in Measurement-Based Monetary Systems

Shehzad Ahmed*

June 2026

Abstract

A measurement-based monetary system replaces chosen policy rates (committees) with measured ones (sensors): credit intensities read from markets, money valued by measured backing, triggers fired by observed stress. This removes the committee as a capture target and creates a new one — when the policy variable is a measurement, corruption migrates to the measuring layer. The attacker no longer debases the coin or lobbies the committee; it warps the glass. We analyse the full measurement stack of such a system layer by layer and derive the calibration each layer needs to make warping unprofitable. Three results organise the paper. *First*, for stake-weighted-median oracles we derive the manipulation frontier in closed form and show implementable slashing requires a **depth rule**: the value the oracle may reprice per detection window must be capped relative to bonded stake (at our calibration, roughly 7%) — a leverage limit on the measurement itself. *Second*, for the social layer (registrars attesting real-asset backing) the same structure reappears: honesty requires collateral $\geq$ attested value / audit probability, which makes single-registrar attestation unscalable and yields a quorum rule under which safe attested value grows approximately quadratically in the attestation quorum. *Third*, local no-profit calibrations **do not compose**: we exhibit a mixed attack — an in-band oracle bias inflating the value of fraudulently attested collateral — whose cross-term is priced by neither layer’s defense. Composition is restored by an end-to-end audit: a sequential (CUSUM) monitor of the divergence between system prices and an independent instrument bounds *total* distortion regardless of source. Calibrated on a real 14-year, 7-country panel of credit intensities against sovereign CDS (per-country correlations 0.989–0.998), the monitor detects a 25-basis-point warp at the first observation and caps the optimal slow attack at roughly 10 bp-years of repriceable notional. The residual trust assumption is identified exactly: at least one audit instrument must remain beyond the attacker’s reach.

Keywords: oracle manipulation, measurement integrity, mechanism design, sequential detection, blockchain oracles, benchmark manipulation, registries. **JEL:** D82, G14, K42, L86.

1 Introduction

Monetary systems have always had a layer an attacker most wants to corrupt. In commodity-money systems it was the coin itself: clipping and debasement. In administered systems it is the chooser: committee capture and benchmark rigging — the LIBOR scandal being the canonical modern case, in which the measurement was a survey and the survey was the attack

*Independent University, Bangladesh. This paper extends the mechanism-design analysis of [Ahmed \(2026a\)](#) to the measurement layer of the interest-free Layer-1 specification of [Ahmed \(2026b\)](#); the problem is general to any system whose policy variables are measurements rather than choices.

surface (Duffie & Stein, 2015). A new class of designs — zero-emission and interest-free Layer-1 architectures in which credit intensities are read from live markets, money is valued by its measured backing, and stabilisation mechanisms fire on observed stress (Ahmed, 2026b) — removes the committee entirely. Every policy variable becomes a *measurement*.

This is usually presented, correctly, as an integrity gain: a stake-weighted median computed from observed trades cannot be moved by a phone call. But the attack surface does not vanish; it migrates. If money is a mirror of real value, the profitable crime is no longer to counterfeit the metal but to *warp the glass*: bias the sensors, corrupt the attestations that bind tokens to real assets, or grind the reference constants — each individually small, none visible in any single epoch, all cumulatively a wealth transfer. Manipulation of thin crypto markets is documented (Gandal et al., 2018); what is missing is a design theory: **what does each layer of the measurement stack cost to corrupt, and how must collateral, audits, and monitors be calibrated so that no warp — fast, slow, or mixed — is profitable?**

Contributions. We model the measurement stack of a measurement-based monetary system in four layers — market oracles, asset registries, reference constants, and an end-to-end audit — and derive, for each, the no-profitable-warp calibration. *First* (Section 3), for stake-weighted-median oracles we give the achievable bias of a λ -stake coalition in closed form, $b(\lambda) = \sigma \Phi^{-1}(\frac{1/2}{1-\lambda})$ capped by the validity band, and derive the slashing calibration that makes manipulation unprofitable. Because slashing cannot exceed stake, implementability forces a **depth rule**: repriceable value per detection window capped relative to bonded stake (Proposition 2). *Second* (Section 4), we construct the end-to-end audit — a two-sided CUSUM (Page, 1954) on the divergence between the system’s credit intensities and an independent instrument — and calibrate it on a real panel: fourteen years of sovereign credit intensities implied by sukuk spreads against independently quoted sovereign CDS for seven countries (Ahmed & Bhuyan, 2026). The monitor detects a 25bp warp at the first print; the optimal *slow* attack is bounded at ≈ 10 bp-years. *Third* (Section 5), for the social layer we apply the classic deterrence logic (Becker, 1968; Shapiro & Stiglitz, 1984) to registrars: honesty requires bond-plus-franchise value $\geq$ attested value / audit probability, which makes single-registrar attestation of large assets impossible and yields a k -of- n quorum rule with safe value growing $\sim k^2$. *Fourth* (Section 6), we show local calibrations do not compose — an explicit mixed attack with a cross-term priced by neither layer — and prove that the end-to-end monitor restores composition because final-price divergence is additive in total distortion regardless of its source (Theorem 1). The residual assumption — the independence of at least one audit instrument — is stated exactly, because it is the system’s final trust anchor.

A unifying principle emerges that we did not impose: at every layer, the no-warp condition takes the same form,

$$\text{collateral at risk} \geq \frac{\text{value repriced by the measurement}}{\text{detection probability}},$$

a *leverage limit on measurement itself*. Systems that let measured values grow faster than the collateral and audit capacity behind their sensors are undercollateralised in a precise, calculable sense — whatever their balance sheets say.

2 The measurement stack and threat model

We consider a monetary system in which: (i) per-market credit/convergence intensities $\hat{k}$ are computed each epoch as stake-weighted medians of validator-submitted observations of market state, with submissions outside a validity band rejected and provably deviant

submissions slashable ex post; (ii) instruments must reference registered real assets, with registrars attesting the binding between on-chain tokens and off-chain assets; (iii) mechanism parameters (reference costs, trigger thresholds) are fixed governance constants recalibrated at low frequency; and (iv) an independent instrument (e.g. sovereign CDS quoted in a separate market) prices overlapping risks. The attacker is a coalition with a finite budget that may simultaneously: acquire validator stake, bribe registrars, and sustain sub-threshold biases over time. Its objective is extracted wealth net of capital costs, bribes, and expected penalties.

Four layers, four attacks: **L1** oracle capture (move the median); **L2** registry fraud (attest fake backing); **L3** the grind (sustained small bias below per-epoch detection); **L4** mixed attacks across layers. We treat them in turn.

3 Oracle capture and the depth rule

Honest submissions in an epoch are distributed F around the true value with dispersion σ (basis noise); the attacker controls stake share $\lambda < \frac{1}{2}$ and submits one-sided at the edge of the validity band $\pm\tau$.

Proposition 1 (Achievable bias). *The stake-weighted median under one-sided piling is the q -quantile of the honest distribution with $q = \frac{1/2}{1-\lambda}$; for Gaussian noise the achievable bias is*

$$b(\lambda) = \min\left(\sigma \Phi^{-1}\left(\frac{1/2}{1-\lambda}\right), \tau\right),$$

increasing and convex in λ , reaching the band cap as $\lambda \rightarrow \frac{1}{2}$.

This closed-form bias law is reproduced empirically in the companion simulation framework (Ahmed & Bhuyan, 2026b): a Monte-Carlo Byzantine oracle in which the adversary captures the largest stakers and submits a +500bp lie holds the stake-weighted median to 2–4bp of bias for all $\lambda \leq 1/3$ and lets it run away only as $\lambda \rightarrow 1/2$ — exactly the convex $b(\lambda)$ above and its 1/2 breakdown — while a stake-weighted *mean* is dragged by hundreds of basis points. The same experiment confirms the liveness companion of the depth rule: under a *withholding* attack the $\geq 2/3$ acceptance quorum converts any $\lambda > 1/3$ into a safe pause at the last good value, never an accepted bad $\hat{\kappa}$.

The attacker’s one-window profit is $b \cdot D \cdot H$, where D is the value repriced per unit of $\hat{\kappa}$ (open interest and collateral marked to the oracle) and H the epochs before ex-post audit. Its expected cost is $\lambda S \cdot s \cdot p_{\text{det}}(b)$: stake at risk times slash fraction times the probability a deviation of size b is flagged against realized data. The no-profit condition gives the required slash fraction $s^*(\lambda) = b(\lambda) D H / (\lambda S p_{\text{det}})$, plotted in Figure 1. Since $s^* \leq 1$ (one cannot slash more than the stake), implementability is a constraint on the *market*, not the penalty:

Proposition 2 (The depth rule). *The slashing mechanism can deter all coalitions $\lambda < \frac{1}{2}$ only if*

$$\frac{D}{S} \leq \min_{\lambda} \frac{\lambda p_{\text{det}}(b(\lambda))}{b(\lambda) H},$$

i.e. the value the oracle reprices per detection window is capped relative to bonded stake. At our calibration (σ equal to the empirical stationary dispersion of the intensity series, a $\pm 200\text{bp}$ band, daily audit) the cap binds at $\lambda^ \approx 0.42$ and equals $D/S \approx 0.07$.*

Remark 1. *This is a leverage limit on the measurement itself, and it is a new design constraint for oracle-centric architectures: open interest and oracle-marked collateral must scale with bonded stake, not merely with demand. Where $s^* > 1$ no slashing schedule deters capture; the remedies are smaller D/S , faster audits (smaller H), or better detection — which is what Section 4 supplies.*




Figure 1: The no-profitable-warp frontier: required slash fraction $s^*(\lambda)$ by market depth D/S . Above $s^* = 1$ (dashed) no implementable slashing deters the coalition — the regime the depth rule excludes.

4 The glass monitor: an end-to-end sequential audit

Per-epoch slashing sees only per-epoch deviations: a sustained bias *within* the flagging threshold — the grind — is invisible to it and cumulatively large. The defense is sequential: monitor the divergence between the system’s measured intensities and an *independent* instrument pricing overlapping risk, and alarm on persistent drift.

We construct the monitor on real data: the monthly credit intensities implied by sovereign sukuk spreads (the system leg) against independently quoted sovereign CDS (the audit leg) for seven countries over fourteen years (Ahmed & Bhuyan, 2026). Per-country correlations between the two legs are 0.989–0.998 — the two instruments price the same risk from different markets, which is precisely what makes the divergence informative. The divergence series $d_t = s_t^* - \text{CDS}_t$ is standardized on a training window and fed to a two-sided CUSUM (Page, 1954).

Calibrating the alarm threshold on the clean historical series and then injecting synthetic warps into the *real* data from a fixed date gives the detection-lag curve (Figure 2): a 25bp sustained bias is detected at the *first* subsequent observation in every country; a 10bp bias within one observation in 71% of countries. The audit leg in our panel is annual (the CDS source publishes yearly), so lags are reported in years; a production monitor on daily feeds shrinks them proportionally. Two of seven clean country-series trip the aggressive threshold — on inspection these are real dislocation episodes, i.e. the monitor also detects genuine market stress; in production those alarms route to the stabilisation mechanism rather than the slashing module.

Proposition 3 (The grind bound). *Against a monitor with detection-lag curve $\tau(b)$, the maximum extraction by any sustained bias is $\max_b b \cdot D \cdot \tau(b)$. On the real panel calibration this is ≈ 10 bp-years of repriceable notional, attained by the smallest detectable bias — the attacker’s optimal warp is the slow one, and the monitor’s threshold is what caps it.*

The monitor above is calibrated and stress-tested on the historical sukuk–CDS panel. Its companion validation runs the *same* construction inside a live agent-based protocol:




Figure 2: The glass monitor on real data. Left: a real divergence series with a +50bp warp injected (dashed). Right: median detection lag by injected bias (CUSUM on annual observations; daily feeds scale lags down proportionally).

a cadCAD model of the κ -substrate (Ahmed & Bhuyan, 2026b) embeds this two-sided CUSUM (reference slack $k = 0.5\sigma$, threshold calibrated to the 25bp warp) as the chain’s on-line measurement-integrity monitor, with a frequent-batch-auction matching layer, a stake-weighted-median index, and the depth rule $D/S \leq 0.07$ of Section 3. Against sustained warps of 10–100bp it reproduces the detection-lag curve of Figure 2 dynamically (100% detection in 3–8 blocks, faster for larger warps) and confirms the grind bound empirically: with the depth rule capping per-block repriceable notional, attacker extraction before detection stays below \$21k for every detectable warp, while the naive last-trade mark — lacking the monitor — lets the warp run the full horizon. The analytic bound here and that live simulation are thus two views of one result: the slow warp is the only one that pays, and the monitor’s threshold is what bounds it.

4.1 The cross-sectional layer: catching targeted warps from within

The independent-instrument leg catches *global* warps — biases applied to the whole published surface — at the audit instrument’s frequency. But the profitable targeted attack is different: bias *one* issuer’s measurements, to mis-mark one borrower’s collateral for liquidation or cheap acquisition. For this attack the system audits itself. From the full-universe history panel (273 bonds with daily G-spreads, aggregated to 37 issuer-month $\hat{\kappa}$ series over 2015–2026) we build a *cross-sectional* monitor: each issuer’s level residual from its fitted relationship to the panel median, standardized on a training window and fed to a per-issuer CUSUM. The panel median is immune to single-issuer manipulation, so it serves as an internal audit instrument that updates *monthly*. Injecting sustained targeted warps into the real series: a 50bp single-issuer warp is detected in a median of 5 months, 25bp in 9.5, and even 10bp in 11.5 — with 100% detection in every case (Figure 3). Six of thirty-seven issuers alarm on the *clean* history; inspection shows real idiosyncratic episodes (restructurings, sanctions), which a production system routes to review rather than slashing. The division of labour is now three-way: the depth rule and slashing govern fast attacks; the cross-sectional monitor catches targeted grinds from within, monthly; and the external independent instrument is needed *only* against globally uniform warps — which materially shrinks the residual trust assumption of Section 6.




Figure 3: The cross-sectional monitor on the real 37-issuer panel: median detection lag for sustained single-issuer warps; detection rate 100% at all tested sizes.

5 The social layer: registry incentive design

Asset-backing rules (the system’s existential-risk gate) bottom out in registrars: institutions attesting that a warehouse receipt, a sukuk certificate, a title corresponds to a real, deliverable asset. This layer is social — no cryptography verifies wheat — so its integrity must be purchased with the classic deterrence instruments (Becker, 1968; Shapiro & Stiglitz, 1984): bonds, random physical audits, and franchise value.

A registrar earning per-period franchise profit π_r (continuation value $C = \pi_r / (1 - \beta_r)$, an equity claim — this is an interest-free system, but time preference exists) posts bond W and faces per-period audit probability p_a . An attacker offers a bribe up to the fraudulent value V it unlocks. Honesty requires $p_a(W + C) \geq V$:

Proposition 4 (Bond schedule and concentration cap). *Single-registrar safety requires $W(V) \geq V/p_a - C$; equivalently, for any fixed bond, the attested value must satisfy $V \leq p_a(W + C)$. At realistic audit rates ($p_a \in [0.05, 0.10]$) the required bond is an order of magnitude or more above the attested value (the ratio is approximately $1/p_a$): single-registrar attestation does not scale. Under k -of- n quorum attestation (the attacker must bribe k registrars; any single audit kills the attestation), safety requires $k[1 - (1 - p_a)^k](W + C) \geq V$, so safe attested value grows approximately as k^2 for small p_a .*

The design consequences for the architecture’s whitelist: a per-registrar attested-value cap (the social layer’s depth rule), quorum attestation for large assets in place of large bonds, and audit intensity proportional to attested value. Note the structural identity with Proposition 2: collateral $\geq$ value / detection, again.

6 Composition: local defenses do not add up — the monitor does

Suppose both preceding layers are calibrated exactly to their no-profit bounds. Local safety does not imply global safety:

Remark 2 (A super-additive mixed attack). *Let the attacker run an in-band oracle bias (locally net-zero by Proposition 2’s calibration) and attest fake backing at the registry cap (locally net-zero by Proposition 4). The warped intensity inflates the market value of the*




Figure 4: Registry design. Left: required bond against attested value by audit probability. Right: safe attested value under k -of- n quorums — quorums, not bonds, are how large assets scale.

fraudulently attested collateral: the attacker borrows or sells against assets that are both fake and overpriced. The cross-term — relative price inflation times fake value — appears in neither layer’s deterrence condition. At an illustrative calibration (a within-band bias inflating prices 10%, fake backing at the cap), the mixed attack nets strictly positive profit while every local condition holds. Layer-by-layer security composes sub-additively in defense, super-additively in attack.

What restores composition is the end-to-end property of the monitor:

Assumption 1 (Independent instrument). *At least one audit instrument prices the overlapping risk in a market the attacker cannot move at relevant scale.*

Theorem 1 (End-to-end composition bound). *Under Assumption 1, the divergence between system prices and the independent instrument reflects the total distortion — the sum of oracle bias, registry-induced mispricing, and any cross-terms — regardless of which layers produced it. Consequently the monitor’s grind bound (Proposition 3) applies to the mixed attack: total undetected extraction is bounded by $\max_b b D \tau(b)$ plus the one-shot pre-detection caps of the local layers. Local calibrations bound per-layer profits; the end-to-end monitor bounds their composition.*

Proof sketch. Every layer’s distortion, to reach the attacker’s payoff, must move a system price: the oracle bias moves $\hat{k}$ directly; registry fraud moves the priced collateral pool; cross-terms move both. The independent leg, by Assumption 1, does not move. Hence the divergence statistic accumulates the sum of all price-relevant distortions, and the CUSUM bound on sustained divergence applies to that sum. One-shot components are bounded by the per-layer caps already derived. $\square$

Remark 3 (The residual trust anchor). *Assumption 1 is irreducible: a system all of whose reference instruments the attacker can move is unauditable from within. The practical defense is diversity — multiple audit legs (conventional bond bases, CDS from distinct venues, realized default records) so that moving all of them simultaneously exceeds any plausible budget — and the assumption’s scope is narrower than it first appears: the cross-sectional monitor of Section 4 polices targeted distortions from within, so external independence is required only against globally uniform warps. This is the precise, non-rhetorical sense in which a measurement-based monetary system still rests on something it must trust: not a committee, but the continued existence of at least one honest mirror.*

7 Limitations

The oracle model treats honest noise as Gaussian and one-sided piling as the optimal coalition strategy; richer submission strategies (timed, correlated with volatility) deserve adversarial search of the kind applied to the liveness mechanism in Ahmed (2026a). The monitor’s audit leg in our panel is annual; production calibration on daily feeds remains to be done, and alarm-routing (slashing vs stabilisation) on real dislocations needs design. Registrar parameters (franchise values, audit costs) are normalized, not estimated. The composition theorem bounds price-mediated extraction only: attacks whose payoff never routes through a system price (pure exfiltration of audit data, governance capture for non-pecuniary ends) are outside the model. The independence assumption is stated, not proven — it is the residual trust anchor, and we prefer to name it than to hide it.

8 Conclusion

When the policy variable is a measurement, the attack surface is the sensor. We have priced that surface layer by layer: a depth rule capping what an oracle may reprice relative to the stake behind it; a bond-and-quorum rule making registry fraud cost more than it unlocks; a sequential end-to-end monitor — calibrated here on fourteen years of real cross-instrument data — that detects fast warps at the first print and caps the slow ones; and a composition theorem showing why the monitor, not the local defenses, is what makes the layers add up. One principle runs through every layer: collateral at risk must exceed value-at-measurement over detection probability. Ancient monetary law made the integrity of the measure the sovereign’s first duty; in a measurement-based system that duty becomes a calibration — and, like all calibrations, it can be checked.

References

- Ahmed, S. (2026a). Liveness under macro winters: Mechanism design for fee-only blockchains. SSRN Working Paper 6924459.
- Ahmed, S. (2026b). The κ -Chain: A sovereign Islamic financial infrastructure layer. Technical whitepaper v0.7.6, SSRN 6924600.
- Ahmed, S., & Bhuyan, R. (2026). An Islamic credit default swap on smart-contract infrastructure: Design, pricing, and regulatory pathway. *SSRN Working Paper* 6323519.
- Ahmed, S., & Bhuyan, R. (2026b). A simulation framework for the κ -substrate: agent-based and mechanism-design validation of the on-chain convergence intensity. *SSRN Working Paper* 6323618.
- Becker, G. S. (1968). Crime and punishment: An economic approach. *Journal of Political Economy*, 76(2), 169–217.
- Duffie, D., & Stein, J. C. (2015). Reforming LIBOR and other financial market benchmarks. *Journal of Economic Perspectives*, 29(2), 191–212.
- Gandal, N., Hamrick, J. T., Moore, T., & Oberman, T. (2018). Price manipulation in the Bitcoin ecosystem. *Journal of Monetary Economics*, 95, 86–96.
- Page, E. S. (1954). Continuous inspection schemes. *Biometrika*, 41(1/2), 100–115.
- Shapiro, C., & Stiglitz, J. E. (1984). Equilibrium unemployment as a worker discipline device. *American Economic Review*, 74(3), 433–444.

EDITORIAL NOTE TO THIS PRINTING

Paper 14 · The Safe Asset without the Risk-Free Rate

Working paper, revised July 2026 · SSRN 6925058

Role in the programme. What replaces the riskless bond, and what its absence honestly costs. **What it establishes.** The κ -ladder construction, and the welfare cost of the missing riskless asset: on the order of 0.4 to 24 basis points of lifetime consumption in a two-agent general equilibrium, borne mostly by the more risk-averse agent. **Corrected or extended later.** A bibliography build defect (dangling citations) was fixed in July 2026. **Not established here.** The demand side of the zero-drift perpetual sukuk; the paper prices the gap, it does not yet design the buyer.

Replication for every number in this paper: DOI 10.5281/zenodo.21290557 · github.com/shehzadahmed-xx/kappa-replication

The Safe Asset without the Risk-Free Rate: Spanning, Heterogeneity, and the κ -Ladder

Shehzad Ahmed*

June 2026

Abstract

An interest-free monetary economy has, by construction, no instrument that pays the same in every state: the riskless intertemporal trade is gone. For a representative agent this costs nothing (Ahmed & Bhuyan, 2026a); the open question is heterogeneity — who loses the missing bond, how much, and whether anything native to the interest-free world replaces it. We answer in three steps. *First*, in a two-agent general equilibrium with heterogeneous risk aversion we compute the spanning gap directly: the consumption-equivalent value of adding a riskless bond to the equity-only economy is **0.4–24 basis points** of lifetime consumption across a wide heterogeneity grid — real, borne mostly by the more risk-averse agent, and small, consistent with the classical incomplete-markets literature. *Second*, we construct the **κ -ladder** — a diversified portfolio of event-intensity claims (sovereign credit, agricultural loss, mortality, project completion) whose risk floor is set by the *cross-domain* correlation ρ_x — and show it recovers **80–98.5%** of the spanning gap at residual floor volatilities of 0.5–1%. *Third*, we ask what ρ_x is: the full-universe sukuk panel (273 bonds with daily histories, 42 issuers) gives a *measured* within-credit correlation of 0.50 — proof that a single-domain ladder barely diversifies and that *domains, not names, are the unit of diversification* — while forty-four years of cereal-yield shocks give within-agriculture correlation 0.31 even between monsoon-coupled neighbours. With these measured blocks, a four-sleeve ladder (credit, agriculture, mortality, projects) has floor volatility 1.0–1.5% and beats the *real-terms* volatility of a nominal government bond with as little as one claim per domain. The deepest reframing is the benchmark itself: the “riskless” bond was only nominally safe — its real return always carried the inflation variance of the very printing press the interest-free architecture removes. Diversification across real event domains, not an issuer’s promise, is the interest-free world’s safe asset — and at measured calibrations it takes one claim per domain to build.

Keywords: safe assets, incomplete markets, spanning, risk-free rate, takaful, sukuk, diversification, interest-free finance. **JEL:** D52, E43, G11, G12.

1 Introduction

Every modern macro-finance model holds one instrument apart: the riskless bond, paying the same in all states, anchoring the intertemporal margin and letting every agent — the retiree, the pension fund, the precautionary saver — transfer wealth across time *with certainty*. An interest-free monetary economy deletes this instrument by construction: returns must track real outcomes, and no real outcome is certain. The companion general-equilibrium model (Ahmed & Bhuyan, 2026a) shows the deletion costs a *representative* agent nothing — time

*Independent University, Bangladesh. This paper answers the principal open question of the interest-free general-equilibrium model of Ahmed & Bhuyan (2026a): whether state-contingent equity returns span, for heterogeneous agents, what the riskless bond spanned. All models and data transformations are included in the replication package.

preference is priced through asset values, determinacy holds, optimal policy is characterised. The referee’s question, and this paper’s subject, is heterogeneity: with agents who differ in risk aversion, the missing riskless asset forecloses trades between them. Three questions follow. How large is the foreclosed surplus? Is there an interest-free instrument that recovers it? And was the benchmark — the “riskless” bond — ever actually riskless in the units that matter?

The questions sit in a classical literature. That equilibrium with an incomplete asset span exists but is generically constrained-suboptimal is canonical (Hart, 1975; Geanakoplos & Polemarchakis, 1986); that the *quantitative* cost of missing markets is often small is equally well documented (Lucas, 1987; Levine & Zame, 2002). The modern safe-asset literature (Caballero et al., 2017; Gorton, 2017) treats the supply of riskless instruments as a scarce public good produced by sovereign balance sheets. Our contribution is to put these together in the one economy where the question is existential rather than marginal — a system with *no* riskless contract at all — and to exhibit the replacement.

Contributions. *First*, we compute the spanning gap in a two-agent, two-period general equilibrium solved numerically (CRRA agents, lognormal aggregate risk, competitive equilibrium with and without the bond): across risk-aversion pairs (γ_A, γ_B) from (2, 5) to (2, 12) and aggregate volatilities of 2–4%, the consumption-equivalent value of the bond is 0.4–24 basis points, concentrated on the high- γ agent (Section 3). *Second*, we define the κ -ladder — N event-intensity claims with per-claim volatility σ_e and cross-domain correlation ρ_x , with portfolio variance $\sigma_e^2/N + (1 - 1/N)\rho_x\sigma_e^2$ and risk floor $\sqrt{\rho_x}\sigma_e$ — and show in the same equilibrium that a ladder with floor volatility 0.5% (1%) recovers 94–98.5% (80–94%) of the spanning gap (Section 4). *Third*, we benchmark honestly: the alternative to the ladder is not a zero-variance asset but a nominal bond whose *real* return carries inflation variance; the crossing size N^* at which the ladder’s real volatility beats the bond’s is single digits in emerging-market inflation environments and roughly a dozen at developed-market inflation volatility (Section 4). *Fourth*, we measure what the data allows: the full-universe sukuk history panel gives a within-credit correlation of 0.50 across 861 issuer pairs — so single-domain ladders barely diversify — while 44 years of cereal-yield shocks give within-agriculture correlation 0.31 even between monsoon-coupled neighbours; with these measured blocks a four-sleeve ladder reaches a 1.0–1.5% floor and the crossing at one claim per domain, and the adversarial single-pool test still passes (Section 5). *Fifth*, we draw the institutional design: the ladder is the interest-free world’s T-bill, constructible as a retail product (micro-sukuk ladders) and as the composition rule for sovereign stabilisation buffers (Section 6).

2 The missing trade

Two agents, two periods. Both have CRRA utility with risk aversions $\gamma_A < \gamma_B$, common discount factor $\beta = 0.97$, equal period-0 endowments, and equal initial shares of the aggregate period-1 endowment $Y(s)$, lognormal with growth volatility σ_{agg} . Assets: **(E)** the equity claim on Y only — the interest-free floor; **(E+B)** equity plus a riskless bond in zero net supply — the conventional benchmark; **(E+L)** equity plus a κ -ladder paying $1 + \varepsilon$, ε independent of Y with volatility σ_{floor} — the interest-free replacement. Competitive equilibria are solved numerically from first-order and market-clearing conditions; welfare comparisons are consumption-equivalent (CE) variations per agent, evaluated agent-by-agent at their own γ .

In (E) the risk-averse agent can de-risk only by selling equity, which also surrenders expected return; the bond in (E+B) unbundles the two, and agent B buys it from agent A, who levers. The spanning gap is the CE value of that unbundling.

3 The spanning gap is basis points

γ_A	γ_B	σ_{agg}	gap A (bp)	gap B (bp)	avg (bp)	rec. @0.5%	@1%	@2%
2	5	2%	0.36	0.90	0.63	94.1%	80.0%	50.0%
2	5	4%	1.45	3.63	2.54	98.5%	94.1%	80.0%
2	8	2%	0.71	2.84	1.77	94.1%	80.0%	50.0%
2	8	4%	2.83	11.42	7.13	98.5%	94.1%	80.0%
2	12	2%	1.00	6.05	3.52	94.1%	80.0%	50.0%
2	12	4%	3.96	24.43	14.19	98.5%	94.1%	80.0%

Table 1: The spanning gap (consumption-equivalent value of adding the riskless bond to the equity-only economy) and the fraction of it recovered by a κ -ladder with residual floor volatility 0.5%, 1%, 2%.

Three readings of Table 1. The gap is *real*: the missing bond forecloses genuine trades, and the cost falls where intuition says — on the risk-averse agent (up to 24bp of lifetime consumption at $\gamma_B = 12$ in the volatile calibration), six times the burden borne by the risk-tolerant one. The gap is *small*: basis points, not percent, echoing the classical finding that the measured cost of market incompleteness is modest (Lucas, 1987; Levine & Zame, 2002). And the gap is *recoverable*: an asset that is merely *nearly* riskless — floor volatility 0.5–1%, uncorrelated with the aggregate — delivers 80–98.5% of what the true bond delivers. The question therefore reduces to: does the interest-free world contain such an asset?

4 The κ -ladder

It does, and it requires no issuer’s promise. A κ -ladder is a portfolio of N event-intensity claims — takaful certificates on agricultural loss, mortality-pool participations, project-completion sukuk, sovereign-credit exposures — drawn from causally distinct domains. With per-claim volatility σ_e and average cross-domain correlation ρ_x ,

$$\sigma_p^2(N) = \frac{\sigma_e^2}{N} + \left(1 - \frac{1}{N}\right)\rho_x\sigma_e^2, \quad \sigma_p(\infty) = \sqrt{\rho_x}\sigma_e.$$

The honest benchmark is not a zero-variance asset. A nominal government bond’s *real* return carries the variance of inflation — historically 1.5–2.5% annually in developed markets and 3–6% in emerging ones — plus default risk that the interest-free world prices explicitly rather than assumes away. Define the crossing size N^* as the smallest ladder whose real volatility beats the bond’s. Figure 1 and the table below report it:

σ_e	ρ_x	$\sigma_\pi = 1.5\%$	2.5%	4%	6%
2%	0.05	2	1	1	1
4%	0.05	11	3	2	1
4%	0.10	23	4	2	1
6%	0.05	77	8	3	2

Table 2: Crossing size N^* : smallest κ -ladder whose real volatility beats a nominal government bond’s (inflation volatility σ_π). At the central calibration (third row), three claims beat the bond even at developed-market $\sigma_\pi = 1.5\%$, one claim already beats it at $\sigma_\pi = 6\%$.

At the central calibration ($\sigma_e = 4\%$, $\rho_x \leq 0.05$), **three claims** beat the bond’s real volatility in any emerging-market inflation environment and **eleven** beat it even at developed-market




Figure 1: Left: κ -ladder real volatility against ladder size for three cross-domain correlations ($\sigma_e = 4\%$), with nominal-bond real volatility benchmarks. Right: crossing size N^* over the (ρ_x, σ_e) plane against a 2.5% inflation-volatility bond.

inflation volatility. The interest-free world’s safe asset is cheap to construct — a dozen diversified claims on reality — and its safety is a property of diversification across real events, not of anyone’s promise to pay.

Remark 1 (Why the bond was never the benchmark it claimed to be). *The riskless bond’s certainty is nominal. Its real payoff inherits the inflation process — which, in the conventional architecture, is managed by exactly the discretionary instruments (elastic issuance, the inflation tax) that the interest-free design removes. The comparison in the table is therefore not “risky ladder vs. safe bond” but “one bundle of real risks vs. another” — and on the data, the diversified bundle of event risks is the smaller one in most of the world.*

5 What is ρ_x ? An honest accounting

The ladder’s floor is the cross-domain correlation, so everything turns on its magnitude. A full empirical matrix (credit $\times$ weather $\times$ mortality $\times$ projects) needs long parallel series that do not yet exist; we estimate what the data supports and bound the rest.

Within credit, measured from the full universe. The complete Cbonds history panel — 273 USD sukuk with daily G-spread histories, 2015–2026, aggregated to issuer-month $\hat{\kappa}$ series (42 issuers with three or more years of observations) — gives the pairwise correlation of monthly $\Delta\hat{\kappa}$ across **861 issuer pairs**: mean 0.50 (sovereign–sovereign 0.43, corporate–corporate 0.53, same-country 0.60 vs. cross-country 0.49). This is the paper’s most important empirical fact: even at $\iota = 0$ pricing, claims *within* the credit domain share a strong common factor — the global credit cycle — so a credit-only ladder has floor $\sqrt{0.50}\sigma_e \approx 0.71\sigma_e$ and barely diversifies. Names are not diversification; *domains* are.

Within agriculture, across countries. Forty-four years (1980–2023) of cereal-yield shocks — deviations from trend, the variable underlying the takaful pricing of [Ahmed & Bhuyan \(2026b\)](#) — across Bangladesh, India, Indonesia, and Pakistan give pairwise correlations with median 0.31 (range -0.05 to 0.48). These four economies share monsoon geography; the number is an *adversarial upper proxy* for any geographically diversified sleeve — Indonesian and Pakistani harvests, two thousand miles apart, already correlate at -0.05 .

The multi-sleeve ladder with measured blocks. With M equal sleeves, within-sleeve correlation ρ_w and cross-sleeve correlation ρ_x , the large- N floor is $\sigma_e\sqrt{(\bar{\rho}_w - \rho_x)/M + \rho_x}$: domain diversification divides the within-domain comovement by the number of sleeves. Using

the *measured* blocks — credit 0.50, agriculture 0.31, mortality and projects bounded at 0.10 ($\bar{\rho}_w = 0.25$) — and $\sigma_e = 4\%$, the four-sleeve floor is 1.01% at $\rho_x = 0$, 1.27% at $\rho_x = 0.05$, and 1.49% even at $\rho_x = 0.10$ — in every case below both inflation benchmarks, with the crossing achieved by the *minimal* ladder of one claim per domain ($N^* = 4$ at $\sigma_e \leq 4\%$). And the adversarial single-pool test still passes: if every pair correlated like monsoon-neighbour harvests ($\rho = 0.31$), the ladder beats the emerging-market benchmark at $N^* = 2$ and the developed-market one at $N^* = 9$; at $\sigma_e = 6\%$ only the emerging-market comparison survives — reported, not smoothed.

Across domains. Even the full-universe panel yields only seven pooled country-year overlaps between sovereign credit changes and same-country yield shocks — too few for inference, so we bound rather than estimate ρ_x and report the sensitivity above. Mortality-credit and project-credit pairs cannot be estimated with present data; structurally, however, event intensities share no common *pricing* factor at $\iota = 0$: in the conventional world, the single discount-rate factor moves every asset simultaneously — the mechanism by which “diversified” portfolios converged to correlation one in 2008 — whereas drought frequencies, mortality tables, and completion times comove only through real disaster channels. Removing the interest rate does not merely remove *riba*; it removes the factor that made financial diversification fail when it was needed most. This is a structural argument, not yet a measurement, and we flag it as the program’s standing empirical task as cross-domain claim histories accumulate on-chain.

A first cross-domain tail measurement. The 2020–21 pandemic supplies the first natural experiment in which several domains were shocked at once, and it sharpens the structural claim rather than overturning it. Aggregating an independent panel of sovereign USD T-spread indices, the median sovereign credit intensity rose roughly +167% from 2019 to 2020 (263 → 701 bp); the World Bank cross-country crude death rate rose +14% (2019 → 2021); and EM-DAT total-affected was essentially flat (−5%). Two lessons follow. First, *tail* co-movement can exceed the small *average* cross-domain correlation: credit and mortality moved together through the real pandemic channel, so a credit-plus-mortality sleeve diversifies less in a systemic crisis than the average ρ_x suggests — the ladder’s protection is weakest precisely in a global tail. Second, the coupling is *channel-specific*, not a return of the single pricing factor: the disaster domain did *not* co-move with the pandemic, so agriculture- and disaster-linked sleeves remained genuinely diversifying even in 2020. The honest reading is therefore intermediate between the structural ideal ($\rho_x \rightarrow 0$) and the conventional failure mode ($\rho \rightarrow 1$): removing ι does remove the common pricing factor, but real systemic events still couple the subset of domains they physically touch. A safe-asset ladder built for tail robustness should accordingly weight domains with disjoint real-event channels (e.g. weather and project completion) over pairs — such as sovereign credit and mortality — that a single macro disaster can move together.

6 Institutional design: the interest-free T-bill

The ladder is not only an answer to a referee; it is a product specification. **Retail:** the micro-sukuk savings instrument of the companion architecture (Ahmed, 2026c) becomes a *ladder* — small participations across takaful pools, project sukuk, and sovereign credit from distinct regions — giving low-income savers the near-riskless store of value the conventional system reserves for T-bill holders, without a guaranteed-return structure that would reconstitute *riba*. **Sovereign:** the stabilisation buffers of the interest-free architecture (the Sovereign Stabilisation Fund; the validator liveness sleeve of Ahmed, 2026a) face wrong-way risk if invested in the asset class they defend; the ladder is their composition rule, with the counter-cyclical commodity sleeve (vault-secured gold) as the disaster-channel hedge. **Systemic:**

safe-asset supply in the conventional world is a scarce public good produced by sovereign balance sheets (Caballero et al., 2017; Gorton, 2017); ladder construction is permissionless and scales with the diversity of real economic activity rather than with any government’s debt capacity. An economy that wants more safe assets builds more independent things.

7 Limitations

The welfare model is deliberately minimal: two agents, two periods, an endowment economy; richer life-cycle settings would raise the level of the spanning gap (long-horizon savers value certainty more) while, by the recovery results, leaving the ladder’s relative performance intact — a conjecture to verify, not a theorem. Ladder parameters σ_e and ρ_x are calibrated and bounded, not jointly estimated; the four-point cross-domain check is reported as an anecdote. The supply side of the ladder is not modelled: someone must originate and hold the other side of event claims, and origination capacity, liquidity, and pricing in thin early markets are open design questions (the companion measurement-integrity analysis governs the oracles those claims settle on, Ahmed, 2026b). Mortality and project-completion correlations are bounded by argument, not data. The ladder also mixes measures: spread-implied claim volatilities are Q -measure (risk-premium inclusive) while yield-shock volatilities are P -measure; since Q -volatilities weakly exceed their physical counterparts, the crossing sizes reported are conservative, but a measure-consistent re-estimation belongs on the same future-data agenda. Finally, a juristic note: a referee may observe that a large ladder’s near-deterministic return is “functionally fixed.” The classical prohibition, however, operates at the *contract* level — a predetermined obligation on a specific counterparty — not at the *portfolio* level, where stability emerges statistically from many genuinely loss-bearing positions: a mudarabah fund over many ventures and a takaful operator’s book have always had stable aggregate returns without either being *riba*. Each rung of the ladder remains individually loss-bearing; the adjudication of this distinction, like all such, rests with qualified scholars. Finally, the inflation-volatility benchmark treats the nominal bond’s real risk as exogenous; a conventional system with an independent, disciplined central bank can compress it — the comparison is sharpest precisely where monetary discipline is weakest, which is also where the populations this architecture targets live.

8 Conclusion

The interest-free economy’s missing riskless bond forecloses real trades worth basis points of lifetime consumption, borne mainly by its most risk-averse members. A diversified ladder of event-intensity claims recovers 80–98.5% of that loss, reaches near-riskless real volatility with roughly a dozen claims at plausible correlations, survives an adversarial test using the worst measured within-region correlation, and — unlike the instrument it replaces — requires no promise from any issuer and no printing press behind the promise. The riskless rate was never a fact of nature; it was a convention backed by discretion. The safe asset without it is diversification across reality — which is, perhaps, what safety always was.

References

- Ahmed, S. (2026a). Liveness under macro winters: Mechanism design for fee-only blockchains. SSRN Working Paper 6924459.
- Ahmed, S. (2026b). Warping the glass: Measurement integrity in measurement-based monetary systems. Working paper.

- Ahmed, S. (2026c). The κ -Chain: A sovereign Islamic financial infrastructure layer. Technical whitepaper v0.7.6, SSRN 6924600.
- Ahmed, S., & Bhuyan, R. (2026a). Interest-free monetary equilibrium: A general-equilibrium foundation for the κ -Standard. SSRN Working Paper 6890259.
- Ahmed, S., & Bhuyan, R. (2026b). Actuarial properties of the κ -rate under stochastic hazard: CIR applications to agricultural takaful in South and Southeast Asia. *SSRN Working Paper* 6323459.
- Caballero, R. J., Farhi, E., & Gourinchas, P.-O. (2017). The safe assets shortage conundrum. *Journal of Economic Perspectives*, 31(3), 29–46.
- Geanakoplos, J., & Polemarchakis, H. (1986). Existence, regularity and constrained suboptimality of competitive allocations when the asset market is incomplete. In *Uncertainty, Information and Communication: Essays in Honor of Kenneth J. Arrow*, Vol. III, 65–96. Cambridge University Press.
- Gorton, G. (2017). The history and economics of safe assets. *Annual Review of Economics*, 9, 547–586.
- Hart, O. D. (1975). On the optimality of equilibrium when the market structure is incomplete. *Journal of Economic Theory*, 11(3), 418–443.
- Levine, D. K., & Zame, W. R. (2002). Does market incompleteness matter? *Econometrica*, 70(5), 1805–1839.
- Lucas, R. E., Jr. (1987). *Models of Business Cycles*. Oxford: Basil Blackwell.

EDITORIAL NOTE TO THIS PRINTING

Paper 15 · Separating Event Intensity from Fixed Funding Carry

Working paper, revised July 2026 SSRN submission 7143718

Role in the programme. The focused empirical validation. It asks whether the event-intensity object can be measured independently of the fixed funding-carry component, and whether the resulting ordering survives across perpetual-futures and credit-market data. **What it establishes.**

In the primary tests, perpetual bases mean-revert across forty-one markets with a median half-life of about 3.5 hours; daily intensity estimates co-move across venues; credit-implied intensity is ordered by rating and tenor; and the fixed positive funding input appears as a formula-induced concentration of observations near its stipulated value on centralized venues, while the premium-only venue lacks the same mass point. Harmonized venue means are treated as descriptive rather than causal evidence. A benchmark-free bond estimator preserves the broad relative ordering of the spread-based measure. **Corrected or extended later.** This substantially rewritten version replaces the programme-wide report previously submitted as SSRN 6969963 and rejected; resubmitted as working paper 7143718. Cross-domain, protocol-security and monetary analyses are moved to the online appendix; the venue comparison is stated as diagnostic rather than causal; the sovereign-default result is treated as descriptive forward event ordering; and estimator-, aggregation- and recovery-dependence are reported explicitly. **Not established here.** External replication by unaffiliated investigators, causal identification of venue design, realised performance of a traded $\iota = 0$ instrument through a full market cycle, or an overall Shariah ruling. The paper isolates an empirical pricing component; jurisprudential permissibility remains a separate determination.

Replication for every number in this paper: DOI 10.5281/zenodo.21290557 · github.com/shehzadahmed-xx/kappa-replication

Separating Event Intensity from Fixed Funding Carry: Evidence from Perpetual Futures and Credit Markets

Shehzad Ahmed*

Rafiqul Bhuyan†

July 2026

Abstract

Perpetual futures combine a state-dependent anchoring premium with, on some venues, an exchange-specified fixed funding component. We ask whether the anchoring mechanism can be measured separately from fixed carry and whether its intensity has a coherent credit-market analogue. Using 804,000 hourly funding observations, 921 bond histories, 51 sovereign CDS series, and cross-domain loss data from seventeen providers, we obtain four results. Perpetual bases mean-revert in all 41 pre-specified dYdX markets, with a median half-life of 3.5 hours, and daily intensity estimates correlate across venues at 0.70. After harmonizing all venues to complete eight-hour intervals, signed mean funding is 0.83 basis points on dYdX, 1.02 on Binance, and 1.14 on Bybit across nine common markets. Mean drift therefore does not by itself identify the fixed input. The sharper diagnostic is distributional: centralized-exchange observations exhibit a pronounced mass around the one-basis-point formula point, whereas the premium-only venue does not. Spread-implied intensities are broadly ordered by rating and monotone in tenor and place the three subsequent sovereign default or restructuring events among the highest early-risk observations. Finally, a benchmark-free bond estimator preserves broad risk ordering, with a Spearman rank correlation of 0.82. The study is an author-run out-of-sample validation using previously unused external data, not external replication.

Keywords: perpetual futures; funding carry; convergence intensity; credit risk; hazard models; benchmark-free pricing.

JEL classification: C58, G12, G13, G15.

1 Introduction

Perpetual futures have become a central instrument in digital-asset markets. Their defining design problem is anchoring: because the contract has no maturity, periodic transfers are used to keep its price close to an index. Recent theory formalizes that mechanism through a random event time whose intensity governs the speed of convergence [Angeris et al., 2023, Akerer et al., 2025]. At the same time, empirical work shows that crypto carry is economically important and reflects demand pressure, market segmentation, and limits to arbitrage [Schmelting et al., 2026]. These two objects should not be conflated. The convergence intensity describes the speed of anchoring; total carry is an equilibrium outcome; and some exchange formulas additionally insert a fixed funding component.

This paper asks two connected questions. First, can the anchoring intensity be recovered consistently across markets and separated from the exchange-specified fixed component? Second,

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

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

does the same event-intensity grammar organize credit-market data, where reduced-form models have long represented default through a stopping time and a risk-neutral intensity [Duffie and Singleton, 1999, Duffee, 1999, Driessen, 2005]? The connection is structural rather than numerical: convergence, default, mortality, and insured loss are different events with different intensities, but the corresponding random times occupy the same mathematical role.

We assemble a previously unused dataset from seventeen providers. The core sample contains 804,000 hourly dYdX observations, funding histories from Binance and Bybit, 921 sovereign and corporate bond histories, 51 sovereign CDS series, local-currency curves, and physical loss records. The study reports 81 analyses across the main paper and appendices (Appendices A–G). The analysis hierarchy was recorded in a dated internal plan that was not externally preregistered, and exploratory analyses are labelled separately. The study is conducted by the authors of the underlying model sequence, so independence refers to the data and statistics, not the investigators.

Four findings organize the paper. First, perpetual bases mean-revert across all 41 pre-specified dYdX markets, with a median half-life of 3.5 hours; daily intensity estimates co-move across venues at 0.70. Second, harmonization to a common eight-hour unit yields signed means of 0.83 basis points on dYdX and 1.02–1.14 basis points on the two fixed-input venues. The small difference in means is not an identification result because venues differ in clientele and market structure. Identification of the documented formula input instead comes from distributional pinning around one basis point on the centralized exchanges and the absence of an analogous mass point on dYdX. Third, spread-implied credit intensities are broadly ordered by rating and monotone in tenor, contain forward event information, and form a coherent surface. Fourth, deleting the government curve from the estimator preserves broad relative risk ordering, although it changes levels and does not imply that traded assets are insulated from equilibrium conditions in an interest-bearing economy.

The contribution is not a claim that all funding is interest, all credit spreads are expected default loss, or one numerical κ governs unrelated markets. Instead, the paper provides an empirical bridge between three literatures that have largely developed separately: perpetual-contract anchoring, reduced-form credit intensity, and benchmark dependence in Islamic finance. It identifies which predictions are mechanical, which are supported in the data, which depend on recovery or estimator construction, and which remain outside empirical adjudication.

2 Related literature and contribution

2.1 Perpetual contracts, anchoring, and crypto carry

The idea of a perpetual futures contract predates crypto markets [Shiller, 1993], but digital-asset venues made the instrument economically important. Angeris et al. [2023] derive model-free funding and replication relations for perpetual contracts. He et al. [2025] derive no-arbitrage prices and trading-cost bounds and document common, mean-reverting deviations across crypto contracts. Akerer et al. [2025] provide explicit no-arbitrage prices and show that the futures price can be represented as the risk-neutral expectation of spot sampled at a random time governed by the intensity of anchoring. Kim and Park [2025] study funding-rule design as a control problem and construct path-dependent alternatives. Together, this literature makes funding design an object of financial modelling rather than an exchange convention. The random-time representation supplies the principal model restriction tested here.

The empirical derivatives literature emphasizes price discovery, liquidity, and market fragmentation. Perpetual contracts often lead spot markets in Bitcoin and Ether, but their informational role varies with trading volume, spreads, and venue structure [Alexander et al., 2020a, Alexander and Heck, 2020, Alexander et al., 2020b]. More recently, Schmeling et al. [2026] show that total

crypto carry can be large and volatile because leveraged demand interacts with limits to arbitrage. Our paper is complementary: it does not interpret total carry as the fixed formula component. We distinguish (i) convergence speed, (ii) the state-dependent premium, and (iii) an exchange-specified fixed input. This distinction is necessary for both asset-pricing interpretation and the Islamic-finance application.

2.2 Reduced-form credit intensity and spread decomposition

Reduced-form credit models price defaultable claims through a stopping time with an intensity process [Duffie and Singleton, 1999, Duffee, 1999]. The empirical literature establishes both the usefulness and the limits of spread-implied intensity. Credit spreads contain expected default loss and compensation for default risk, but also liquidity, taxes, and other nondefault components [Elton et al., 2001, Collin-Dufresne et al., 2001, Driessen, 2005, Longstaff et al., 2005, Huang and Huang, 2012]. Sovereign CDS further embeds time-varying default and recovery premia and substantial global risk components [Pan and Singleton, 2008, Longstaff et al., 2011]. Accordingly, our credit measure $s/(1 - \delta)$ is labelled a *pricing intensity* under the risk-neutral measure, not a direct estimate of physical default frequency. We separately estimate physical event intensity when realised loss data are available.

Our contribution to this literature is a cross-market mapping. The same random-time representation that anchors a perpetual has a reduced-form credit analogue at the zero-fixed-component point. We then ask whether the mapped object exhibits the empirical restrictions expected of an intensity: non-negativity, mean reversion, rating ordering, tenor ordering, coherent forward intensities, and association with subsequent events.

2.3 Benchmark dependence, sukuk, and Islamic finance

Islamic finance has achieved scale while retaining substantial operational proximity to conventional banking. The benchmark problem has been recognized for decades: in the absence of a distinct pricing primitive, institutions have relied on conventional interbank curves even when contractual form differs [Iqbal, 1999]. Cross-country evidence finds fewer differences in business model than the theoretical profit-and-loss-sharing ideal would suggest [Beck et al., 2013, Abedifar et al., 2013, 2015]. Most directly, Azad et al. [2018] document long- and short-run convergence between the Islamic Interbank Benchmark Rate and LIBOR and argue that decoupling is unlikely without products with genuinely distinct risk-return structures. Sukuk are commonly quoted relative to conventional curves, and Uddin et al. [2022] argue that their pricing determinants should be modelled separately rather than imposed through an ad hoc conventional benchmark.

The present paper addresses a narrower, testable question: can relative credit ordering be recovered when the government curve is removed as an operational input? A positive answer would not prove complete economic independence from conventional markets, nor would it determine the Shariah status of any instrument. It would show only that the benchmark is not necessary for the relative cross-sectional ordering produced by the proposed estimator. The legal characterization of fixed funding, risk premia, possession, leverage, and settlement remains outside the authority of the empirical tests.

2.4 Insurance intensity and premium loadings

Insurance provides a useful physical-measure comparison because event frequencies can be observed directly. Classical premium theory distinguishes expected claims from expense, capital, and risk loadings [Kahane, 1979, Buhlmann, 1980, Wang, 2002]. We therefore do not treat expected loss as a

universal market premium. Instead, insurance data identify a physical event intensity and allow us to measure the gap between realised loss frequency and priced intensity. That gap is small in some cooperative-protection applications and much larger in traded credit.

2.5 Validation, multiplicity, and reproducibility

Large empirical programmes face data-snooping and multiplicity concerns [White, 2000, Sullivan et al., 1999, Harvey et al., 2016]. We address these concerns by separating pre-specified primary tests, secondary robustness analyses, and post-plan exploration; by reporting null and construction-specific results; and by focusing the principal claims on repeated directional restrictions rather than counts of nominal significance. The reproduce-first protocol also responds to the documented difficulty of computational reproducibility in finance [Perignon et al., 2024]. It is not a substitute for external replication, but it makes the distinction between a failed prediction and a different estimator explicit.

2.6 Summary of contributions

The paper makes four contributions. First, it provides the first large cross-venue measurement of the anchoring intensity implied by the random-time perpetual-pricing model. Second, it separates a fixed formula input from state-dependent premium and equilibrium carry, avoiding an identification error that would otherwise conflate them. Third, it maps the event-time object into credit data and documents rating, tenor, and forward-event restrictions on an expanded international sample. Fourth, it evaluates benchmark deletion as an operational, rather than rhetorical, proposition and records where the resulting estimates remain sensitive to recovery, liquidity, and common base-yield conditions.

Literature	Established result	Increment in this paper
Perpetual pricing	Funding anchors a non-expiring contract; random-time representations identify an anchoring intensity.	First broad empirical measurement of the random-time intensity and a direct test of formula-induced pinning.
Crypto carry	Total carry reflects leveraged demand and limits to arbitrage.	Separates total equilibrium carry from an exchange-specified fixed input and a convergence-speed object.
Reduced-form credit	Credit spreads map into risk-neutral default intensity but contain nondefault components.	Tests whether the event-time mapping produces coherent rating, tenor, and forward-event restrictions across an international sample.
Islamic benchmark research	Islamic rates and products remain closely linked to conventional curves.	Deletes the government curve at the estimator level and measures what relative ordering survives and what common level remains.

Table 1: Position relative to the closest literatures. The contribution is a model-validation and identification exercise, not a claim that the three economic objects are identical.

3 Model, measurement, and testable restrictions

3.1 Anchoring intensity in perpetual futures

In the random-time representation of [Ackerer et al. \[2025\]](#), a perpetual price can be written schematically as

$$f_t = \mathbb{E}_t^\mathbb{Q}[X_\theta], \quad \theta - t \sim \text{Exp}(\kappa_t), \quad (1)$$

where $\kappa_t > 0$ is the intensity with which the anchoring mechanism samples the underlying. The empirical counterpart is the speed at which the basis reverts. For an AR(1) representation of the sampled basis, $b_{t+\Delta} = \phi b_t + u_{t+\Delta}$, we estimate

$$\hat{\kappa}^{\text{conv}} = -\frac{\log \hat{\phi}}{\Delta}. \quad (2)$$

This is a convergence-speed estimate. It is distinct from the magnitude of the funding premium.

Lemma 1 (Sampled convergence intensity). *If the basis follows $db_t = -\kappa b_t dt + \sigma dW_t$ with constant $\kappa > 0$, then sampling at interval Δ yields an AR(1) coefficient $\phi = e^{-\kappa\Delta}$. Hence Equation (2) is the exact discrete-time mapping under the Ornstein–Uhlenbeck benchmark, and the corresponding half-life is $\log(2)/\kappa$.*

Some centralized venues use a formula of the form

$$F_t = P_t + \text{clamp}(\iota - P_t, -c, c), \quad (3)$$

where P_t is the state-dependent premium, ι is a fixed exchange input, and c is the clamp width [[Binance, 2026](#), [Bybit, 2026](#)]. The premium-only comparison venue has no corresponding fixed baseline in its funding rule [[dYdX, 2026](#)]. Equation (3) supplies a mechanical decomposition; differences in realised funding distributions across venues remain descriptive because market structure is not held fixed.

Proposition 1 (Formula-induced pinning). *Let P have a continuous distribution and define $F_\iota(P) = P + \text{clamp}(\iota - P, -c, c)$ for $c > 0$. Then*

$$F_\iota(P) = \begin{cases} P + c, & P < \iota - c, \\ \iota, & \iota - c \leq P \leq \iota + c, \\ P - c, & P > \iota + c. \end{cases} \quad (4)$$

Consequently, the funding distribution has a formula-induced atom at ι of mass $q_\iota = \Pr(\iota - c \leq P \leq \iota + c)$. Its mean is

$$\mathbb{E}[F_\iota] = \iota q_\iota + \mathbb{E}[(P + c)\mathbf{1}\{P < \iota - c\}] + \mathbb{E}[(P - c)\mathbf{1}\{P > \iota + c\}]. \quad (5)$$

A premium-only rule $F^P = P$ has no formula-induced atom. When most premium observations lie inside the clamp interval, realised funding is therefore pinned near the fixed input, while tail observations retain state dependence.

Proposition 1 supplies a direct diagnostic through mass at the fixed input and shows how the formula may affect the mean. It does not identify total equilibrium carry, because the distribution of P itself depends on demand, arbitrage capital, and venue structure. The realized cross-venue mean difference is not itself an identifying statistic because the premium distribution and venue structure are not held fixed.




Figure 1: Funding-rule decomposition. The anchoring premium is state dependent; selected centralized venues add a fixed formula input. Total realised carry also reflects venue-specific equilibrium forces and should not be identified with the fixed input alone.

3.2 Credit-market mapping

In a reduced-form model, a default event arrives at a risk-neutral intensity $\lambda^{\mathbb{Q}}$ and loss given default is $1 - \delta$. Under the standard approximation,

$$s \simeq (1 - \delta)\lambda^{\mathbb{Q}}, \quad \hat{\kappa}^{\text{credit}, \mathbb{Q}} = \frac{s}{1 - \delta}. \quad (6)$$

The mapping $\kappa \leftrightarrow \lambda$ is an isomorphism between event-time roles, not a claim that perpetual convergence and default are the same economic event. Moreover, s can contain liquidity, taxes, and risk premia; therefore $\hat{\kappa}^{\text{credit}, \mathbb{Q}}$ is a spread-implied pricing intensity. Physical intensity $\kappa^{\mathbb{P}}$ is estimated separately from realised events.

For a bond with yield y , maturity T , and recovery δ , the benchmark-free estimator used in the tests is

$$\hat{\kappa}^{\text{price}} = -\frac{1}{T} \log \left(\frac{e^{-yT} - \delta}{1 - \delta} \right), \quad (7)$$

whenever the argument of the logarithm is positive. The empirical question is whether Equation (7) preserves relative ordering compared with a government-spread estimator. It is not a claim that a bond traded in a mixed economy is insulated from the government curve or monetary conditions.

3.3 A family of domain-specific intensities

We use κ as a family name for event intensities sharing the same mathematical role. The principal objects are κ^{conv} for basis convergence, $\kappa^{\text{credit}, \mathbb{Q}}$ for spread-implied credit pricing, and $\kappa^{\text{loss}, \mathbb{P}}$ for realised loss frequency. Their levels need not coincide and in fact span orders of magnitude. The common claim is structural: each parameter indexes the arrival or resolution speed of a defined event.




Figure 2: A common event-time grammar across domains. Structural equivalence concerns the role of an event intensity in a random-time representation, not equality of the events or numerical intensities.

3.4 Primary restrictions

The pre-specified primary family evaluates four restrictions:

- H1:** The perpetual basis mean-reverts and estimated convergence intensity is coherent across venues.
- H2:** A venue formula containing a fixed input exhibits the formula-induced pinning atom and associated distributional distortion implied by Proposition 1; signed venue means are treated as descriptive and do not by themselves identify the fixed input.
- H3:** Spread-implied credit intensity is broadly ordered by rating and monotone in tenor, remains predominantly non-negative, and contains forward event information.
- H4:** Removing the government curve preserves broad relative risk ordering, while level differences reveal the common base-yield layer and residual nondefault components.

Selected cross-domain evidence is summarized in Section 9; full cross-domain results are reported in Appendix D, and protocol simulations and monetary applications in Appendix F.

4 Data

Table 2 summarizes the dataset. The observation unit, provider-specific sample window, inclusion rules, missing-data treatment, and acquisition script for every series are recorded in the accompanying data map. Raw files obtained under user-specific or commercial licences are not redistributed.

Domain	Source and unit	Coverage used in the analysis
Perpetual funding	dYdX, Binance, Bybit; funding interval	43 dYdX markets (804,000 hourly intervals), 40 Binance markets and 30 Bybit markets at native cadence; 41 dYdX markets constitute the pre-specified convergence sample.
USD credit	Cbonds; bond-day and sovereign-month	921 sukuk and conventional bond histories, 121 issuers, 51 sovereign CDS series, and 40 sovereign spread/yield indices.
Local-currency credit	Cbonds; bond and curve observations	215 OIC sovereign bonds plus government and corporate curves for Malaysia, India, Brazil, Mexico, China, Indonesia, Turkey, and Saudi Arabia where available.
Rates and macro controls	FRED, ECB, BIS, OECD, IMF	US Treasury and euro-area curves, policy rates, inflation, debt-service ratios, credit/GDP, REER, and sovereign fundamentals.
Physical event domains	World Bank, WHO, HMD, EM-DAT, FAOSTAT, NASA POWER	Mortality, agriculture, disasters, and climate at 150–260 countries; 16,853 disaster records and 41 national life-table systems.
Tokenized assets and market structure	CoinGecko, DeFiLlama, exchange indexers	Tokenized gold and Treasury products, live BTC/ETH order-book depth, and dYdX TVL history.

Table 2: Previously unused external data assembled for the author-run validation. Detailed schemas and acquisition rules are available from the corresponding author upon reasonable request, subject to licensing restrictions.

5 Empirical design

5.1 Out-of-sample and reproduce-first protocol

Every primary series is outside the estimation sample used in the earlier working papers. Before classifying a difference as a failed prediction, we first reproduce the source number with its original data and estimator and then determine whether the new statistic measures the same object. This protocol is particularly important because convergence speed, funding magnitude, cross-sectional ranking, and time-series co-movement are not interchangeable statistics.

The investigators are not independent of the prior work. Accordingly, this paper is an author-run out-of-sample validation on previously unused external data. The design is stronger than an in-sample restatement but weaker than unaffiliated replication.

5.2 Multiplicity and analysis hierarchy

The empirical programme comprises 81 analyses reported in the accompanying analysis ledger. We distinguish: (i) the four primary restrictions in Section 3; (ii) secondary tests of recovery, regime, estimator, and aggregation sensitivity; and (iii) separately labelled exploratory analyses. We report effect sizes, confidence intervals, exact p -values where available, null results, and construction dependence.

False-discovery-rate methods provide a reference for interpreting related families [Benjamini and Hochberg, 1995], but the main conclusions are not selected from a large set of marginal rejections: they rest on repeated mean reversion, broad rating and tenor ordering, formula-induced distributional pinning, forward event ordering, and rank preservation across distinct samples.

5.3 Identification boundaries

Three boundaries are load-bearing. First, the formula decomposition in Equation (3) is mechanical, while the cross-venue comparison is not causal. Second, spread-implied intensity is a pricing object, not a pure physical hazard. Third, benchmark deletion tests operational dependence and relative ordering, not general-equilibrium independence. These distinctions govern the wording of every result below.

6 Perpetual-futures evidence

6.1 Convergence speed and cross-venue coherence

Across the 41 pre-specified dYdX markets, the basis mean-reverts in every case. The median estimated half-life is approximately 3.5 hours, corresponding to an annualized convergence-speed intensity of order 10^3 . The magnitude is economically plausible for a mechanism designed to close an intraday basis and is categorically different from annual credit or mortality intensities. In the separately labelled 43-market exploratory extension, the median half-life is 3.7 hours.




Figure 3: Distribution of estimated perpetual-basis half-lives in the 43-market exploratory extension. The median is 3.67 hours. The corresponding median in the 41-market pre-specified sample is approximately 3.5 hours. The figure displays convergence speed, not the magnitude of the funding premium.

Daily estimates formed from funding magnitude co-move across venues at a mean correlation of 0.70 across 15 common markets. CEX-to-CEX correlation is higher (0.85) than DEX-to-CEX correlation (0.65). This pattern supports a common market-wide anchoring signal while leaving room for venue design, clientele, and the fixed component to affect measured funding.

6.2 Fixed formula input and signed funding drift

Table 3 reports signed mean funding after converting hourly dYdX settlements into complete non-overlapping eight-hour sums, with a per-market breakdown for the three most liquid cross-venue

pairs. Per-market means vary across markets—for example, ETH yields 1.21 bp on dYdX versus 1.32 bp on Bybit—but the within-market differences are smaller than the venue-level equal-weight means (0.83 bp for dYdX, 1.02 bp for Binance, 1.14 bp for Bybit) suggest. The informative contrast is distributional: the centralized-exchange series cluster around the documented one-basis-point formula point, while dYdX does not exhibit the same mass point.

Table 3: Signed mean funding by market and venue

Market	dYdX	Binance	Bybit
basis points per complete eight-hour interval			
BTC	1.04	1.07	1.20
ETH	1.21	1.29	1.32
XRP	– [†]	– [†]	– [†]
Nine-market equal-weight mean	0.83	1.02	1.14

All observations are expressed per complete non-overlapping eight-hour interval. Hourly dYdX settlements are summed into eight-hour windows; Binance and Bybit observations use the same unit. The venue means are descriptive and do not identify the causal contribution of the fixed formula input to total carry. [†]XRP is present on all three venues in the funding archive but was recovered on fewer than 50 eight-hour windows on at least one venue and is excluded from the headline common-market comparison; the full nine-market set comprises ADA, AVAX, BTC, DOGE, ETH, LINK, LTC, NEAR, and SOL.




Figure 4: Signed mean funding across nine common markets after aggregation to complete non-overlapping eight-hour intervals. Venue means are descriptive; the figure does not identify the fixed formula input.

The full funding distributions are also distinct: the CEX-versus-dYdX comparison gives Cohen’s $d = 0.64$ and Kolmogorov–Smirnov distance near 0.6. On Binance, 44–48% of observations in the selected liquid markets lie at or immediately around 0.0001, the point mass predicted by Proposition 1; the premium-only venue is centered near zero, has no comparable formula-induced plateau, and produces negative observations more frequently. These facts support the mechanical pinning restriction, but total funding remains shaped by demand and arbitrage frictions as emphasized by Schmeling et al. [2026].




Figure 5: Distribution of eight-hour funding on four common markets. The dashed line marks the documented one-basis-point centralized-exchange formula input. The identifying evidence is the CEX mass around that point, not a large difference in venue means.

6.3 Transfer simulation

A balanced long-short Monte Carlo with 2,000 seeds provides a mechanism check rather than market evidence. At $\iota = 0$ the confidence interval for cumulative net transfer straddles zero; the small positive mean is attributable to discretization. Introducing a positive fixed component creates a transfer that rises approximately linearly with ι . The result validates the code-level decomposition but does not determine the legal or economic character of the entire perpetual contract.

7 Credit-market evidence

7.1 Rating and tenor ordering

Spread-implied pricing intensity is broadly ordered by credit quality. Across 736 rated bonds in the canonical build, median $\hat{\kappa}^{\text{credit}, \mathbb{Q}}$ rises from 1.39% for AAA to 9.50% for CCC, with a small AA-A inversion and a pronounced step at the investment-grade/high-yield boundary (Figure 6). Table 4 reports the interquartile ranges and sample sizes for each bucket. The ordering is invariant to a common recovery assumption because $1/(1 - \delta)$ rescales all observations.

Multi-tenor CDS curves slope upward for all 12 sovereigns in the primary sample. In the expanded sukuk panel of 179,618 bond-days and 96 issuers, an issuer-and-month fixed-effects specification yields a maturity slope of 2.85 basis points per year for investment-grade sukuk ($t = 4.1$ with issuer-clustered standard errors) and 2.86 for the full sample ($t = 4.7$). A rating-by-tenor surface (Figure 7) is broadly ordered by rating and monotone in tenor; all fifteen adjacent forward-intensity calculations are non-negative. Principal-component analysis produces the familiar level, slope, and curvature structure, with the first component accounting for 98.7% of variation [Litterman and Scheinkman, 1991].




Figure 6: Median spread-implied pricing intensity by credit rating. The cross-section is broadly ordered, with a small AA–A inversion; the pronounced investment-grade/high-yield step is the principal result.

Table 4: Spread-implied pricing intensity by rating bucket

Rating	Bonds	Median (%)	Q1 (%)	Q3 (%)
AAA	56	1.39	0.56	1.55
AA	188	1.92	1.54	2.45
A	109	1.91	1.42	2.38
BBB	200	2.50	1.80	2.87
BB	107	6.15	2.87	7.53
B	54	6.74	5.87	7.75
CCC	22	9.50	6.36	11.12




Figure 7: Rating-by-tenor surface for spread-implied pricing intensity, reconstructed from CDS data using the `e30_foundational.py` methodology (2023+ means, $\delta = 0.60$). The surface is broadly ordered by rating and monotone in tenor; absolute levels depend on recovery and spread composition.

7.2 Dynamics and forward event ordering

All 20 sovereign spread-implied series are mean-reverting, with AR(1) coefficients between 0.77 and 0.98. Negative observations occur in 7 of 54,047 sovereign spread observations (0.01%), concentrated in Sri Lanka during default; CDS and absolute-basis measures are non-negative by construction.

The default evidence is informative but small-sample. Using mean intensity over 2016–2018, the three issuers that subsequently defaulted or restructured—Ukraine, Lebanon, and Sri Lanka—occupy the top three early-risk ranks among 20 sovereign indices. In a separate exploratory event-time construction, all three fall in the top decile 12 months before their event. These results are descriptive forward ordering, not an estimate of predictive accuracy from a sufficiently large event panel.

7.3 Physical intensity, pricing intensity, and nondefault components

The distinction between $\mathbb{P}$ and $\mathbb{Q}$ is central. The spread-implied measure contains compensation for risk and can also absorb liquidity, tax, and convenience components. The empirical credit literature shows that these nondefault components are material, especially for high-grade bonds [Elton et al., 2001, Longstaff et al., 2005, Huang and Huang, 2012]. We therefore interpret the rating and tenor results as restrictions on a pricing-intensity surface, not evidence that the full spread equals expected loss.

This distinction is visible in the physical-versus-pricing decomposition. In USDA crop insurance, aggregate premium is 9.6% of liability and realised loss cost is 7.9%, a loss ratio of 0.82. In corporate and sovereign credit, the gap between priced intensity and realised default frequency is much larger. The comparison is consistent with standard insurance theory: premiums contain expected loss plus




Figure 8: Early sovereign pricing-intensity ranking and subsequent default or restructuring events. The figure is descriptive because the event count is three.

loadings, and the loading varies by market and risk-bearing technology.

8 Benchmark dependence and relation to monetary rates

8.1 Benchmark-free relative ordering

The benchmark-free estimator in Equation (7) is computed for 856 bonds in the canonical sample. Its ranking correlates 0.82 with the spread-based ranking at the measured corporate recovery convention. The rank correlation is similar across instrument types: 0.82 for sukuk and 0.81 for conventional bonds, indicating that the ordering preservation is not driven by a single instrument category. Subgroup estimates use the subset of bonds with an unambiguous sukuk or conventional classification at the snapshot date used for the subgroup analysis; the full canonical estimator contains 856 bonds. Most of the level difference is captured by a common shift: the median gap is 473 basis points of intensity, compared with 472 basis points obtained by mapping the Treasury yield through the same recovery transformation. Year-by-year, the median gap correlates 0.93 with the 10-year Treasury yield.




Figure 9: Benchmark-free versus spread-implied pricing intensity. The diagonal is a visual reference; the estimators preserve broad ordering but are not identical.

The result has a narrow interpretation. Relative issuer ordering and much of the curve shape can be recovered without supplying the Treasury curve to the estimator. The residual ranking differences show that the two constructions are not identical. Moreover, an asset traded in a dollar

system remains exposed to equilibrium discounting, liquidity, and opportunity-cost conditions even if its formula omits the benchmark.

8.2 Why event intensity is not the monetary rate

Five measurements separate the objects. First, estimated intensities span approximately six orders of magnitude: mortality is of order 10^{-3} to 10^{-2} per year, credit is of order 10^{-2} to 10^{-1} , and liquid-perpetual convergence is of order 10^3 . Second, the ratio of policy rates to sovereign credit intensity varies from 1.6 to 6.6 across the sample. Third, credit intensity correlates only 0.41 with inflation. Fourth, euro-area risk-free yields were negative for 1,644 days while credit intensity remained economically meaningful. Fifth, tokenized Treasury products and tokenized-gold perpetuals coexist on-chain while paying economically different objects: base-rate carry in the former and anchoring premium in the latter.

These facts do not imply that r and κ are statistically unrelated. Both respond to macroeconomic risk, and the benchmark-free level gap tracks the government curve. The supported conclusion is separation of roles: the event-intensity object is not numerically or economically identical to the monetary rate, even though equilibrium links remain in a mixed system.

9 Cross-domain evidence and robustness

Across agriculture, mortality, disaster, climate, macro, and credit channels, the six-domain panel spans 17,417 unique country-year rows (41,397 nonmissing observations). Domain coverage varies considerably — credit has 220 observations, agriculture 9,643 — and only six of the 15 possible domain pairs have sufficient overlapping observations to estimate a pairwise correlation. Among those six estimable pairs, the mean absolute off-diagonal correlation is 0.075. Within-domain dependence is much higher. Table 5 reports domain-level coverage, and Figure 10 displays the pairwise correlation matrix. Pairwise sample sizes and eligibility status of all 15 pairs are reported in Appendix D.

Table 5: Canonical cross-domain panel coverage

Domain	Nonmissing observations	Countries	Year range
Agriculture (cereal yield)	9,643	227	1980–2024
Mortality (death rate)	11,915	264	1980–2024
Disaster count	3,660	223	2000–2026
Climate (temperature)	5,064	211	2000–2023
Macro (GDP growth)	10,895	261	1980–2024
Credit (private credit/GDP)	220	20	2016–2026
Unique country-year rows			17,417
Nonmissing domain-year obs.			41,397

During COVID-19, sovereign credit intensity and mortality rose together, while disaster impact did not; in the 2022 financial and war shock, mortality and cereal yields did not co-move with credit in the same way. Average diversification therefore weakens when domains share a causal real-event channel rather than during every financial disturbance.

Appendix C reports credit-market robustness and recovery sensitivity, Appendix E reports benchmark-free estimator diagnostics, Appendix F reports manipulation tests, protocol simulations,




Figure 10: Cross-domain dependence matrix for the canonical six-domain panel: agriculture, mortality, disaster, climate, macro, and credit. Only six of 15 possible domain pairs have sufficient overlapping observations for estimation; blank cells indicate pairs without adequate coverage. The mean absolute off-diagonal correlation among estimable pairs is 0.075. Sample sizes and eligibility for all 15 pairs are reported in Appendix D.

and monetary extensions, and Appendix G reports the complete analysis ledger. Null results are retained: the macro debt-growth claim does not survive within-country lagged controls; credit intensity does not lead exchange rates; the apparent jump asymmetry is outlier-driven; and several headline magnitudes vary with aggregation and estimator construction.

10 Primary findings and identification status

Table 6 summarizes the empirical restrictions, their evidentiary status, and the principal identification boundaries.

Restriction	Assessment	Evidence and boundary
Perpetual basis mean reversion	Supported	41/41 pre-specified markets; median half-life 3.5 hours.
Cross-venue coherence	Supported	Mean daily correlation 0.70 across 15 common markets.
Fixed-input pinning	Supported mechanically and descriptively	44–48% near the one-bp formula point; signed means are 0.83 bp (dYdX), 1.02 bp (Binance), and 1.14 bp (Bybit); venue effects are not controlled causally.
Rating ordering	Supported with qualification	Broadly ordered from AAA through CCC; small AA–A inversion and pronounced investment-grade/high-yield step.
Tenor ordering	Supported	Upward CDS curves in 12/12 sovereigns; positive panel maturity slope.
Forward event information	Descriptive support	Three subsequent events rank highest in the early sovereign cross-section; only three events.
Benchmark-free ordering	Supported with residual differences	Spearman 0.82 in the 856-bond canonical sample; level gap tracks the Treasury curve.
Physical versus pricing intensity	Supported as a distinction	Small loading in crop insurance; much larger gap in credit; not a legal conclusion.
General-equilibrium independence	Not established	Formula deletion does not remove opportunity costs or mixed-system equilibrium effects.
External replication	Not established	Data are external to the source studies; investigators are not independent.

Table 6: What the evidence supports and what remains outside identification.

11 Economic interpretation and Islamic-finance boundary

The empirical contribution is the separation of a measurable event-intensity object from an exchange-specified fixed funding input. In Islamic-finance terms, this isolates the component relevant to a charge stipulated for time alone. It does not establish that a perpetual, credit-protection contract, sukuk, or insurance product is permissible overall. Questions of excessive uncertainty, gambling, possession, leverage, settlement, asset backing, and the legal status of risk premia require qualified jurisprudential judgment.

The physical-versus-pricing decomposition helps state that question precisely. Where priced intensity closely tracks realised loss, as in the crop-insurance comparison, the economic object is primarily compensation for observed claims plus a modest loading. Where priced intensity substantially exceeds physical event frequency, as in credit, the residual includes risk compensation and market frictions. The data can measure that wedge and test its behavior; they cannot convert it into a legal ruling.

12 Conclusion

The evidence supports a narrow but consequential distinction. The anchoring intensity in perpetual futures is observable, mean-reverting, and coherent across venues. A fixed exchange input appears as a formula-induced concentration of observations near its stipulated value. Harmonized venue means are descriptively different but too close to identify the causal contribution of the formula to total carry. The mapped credit-market object is broadly ordered by rating and monotone in tenor, predominantly non-negative, and informative about subsequent events. Relative credit ordering can also be recovered without using a government curve as an operational input, but levels and equilibrium exposure remain connected to the monetary environment.

The study therefore validates several restrictions of an event-time pricing architecture without establishing a universal pricing primitive, a causal venue experiment, or an overall Shariah verdict. The decisive next tests are unaffiliated replication, prospective analysis on a publicly frozen sample, and traded evidence from instruments designed ex ante around the separated components.

Data and code availability

The data and code supporting the findings are available from the corresponding author upon reasonable request, subject to applicable third-party licensing restrictions.

Declaration of generative AI and AI-assisted technologies

During preparation of the manuscript, the authors used generative AI tools for code review and debugging assistance, data-formatting support, consistency checks, and language editing. All analyses were executed by the authors, who independently reviewed and verified the source data, code, statistical outputs, interpretations, and conclusions and accept full responsibility for the manuscript.

References

- Abedifar, P., Molyneux, P., and Tarazi, A. (2013). Risk in Islamic banking. *Review of Finance*, 17(6), 2035–2096.
- Abedifar, P., Ebrahim, S. M., Molyneux, P., and Tarazi, A. (2015). Islamic banking and finance: Recent empirical literature and directions for future research. *Journal of Economic Surveys*, 29(4), 637–670.
- Akerer, D., Hugonnier, J., and Jermann, U. (2025). Perpetual futures pricing. *Mathematical Finance*, 36(3), 481–499. <https://doi.org/10.1111/mafi.70018>.
- Ahmed, S., and Bhuyan, R. (2026a). The interest parameter in perpetual futures: Shariah analysis and empirical evidence from centralized and decentralized exchanges. Working paper.
- Ahmed, S., and Bhuyan, R. (2026b). From perpetual contracts to Islamic credit: The random stopping time equivalence and a foundation for sukuk, takaful, and credit protection. Working paper, SSRN 6322858.
- Alexander, C., Choi, J., Park, H., and Sohn, S. (2020a). BitMEX bitcoin derivatives: Price discovery, informational efficiency, and hedging effectiveness. *Journal of Futures Markets*, 40(1), 23–43.
- Alexander, C., and Heck, D. F. (2020). Price discovery in Bitcoin: The impact of unregulated markets. *Journal of Financial Stability*, 50, 100776.
- Alexander, C., Choi, J., Massie, H. R. A., and Sohn, S. (2020b). Price discovery and microstructure in ether spot and derivative markets. *International Review of Financial Analysis*, 71, 101506.
- Angeris, G., Chitra, T., Evans, A., and Lorig, M. (2023). A primer on perpetuals. *SIAM Journal on Financial Mathematics*, 14(1), SC17–SC30.
- Azad, A. S. M. S., Azmat, S., Chazi, A., and Ahsan, A. (2018). Can Islamic banks have their own benchmark? *Emerging Markets Review*, 35, 120–136. <https://doi.org/10.1016/j.ememar.2018.02.002>.
- Beck, T., Demirguc-Kunt, A., and Merrouche, O. (2013). Islamic vs. conventional banking: Business model, efficiency and stability. *Journal of Banking & Finance*, 37(2), 433–447.
- Benjamini, Y., and Hochberg, Y. (1995). Controlling the false discovery rate: A practical and powerful approach to multiple testing. *Journal of the Royal Statistical Society: Series B*, 57(1), 289–300.
- Binance. (2026). Introduction to Binance Futures funding rates. Official product documentation, accessed July 2026.
- Brunnermeier, M. K., and Pedersen, L. H. (2009). Market liquidity and funding liquidity. *Review of Financial Studies*, 22(6), 2201–2238.
- Buhlmann, H. (1980). An economic premium principle. *ASTIN Bulletin*, 11(1), 52–60.
- Bybit. (2026). Funding fee calculation. Official product documentation, accessed July 2026.
- Collin-Dufresne, P., Goldstein, R. S., and Martin, J. S. (2001). The determinants of credit spread changes. *Journal of Finance*, 56(6), 2177–2207.

- dYdX. (2026). Funding and market parameters. Official v4 protocol documentation and source code, accessed July 2026.
- Driessen, J. (2005). Is default event risk priced in corporate bonds? *Review of Financial Studies*, 18(1), 165–195.
- Duffee, G. R. (1999). Estimating the price of default risk. *Review of Financial Studies*, 12(1), 197–226.
- Duffie, D., and Singleton, K. J. (1999). Modeling term structures of defaultable bonds. *Review of Financial Studies*, 12(4), 687–720.
- Elton, E. J., Gruber, M. J., Agrawal, D., and Mann, C. (2001). Explaining the rate spread on corporate bonds. *Journal of Finance*, 56(1), 247–277.
- Gilchrist, S., and Zakrajsek, E. (2012). Credit spreads and business cycle fluctuations. *American Economic Review*, 102(4), 1692–1720.
- Harvey, C. R., Liu, Y., and Zhu, H. (2016). ... and the cross-section of expected returns. *Review of Financial Studies*, 29(1), 5–68.
- He, S., Manela, A., Ross, O., and von Wachter, V. (2025). Fundamentals of perpetual futures. Working paper, SSRN 4301150, revised June 2025.
- Huang, J.-Z., and 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.
- Iqbal, Z. (1999). The benchmark issue in the Islamic financial system. *American Journal of Islam and Society*, 16(2), 45–53. <https://doi.org/10.35632/ajis.v16i2.2128>.
- Kahane, Y. (1979). The theory of insurance risk premiums: A re-examination in the light of recent developments in capital market theory. *ASTIN Bulletin*, 10(2), 223–239.
- Kim, J., and Park, H. (2025). Designing funding rates for perpetual futures in cryptocurrency markets. Working paper, arXiv:2506.08573.
- Litterman, R., and Scheinkman, J. (1991). Common factors affecting bond returns. *Journal of Fixed Income*, 1(1), 54–61.
- Longstaff, F. A., Mithal, S., and Neis, E. (2005). Corporate yield spreads: Default risk or liquidity? New evidence from the credit default swap market. *Journal of Finance*, 60(5), 2213–2253.
- Longstaff, F. A., Pan, J., Pedersen, L. H., and Singleton, K. J. (2011). How sovereign is sovereign credit risk? *American Economic Journal: Macroeconomics*, 3(2), 75–103.
- Pan, J., and Singleton, K. J. (2008). Default and recovery implicit in the term structure of sovereign CDS spreads. *Journal of Finance*, 63(5), 2345–2384.
- Perignon, C., Akmansoy, O., Hurlin, C., Dreber, A., Holzmeister, F., Huber, J., Johannesson, M., Kirchler, M., Menkveld, A. J., Razen, M., and Weitzel, U. (2024). Computational reproducibility in finance: Evidence from 1,000 tests. *Review of Financial Studies*, 37(11), 3558–3593.
- Schmeling, M., Schrimpf, A., and Todorov, K. (2026). Crypto carry. *Management Science*, published online May 6. <https://doi.org/10.1287/mnsc.2024.05069>.

- Shiller, R. J. (1993). Measuring asset values for cash settlement in derivative markets: Hedonic repeated measures indices and perpetual futures. *Journal of Finance*, 48(3), 911–931.
- Sullivan, R., Timmermann, A., and White, H. (1999). Data-snooping, technical trading rule performance, and the bootstrap. *Journal of Finance*, 54(5), 1647–1691.
- Uddin, M. H., Kabir, S. H., Hassan, M. K., Hossain, M. S., and Liu, J. (2022). Why do sukuk (Islamic bonds) need a different pricing model? *International Journal of Finance & Economics*, 27(2), 2210–2234.
- Wang, S. S. (2002). A universal framework for pricing financial and insurance risks. *ASTIN Bulletin*, 32(2), 213–234.
- White, H. (2000). A reality check for data snooping. *Econometrica*, 68(5), 1097–1126.

A Data construction and sample accounting

This appendix documents the construction of the primary datasets. The study uses seventeen data providers across five domain categories: perpetual-futures markets, fixed-income and credit markets, insurance and loss records, macro-financial time series, and natural-disaster registries.

Perpetual-futures data. Hourly funding-rate observations are drawn from three venues: dYdX (premium-only, v4 protocol via public API), Binance (USD-margined perpetual futures via public REST API), and Bybit (linear perpetuals via public REST API). The sample window begins with each market’s first listing and extends through June 2026. Each observation is a complete non-overlapping eight-hour funding settlement interval, constructed from the raw hourly feed by summing consecutive hourly observations within each UTC eight-hour window. Markets with fewer than 50 complete eight-hour windows are excluded from the primary analysis. The common-market sample comprises the nine contracts listed on all three venues: ADA, AVAX, BTC, DOGE, ETH, LINK, LTC, NEAR, and SOL.

Credit-market data. Bond-level histories are sourced from Cbonds, covering 921 sovereign and corporate USD-denominated bonds with available pricing, rating, and contractual-maturity fields. The sample includes both sukuk and conventional instruments. The benchmark-free snapshot retains the latest valid observation with the price/yield, spread, and maturity fields required by Equation (7). The rating analysis then restricts that snapshot to the 736 bonds with a recognized credit rating. Sovereign CDS curves for 51 sovereigns are sourced from a cross-section of market makers. The rating-tenor surface (Figure 7 of the main paper) is reconstructed from CDS data using `e30_foundational.py`, with 2023+ CDS means and $\delta = 0.60$.

Cross-domain panel. The six-domain panel combines: (i) agriculture (World Bank cereal-yield and crop-production indicators), (ii) mortality (World Bank death-rate and Human Mortality Database life-table series), (iii) disaster (EM-DAT natural-disaster event counts), (iv) climate (World Bank climate-variability indices), (v) macro (World Bank GDP growth, debt-to-GDP, and current-account indicators), and (vi) credit (the sovereign κ estimates from the Cbonds corporate and sovereign panel). Each series is expressed in country-year units; the panel is unbalanced because domain coverage varies substantially across countries and years.

Other data. Additional sources are described in the main paper and the replication archive. These include: USDA Risk Management Agency crop-insurance book (139 commodity lines, \$1.25 trillion liability); Human Mortality Database life tables (41 countries); EM-DAT natural-disaster registry (194 countries, 2000–2024); World Bank World Development Indicators; and dYdX total-value-locked histories for the liveness-trigger analysis.

B Additional perpetual-futures evidence

B.1 Proof of the formula-induced pinning result

Proof of Proposition 1 in the main paper. By definition,

$$\text{clamp}(\iota - P, -c, c) = \begin{cases} c, & \iota - P > c, \\ \iota - P, & -c \leq \iota - P \leq c, \\ -c, & \iota - P < -c. \end{cases}$$

Adding P gives the three branches in Equation (5) of the main paper. Every premium realization in the interval $[\iota - c, \iota + c]$ maps to the single value ι ; because P is continuous, the image distribution therefore has an atom of mass $q_\iota = \Pr(\iota - c \leq P \leq \iota + c)$. Partitioning the expectation over the three regions yields the stated mean. Under the premium-only rule $F^P = P$, a continuous P remains continuous, so no atom is induced by the formula. $\square$

B.2 Cross-venue coherence and distributional diagnostics

The main paper reports cross-venue daily correlation of 0.70 across 15 common markets. The correlation is stable across subperiods; splitting the sample at January 2023 gives pre-2023 mean 0.67 and post-2023 mean 0.72, consistent with convergence in market microstructure across venues.

The pinning distribution (Figure 4 in the main paper) shows a pronounced mass near the one-basis-point formula point on Binance and Bybit that is absent on dYdX. Table 7 reports the exact percentages by venue and market. The pattern is uniform across individual markets: no dYdX market exceeds 5% of observations within ± 0.25 bp of the formula point, while every Binance and Bybit market exceeds 30%.

Table 7: Percentage of eight-hour funding observations within ± 0.25 bp of the one-basis-point formula point, by venue and market

Market	dYdX (%)	Binance (%)	Bybit (%)
BTC	4.2	42.1	38.9
ETH	3.8	44.3	41.2
ADA	2.9	39.8	36.7
AVAX	3.1	41.5	40.2
DOGE	4.0	43.0	37.5
LINK	3.5	40.6	39.1
LTC	2.7	38.2	35.8
NEAR	3.3	42.7	40.9
SOL	4.1	45.1	42.3
Mean	3.5	41.9	39.2

C Credit-market robustness and recovery sensitivity

C.1 Reconciliation of construction-specific estimates

The reproduce-first protocol described in the main paper 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 distinct 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 ≈ 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 $\times$). 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 separates robust directional findings from construction-specific magnitudes.

C.2 Recovery sensitivity

The main paper’s credit estimates use a recovery assumption of 60% for the spread-based estimator, following the convention in the reduced-form credit literature. Table 8 reports the sensitivity of the median pricing intensity to alternative recovery assumptions. The rating ordering is preserved across all recovery values; the level shifts mechanically through $(1 - \delta)^{-1}$.

Table 8: Median pricing intensity (% per year) by rating and assumed recovery rate

Rating	$\delta = 0.40$	$\delta = 0.60$ (baseline)	$\delta = 0.80$
AAA	0.93	1.39	2.78
AA	1.28	1.92	3.84
A	1.27	1.91	3.82
BBB	1.67	2.50	5.00
BB	4.10	6.15	12.30
B	4.49	6.74	13.48
CCC	6.33	9.50	19.00

D Cross-domain panel construction and pairwise sample sizes

D.1 Variable construction

Each of the six domain variables is constructed as follows:

- **Agriculture:** log cereal yield (kg per hectare) from the World Bank, annually per country. Series are detrended by subtracting the country-specific linear trend.
- **Mortality:** death rate (per 1,000 population) from the World Bank. For the Human Mortality Database validation, the central death rate μ_x at each age is used directly.
- **Disaster:** annual natural-disaster event count per country from EM-DAT, expressed as events per year.
- **Climate:** precipitation anomaly (standardised deviation from the 1961–1990 mean) from the World Bank climate portal.
- **Macro:** real GDP growth rate (annual %) from the World Bank, winsorized at the 1st and 99th percentiles.
- **Credit:** sovereign $\hat{\kappa}$ from the Cbonds USD-bond panel (the same estimator used in the main paper’s credit analysis), computed as the most recent annual observation per country with at least one rated bond.

D.2 Overlap rules

A pairwise correlation is estimated only when two domains share at least 20 country-year observations for the same country and year. The threshold is set to ensure a minimum effective sample for a correlation estimate; results are robust to thresholds of 10 and 30 observations.

D.3 Domain coverage

Table 9 reports the number of nonmissing country-year observations for each of the six domains, along with the pairwise sample sizes for all 15 domain pairs. The six-domain panel spans 17,417 unique country-year rows.

Of the 15 possible pairwise correlations, only six domain pairs have at least 20 overlapping country-year observations and are therefore estimable. The six eligible pairs fall into two clusters: the three large World Bank macro domains (mortality, agriculture, and GDP growth) overlap extensively with each other, while the three narrower domains (natural disasters, climate anomalies, and sovereign credit) overlap among themselves but not with the macro panel. The six pairwise correlations and their sample sizes are reported in Figure 10 of the main paper. Specific pairwise sample sizes for all 15 pairs are available in the replication archive (`crossdomain_six_domain_panel.csv`).

Table 9: Domain coverage and pairwise sample sizes

Panel A: Nonmissing country-year observations by domain

Domain	Nonmissing observations
Agriculture (cereal yield)	9,643
Mortality (death rate)	11,915
Disaster (event count)	3,660
Climate (temperature anomaly)	5,064
Macro (GDP growth)	10,895
Credit ($\hat{\kappa}$)	220
Panel total unique country-years	17,417

Panel B: Pairwise overlapping observations

Domain pair	Pairwise n	Eligible ($n \geq 20$)
Mortality $\times$ Macro	10,895	Yes
Mortality $\times$ Agriculture	9,643	Yes
Macro $\times$ Agriculture	9,390	Yes
Climate $\times$ Disaster	3,239	Yes
Credit $\times$ Disaster	177	Yes
Credit $\times$ Climate	160	Yes
Mortality $\times$ Climate	0	No
Mortality $\times$ Disaster	0	No
Mortality $\times$ Credit	0	No
Macro $\times$ Climate	0	No
Macro $\times$ Disaster	0	No
Macro $\times$ Credit	0	No
Agriculture $\times$ Climate	0	No
Agriculture $\times$ Disaster	0	No
Agriculture $\times$ Credit	0	No

E Benchmark-free estimator diagnostics

E.1 Recovery convention sensitivity

The benchmark-free estimator in Equation (7) of the main paper uses a recovery assumption of 10% for the price-based calculation, matching the convention in the corporate-bond literature. Table 10 reports the rank correlation for alternative recovery values.

Table 10: Spearman rank correlation between benchmark-free and spread-based ordering under alternative recovery assumptions

Recovery rate (δ)	Spearman ρ	95% CI
0.05	0.81	[0.78, 0.84]
0.10 (baseline)	0.82	[0.79, 0.85]
0.20	0.82	[0.79, 0.85]
0.30	0.81	[0.78, 0.84]

The rank correlation is essentially invariant to the recovery assumption: the ordering preservation is a property of the cross-sectional κ distribution, not an artifact of the recovery convention.

E.2 Subgroup consistency

The main paper reports that the benchmark-free ordering is consistent across sukuk ($\rho = 0.82$) and conventional bonds ($\rho = 0.81$). The subgroup estimates use the subset of bonds with an unambiguous sukuk or conventional classification at the snapshot date used for the subgroup analysis; the full canonical estimator contains 856 bonds.

Table 11: Benchmark-free ordering by instrument type

Subgroup	N	Spearman ρ	Median κ (%/yr)
Sukuk	273	0.82	3.12
Conventional	595	0.81	3.78
Unclassified	88	—	—

The 88 unclassified bonds are those whose snapshot classification was ambiguous or unavailable at the analysis date. The classification snapshot (868 bonds with unambiguous classification, 956 total identified) draws on a different vintage than the canonical 856-bond analysis sample; the temporal gap in snapshots accounts for the difference in totals, and all 856 canonical bonds are included in the 956 identified in the classification snapshot.

F Protocol and simulation evidence

F.1 Manipulation resistance

Paper 13’s defense of the published κ against tampering has two parts; the oracle depth rule needs live order-book depth, 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 previously unused external data.

F.2 Depth rule and manipulation frontier

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 \$1.3M on dYdX rising to \$26M 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 $\sim 1,700\times$ at 10 bp rising past $10,000\times$ at 200 bp, 95% CIs excluding 1), reproducing the prior working-paper sequence’s $72\text{--}2,739\times$ 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.

F.3 Liveness trigger

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 $\rightarrow$ \$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.

F.4 System-level simulations

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 out-of-sample 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.

F.5 Monetary extensions

The monetary half of the prior working-paper sequence — 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 $\rightarrow$ 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 $\rightarrow$ 0.139 Egypt, bracketing the paper’s range; the index leg runs 1.3–2.6 $\times$ 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 author-run out-of-sample support 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 is larger:

$$S^* \propto \sigma^{2.3}(1 - \rho)^{-2.2}, \quad S_{\text{corp}}^* \sim 25\%, \quad S_{\text{sov}}^* \sim 12\%.$$

The sovereign–corporate ordering therefore survives on a sample with roughly three times the number of 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 856-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, supported on the expanded sample. 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 ≥ 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 Appendix C.1. 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, 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 $\sim 12\%$ (per-issuer persistence) to $\sim 45\%$ (index persistence), so the prior working-paper sequence’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.

G Analysis hierarchy, exploratory tests, and null results

G.1 Analysis hierarchy

The primary validation ledger contains 81 analyses across the main paper and appendices. The analysis hierarchy was recorded in a dated internal plan dated March 2026, which classified each analysis as primary (confirmatory), secondary (descriptive), or exploratory. The plan was not externally preregistered; exploratory analyses are labelled separately in the replication archive. Appendix G.2 separately summarises 30 additional exploratory extensions outside the 81-analysis ledger.

Table 12 summarises the distribution of analyses by category and outcome.

Table 12: Summary of the 81-analysis ledger

Category	Confirmed	Null/refuted	Total
Primary (pre-specified)	16	2	18
Secondary (descriptive)	31	4	35
Exploratory (planned, not pre-specified)	22	6	28
Total	69	12	81

G.2 Exploratory analyses beyond the pre-specified set

Following the 81-analysis ledger, 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 validation exercise that omits failures would not be informative.

κ as a hazard — supported in the reported sample.

- *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).
- *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).
- *Mean-reverting intensity, not a random walk:* forward κ regresses on its own deviation with slope -0.26 , AR(1) half-life ~ 3.4 months (A109).
- *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).
- *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).
- *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).

- $\sim 67\%$ of sovereign- κ variation is one systematic factor (A120).

Interest is removable, with measured consequence — supported in the reported sample.

- 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: 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 in the main paper).
- Across the panel rate risk is 40–52% of realised yield volatility, ballooning in the rate-shock year (A108).
- At $\iota = 0$ the funding distribution has no formula-induced mass point around one basis point, consistent with the premium-only design (the dYdX nine-market equal-weight signed mean is 0.83 bp per complete eight-hour interval); the distribution is continuous rather than pinned at a fixed-input value (A110).

Multi-modal oracle security — supported in the reported sample.

- 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 $\sim \$23\text{M}$ 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 ~ 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 results delimit the supported claims and record the failures alongside them.

G.3 Cross-domain null results

The cross-domain analysis in the main paper and Appendix D reports a mean absolute off-diagonal correlation of 0.075 among six estimable pairs. The null $\rho_x = 0$ is not rejected for any individual pair at conventional significance levels after Bonferroni correction for the planned family of 15 pairs. The small magnitudes are therefore consistent with approximate independence among the six estimable domain pairs in the canonical construction — the assumption that underlies the κ -ladder diversification argument.

EDITORIAL NOTE TO THIS PRINTING

Paper 16 · Riba, Reduced

Working paper, revised July 2026 · SSRN 6970078

Role in the programme. The reduction. Fifty years of argument compressed into one measured quantity. **What it establishes.** The decomposition of priced intensity into realised hazard and risk premium, by domain: about 18% premium share in takaful; about 91% (upper bound) in corporate credit at the pricing convention, about 81% at the measured recovery; at least 75% in sovereign credit, with a conservative floor near 45% that survives every substitution tested. **Corrected or extended later.** The measured-recovery row was added in July 2026 and runs in the conservative direction. **Not established here.** The answer. The question is located to the basis point and reserved; that is the paper's point, not its failure.

Replication for every number in this paper: DOI 10.5281/zenodo.21290557 · github.com/shehzadahmed-xx/kappa-replication

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 Akerer 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 $\sim 18\%$ for takaful (premiums equal expected loss plus a margin); $\approx 91\%$ for corporate credit at the conventional recovery, and still $\approx 81\%$ when re-run at the recovery actually measured on defaulted sukuk ($\delta \approx 0.10$); and the great majority for sovereign credit ($\geq 75\%$ at the sample median, and $\geq \sim 45\%$ even under the most conservative pairing of the Poisson default bound with the tightest investment-grade spread — a floor that also survives substituting the measured low recovery for the convention; point $\sim 85\text{--}95\%$), 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

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

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

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 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,

$$\kappa_Q = \frac{s}{1 - \delta}, \quad (1)$$

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*

$$\Pi = 1 - \frac{\kappa_P}{\kappa_Q} = 1 - \frac{\text{EL}}{s}, \quad \text{EL} = \kappa_P(1 - \delta), \quad (2)$$

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. The recovery robustness is concrete rather than asserted: re-running the

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), $\delta = 0.60$ convention	2.2%	0.19%	$\approx 91\%$ (upper bound)
at measured recovery $\delta \approx 0.10$	0.98%	0.19%	$\approx 81\%$ (upper bound)
Sovereign credit (sukuk)	4.0%	$\leq 1.0\%$ (Poisson UB)	$\geq 75\%$, point $\sim 85\text{--}95\%$

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 (~ 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.” The corporate row is also re-run at the recovery the corpus itself measures: realised recovery on defaulted asset-based corporate sukuk clusters near $\delta \approx 0.10$, not the 0.60 convention [Ahmed & Bhuyan, 2026, P2c], which lowers κ_Q to $s/0.9 = 0.98\%$ and the share to $\approx 81\%$ — the convention *overstates* the corporate share, so the correction runs in the conservative direction. Sovereign recovery is unobserved (no sovereign sukuk has defaulted) and stays at the Moody’s convention; but even if it were as low as the corporate measurement, the median-spread pairing gives $\kappa_Q = 1.6\%/0.9 = 1.78\%$ and $\Pi \geq 1 - 1.0/1.78 \approx 44\%$, so the $\sim 45\%$ floor survives the substitution (the tightest-spread pairing becomes uninformative at low recovery; the median-spread bound replaces it at an almost unchanged level).

corporate row at the recovery actually measured on this market’s own defaulted sukuk ($\delta \approx 0.10$; Table 1) moves its share from $\approx 91\%$ to $\approx 81\%$, and the conservative sovereign floor is attained under both the convention (Poisson bound against the tightest spread at $\delta = 0.60$: 43%) and the measured recovery (Poisson bound against the median spread at $\delta = 0.10$: 44%).

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 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.
- Ahmed, S., & Bhuyan, R. (2026). An Islamic Credit Default Swap on Smart-Contract Infrastructure: Design, Pricing, and Regulatory Pathway.
- 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.

EDITORIAL NOTE TO THIS PRINTING

Paper 17 · The Separation Theorem

Working paper, revised July 2026 SSRN 6970118

Role in the programme. The theorem under everything, stated properly and last. **What it establishes.** Within the event-time pricing class: the pricing operator's pole sits at $\iota = -\kappa$, leaving the riba-free point regular; the structural form of the risk-premium share, $\Pi = |\lambda|/\alpha$; the measure change under a CIR intensity. **Corrected or extended later.** It is itself the corrector: it restates Paper 2's pole discussion in the event parametrisation and reconciles the two (it is the event, not the coordinates, that makes the riba-free point regular). **Not established here.** Anything outside the stated model class; the scope is the theorem's, and the theorem is the scope.

Replication for every number in this paper: DOI 10.5281/zenodo.21290557 · github.com/shehzadahmed-xx/kappa-replication

The Separation Theorem: Interest as a Removable Parameter in No-Arbitrage Pricing

Shehzad Ahmed* Rafiqul Bhuyan†

June 2026 — Preliminary draft

Abstract

The κ -rate programme rests on a single claim: that the interest rate is a *separable and removable* parameter of no-arbitrage pricing, leaving behind a convergence intensity κ that is a hazard, not a discount. This paper gives that claim its rigorous statement and proof, and develops the stochastic machinery it requires. We work in the *event parametrisation* of the perpetual price of Akerer et al. [2025]: an expectation of a payoff at a Cox (doubly-stochastic) stopping time θ with intensity κ , discounted at the interest parameter ι — the form that holds the real-event intensity fixed while the pure time-charge varies, and coincides with AHJ’s funding parametrisation exactly at the riba-free point $\iota = 0$ (away from it, AHJ’s ι enters both the discount and the effective stopping intensity; see §3). Our central result resolves the oldest objection to interest-free pricing—that setting the rate to zero makes prices diverge. We show this is false for the κ -priced operator and true only for the conventional one, and we locate the difference exactly: as a function of ι , the pricing operator is holomorphic on $\{\text{Re}(\iota) > -\kappa\}$ with its only singularity a simple pole at $\iota = -\kappa$; hence for $\kappa > 0$ the point $\iota = 0$ is interior and regular, and the riba-free price $f(0) = \mathbb{E}^{\mathbb{Q}}[X_{\theta}]$ is finite and well defined. The conventional perpetual operator $V(r) = c/r$ has its pole at $r = 0$, which is precisely why interest there is not removable. The convergence intensity $\kappa > 0$ moves the pole off the origin; that displacement *is* the Separation Theorem. We then prove exact multiplicative separation of ι from κ in the event parametrisation, the formal isomorphism with the Duffie–Singleton defaultable claim ($\kappa \leftrightarrow \lambda$, $\iota \leftrightarrow r$ here; equivalently $\kappa_{\text{AHJ}} - \iota \leftrightarrow \lambda$ in the funding parametrisation, the map of the credit-equivalence paper), specify κ as a CIR diffusion and carry it through a change of measure that preserves the Feller condition, derive the affine κ -term-structure, and obtain a corollary identifying the empirically measured risk-premium share with $|\lambda|/\alpha$. The κ -priced monetary base appears as a closing application.

Keywords: separation theorem, convergence intensity, Cox process, removable singularity, Cox–Ingersoll–Ross, change of measure, affine term structure, riba.

1 Setup

Let $(\Omega, \mathcal{F}, (\mathcal{F}_t)_{t \geq 0}, \mathbb{Q})$ be a filtered probability space under the pricing (risk-neutral) measure $\mathbb{Q}$, satisfying the usual conditions, and let $X = (X_t)$ be the adapted spot price of an underlying asset.

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

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

Assumption 1 (Cox convergence time). The convergence intensity $(\kappa_t)_{t \geq 0}$ is non-negative and adapted, and the convergence (or default) time θ is the first jump of a Cox process directed by κ : $\theta = \inf\{t : \int_0^t \kappa_s ds \geq E\}$ with $E \sim \text{Exp}(1)$ independent of $\mathcal{F}_\infty$, so that $\mathbb{Q}(\theta > t \mid \mathcal{F}_\infty^{X, \kappa}) = \exp(-\int_0^t \kappa_s ds)$. We assume $\mathbb{E}^\mathbb{Q}[X_\theta] < \infty$.

Definition 1 (The pricing operator). For an interest parameter $\iota \geq 0$ the no-arbitrage value of the perpetual claim paying X_θ at θ is

$$f(\iota) = \mathbb{E}^\mathbb{Q}[e^{-\iota\theta} X_\theta]. \quad (1)$$

A remark on parametrisation is owed here, because it is where a careless statement of separation would overreach. Equation (1) is exactly the [Duffie & Singleton \[1999\]](#) value of a defaultable claim (intensity κ , discount ι), and it is the *event parametrisation* of the perpetual price of [Ackerer et al. \[2025\]](#): in AHJ’s own funding parametrisation the random-time representation carries the interest parameter in *both* the discount and the effective stopping intensity, $\theta \sim \text{Exp}(\kappa_{\text{AHJ}} - \iota)$ [Theorem 2 of [Ahmed & Bhuyan, 2026, P2](#)], so the two are related by the substitution $\kappa = \kappa_{\text{AHJ}} - \iota$ and coincide identically at $\iota = 0$. Definition 1 deliberately holds the *real-event* intensity κ fixed while the pure time-charge ι varies — the economically meaningful separation experiment, and the one every statement below refers to. In the funding parametrisation, where moving ι also moves the effective intensity, separation is *not* multiplicative for $\iota \neq 0$ (the credit-equivalence paper states this explicitly); nothing in this paper claims otherwise. With a stochastic short rate the factor is $e^{-\int_0^\theta \iota_s ds}$ and all statements below carry over with ι read as the long-run mean.

2 Exact separation

Theorem 1 (Multiplicative separation). *In (1) the interest parameter enters only through the discount functional $D_\iota(\theta) = e^{-\iota\theta}$ and the convergence intensity only through the law of θ . Consequently*

$$f(\iota) = \mathbb{E}^\mathbb{Q}[D_\iota(\theta) X_\theta], \quad D_\iota \equiv 1 \text{ iff } \iota = 0, \quad (2)$$

and the riba-free price is the undiscounted expectation $f(0) = \mathbb{E}^\mathbb{Q}[X_\theta]$, exactly and without remainder.

Proof. $D_\iota(\theta) = e^{-\iota\theta}$ depends on ι alone; the law of θ under Assumption 1 depends on κ alone; X_θ on neither. The integrand factorises as $D_\iota(\theta) \cdot X_\theta$. Setting $\iota = 0$ gives $D_0 \equiv 1$ pointwise, so $f(0) = \mathbb{E}^\mathbb{Q}[X_\theta]$ identically; no limit or approximation is taken. $\square$

Theorem 1 is the algebraic content of separation: interest is a multiplicative overlay on a price whose risk content is carried entirely by θ . The scope is the event parametrisation of Definition 1: in AHJ’s funding parametrisation the same claim would be false for $\iota \neq 0$, because there the interest parameter also shifts the effective stopping intensity — which is precisely the caveat the credit-equivalence paper attaches to its own Separation Theorem [[Ahmed & Bhuyan, 2026, P2](#)]. The two statements are one statement in two coordinate systems, agreeing identically at $\iota = 0$. What Theorem 1 does not yet establish is that $f(0)$ is *finite and non-degenerate*—that removing the overlay does not collapse the price. That is the next, and central, result.

3 The central result: $\iota = 0$ is regular, $r = 0$ is a pole

The classical reason interest-free pricing was thought impossible is that the naive limit is degenerate: a perpetual paying a coupon stream c is worth c/r , which diverges as $r \rightarrow 0$. We show the κ -priced operator does not share this defect, and we identify precisely why.

Theorem 2 (Regularity of the riba-free point). *Suppose $\kappa_t \equiv \kappa > 0$ is constant and X_θ is $\mathbb{Q}$ -integrable. Then:*

- (i) *The map $\iota \mapsto f(\iota)$ extends to a holomorphic function on the half-plane $H_\kappa = \{\iota \in \mathbb{C} : \text{Re}(\iota) > -\kappa\}$, whose only singularity on ∂H_κ is at $\iota = -\kappa$. In particular f is real-analytic on $(-\kappa, \infty)$, the point $\iota = 0$ is interior to H_κ , and the power series of f about 0 has radius of convergence κ .*
- (ii) *In contrast, the conventional perpetual operator $V(r) = \mathbb{E}[\int_0^\infty e^{-rs} c ds] = c/r$ is holomorphic on $\{\text{Re}(r) > 0\}$ with a simple pole at $r = 0$.*
- (iii) *Hence $\iota = 0$ is a regular (removable) point of the κ -priced operator, while $r = 0$ is a non-removable singularity of the conventional operator.*

Proof. (i) Under Assumption 1 with constant κ , $\theta \sim \text{Exp}(\kappa)$. For a bounded payoff, $|X_\theta| \leq M$, dominated convergence gives, for $\text{Re}(\iota) > -\kappa$,

$$f(\iota) = \mathbb{E}[e^{-\iota\theta} X_\theta], \quad |f(\iota)| \leq M \mathbb{E}[e^{-\text{Re}(\iota)\theta}] = M \frac{\kappa}{\kappa + \text{Re}(\iota)} < \infty.$$

The scalar Laplace transform $\mathbb{E}[e^{-\iota\theta}] = \int_0^\infty e^{-\iota t} \kappa e^{-\kappa t} dt = \kappa/(\kappa + \iota)$ is holomorphic on H_κ with a single simple pole at $\iota = -\kappa$; by Morera/dominated differentiation under the expectation, f inherits holomorphy on H_κ , and the distance from 0 to the nearest singularity is κ . For unbounded but integrable X_θ the same conclusion holds on $\{\text{Re}(\iota) \geq 0\}$ by dominated convergence, which already places 0 in the interior of the domain of analyticity. (ii) $\int_0^\infty e^{-rs} c ds = c/r$ for $\text{Re}(r) > 0$, with a simple pole at 0. (iii) Immediate from (i)–(ii): f is analytic at $\iota = 0$ (value $f(0) = \mathbb{E}[X_\theta]$, finite), whereas V has a pole at $r = 0$. $\square$

Remark 1 (The mechanism, in one sentence). The pole of the discounted-expectation operator sits at $\iota = -\kappa$. The convergence intensity $\kappa > 0$ is exactly the margin by which the riba-free point $\iota = 0$ clears that pole. In the conventional perpetual there is no terminating event ($\kappa = 0$), the pole collapses onto the origin, and the discount rate is doing double duty—both discounting and supplying the only convergence—so it cannot be removed. The κ -model separates those two roles: ι discounts, κ converges. Removing ι is harmless precisely because κ , not ι , bounds the horizon ($\theta < \infty$ a.s.).

Remark 2 (Which operator owns the pole). The pole at $\iota = -\kappa$ is a property of the event-parametrised operator (1) — equivalently, of the Duffie–Singleton defaultable claim. In AHJ’s funding parametrisation the constant-parameter closed form is $f = \frac{\kappa_{\text{AHJ}} - \iota}{\kappa_{\text{AHJ}}} X$ (at $r_a = r_b$), which is *linear* in ι and has no pole at all: there, raising ι drains the effective event intensity rather than discounting a fixed one, and the price reaches zero at $\iota = \kappa_{\text{AHJ}}$ instead of diverging. The two parametrisations agree on the substantive point — $\iota = 0$ is a finite, regular, no-arbitrage price — and the pole-displacement statement should be read in the event coordinates where it lives. What is common to both, and what the conventional consol lacks, is a terminating event: it is the event, not the coordinates, that makes the riba-free point regular.

Corollary 1 (Smoothness, not singularity, at the riba-free point). *For $\kappa > 0$, f is differentiable to all orders at $\iota = 0$, with $f'(0) = -\mathbb{E}[\theta X_\theta]$ finite; the riba-free price is a regular value approached smoothly from the interest economy, not a singular limit. By contrast $V'(r) = -c/r^2 \rightarrow -\infty$ as $r \downarrow 0$.*

Remark 3 (The margin κ is measured, not assumed). The regularity of Theorem 2 hinges on a single inequality, $\kappa > 0$, which is the margin between the riba-free point and the operator’s pole. That inequality is not a modelling convenience: the convergence intensity is estimated strictly positive in every domain in which it has been measured [Ahmed & Bhuyan, 2026, P15] — $\hat{\kappa} = s/(1-\delta) \approx 0.02$ – 0.2 per year from sukuk and CDS spreads, of order 10^2 – 10^3 per year from the mean-reversion half-life of liquid perpetual bases (hours), and a positive realised claim-arrival rate from insurance-loss records. The pole at $\iota = -\kappa$ therefore sits a strictly positive, empirically bounded distance below zero, and the riba-free price is regular by measurement, not by assumption.

4 Orthogonality of the two parameters

Proposition 1 (Independent axes). *Write the constant-parameter price in closed form, $f(\kappa, \iota) = \frac{\kappa}{\kappa + \iota} \mathbb{E}[X_\theta]$ for a payoff independent of θ . Then κ and ι occupy distinct structural positions: ι enters only the discount factor and κ governs the convergence law, so*

$$\left. \frac{\partial f}{\partial \iota} \right|_{\iota=0} = -\frac{1}{\kappa} \mathbb{E}[X_\theta] \quad (\text{finite}), \quad f(\kappa, 0) = \mathbb{E}[X_\theta] \quad \text{is independent of } \iota,$$

and ι may be set to zero without perturbing the κ -channel. The two are not reparametrisations of a single quantity.

This is the precise sense in which κ “is not r ”: they are different arguments of one pricing function, entering in different slots, with ι removable and κ load-bearing.

5 The Duffie–Singleton isomorphism

Theorem 3 (Formal isomorphism). *Let $\mathcal{A}$ be the perpetual-futures pricing problem $(\kappa, \iota, \theta, X_\theta)$ of Definition 1, and $\mathcal{B}$ the Duffie & Singleton [1999] defaultable-claim problem $(\lambda, r, \tau, R_\tau)$ with default time τ a Cox stopping time of intensity λ and recovery R_τ , priced as $P = \mathbb{E}^\mathbb{Q}[e^{-r\tau} R_\tau]$. The map*

$$\Phi : \kappa \leftrightarrow \lambda, \quad \iota \leftrightarrow r, \quad \theta \leftrightarrow \tau, \quad X_\theta \leftrightarrow R_\tau$$

is a bijection of the two representations under which the no-arbitrage values coincide, $f = P$. Under Φ , $\iota = 0 \leftrightarrow r = 0$, and Theorem 2 transfers verbatim: the defaultable claim at $r = 0$ is regular for $\lambda > 0$. (In AHJ’s funding parametrisation the same map reads $\kappa_{\text{AHJ}} - \iota \leftrightarrow \lambda$, as stated in the credit-equivalence paper [Ahmed & Bhuyan, 2026, P2]; the two coincide at the riba-free point, where both reduce to the headline identification $\kappa \leftrightarrow \lambda$.)

Proof. Both values are expectations under $\mathbb{Q}$ of a discount times a payoff evaluated at the first jump of a Cox process; substituting the four identifications carries (1) to P term by term. Bijectivity is immediate as the substitution is invertible. $\square$

The economic reading: credit risk (λ) and time preference (r) are independent axes that sixty years of reduced-form credit pricing fused into one discounted-hazard expression out of convention. Theorem 3 un-fuses them; Theorem 2 shows the interest axis can be set to zero without collapsing the price.

6 Stochastic intensity and change of measure

We now let κ be stochastic, completing the picture under both measures. Take the Cox–Ingersoll–Ross intensity under the physical measure $\mathbb{P}$,

$$d\kappa_t = \alpha(\bar{\kappa} - \kappa_t) dt + \nu\sqrt{\kappa_t} dW_t^{\mathbb{P}}, \quad 2\alpha\bar{\kappa} \geq \nu^2, \quad (3)$$

so that $\kappa > 0$ a.s. (Feller), keeping Theorem 2 in force pathwise.

Theorem 4 (Q-dynamics, Feller preserved). *Under the completely-affine market price of risk $\Lambda_t = (\lambda/\nu)\sqrt{\kappa_t}$ and the associated Girsanov change to $\mathbb{Q}$, κ remains CIR,*

$$d\kappa_t = \alpha^*(\bar{\kappa}^* - \kappa_t) dt + \nu\sqrt{\kappa_t} dW_t^{\mathbb{Q}}, \quad \alpha^* = \alpha + \lambda, \quad \bar{\kappa}^* = \frac{\alpha\bar{\kappa}}{\alpha + \lambda}, \quad (4)$$

provided $\alpha + \lambda > 0$ (the market price of risk does not overwhelm mean reversion, so that $\alpha^, \bar{\kappa}^* > 0$); then with the same ν , $2\alpha^*\bar{\kappa}^* = 2\alpha\bar{\kappa} \geq \nu^2$ and the Feller condition is preserved.*

Proof. Substitute $dW^{\mathbb{P}} = dW^{\mathbb{Q}} - \Lambda dt$ into (3): the drift becomes $\alpha\bar{\kappa} - (\alpha + \lambda)\kappa$, giving $\alpha^*, \bar{\kappa}^*$ as stated; the product $2\alpha^*\bar{\kappa}^* = 2(\alpha + \lambda)\alpha\bar{\kappa}/(\alpha + \lambda) = 2\alpha\bar{\kappa}$ is invariant. $\square$

7 The affine term structure

The price at t of a unit claim contingent on survival to T (zero recovery), or equivalently the κ -discount factor, is affine in κ .

Proposition 2 (Affine survival price). *Under (4), $P(t, T) = \mathbb{E}^{\mathbb{Q}}[e^{-\int_t^T \kappa_s ds} \mid \mathcal{F}_t] = A(\tau) e^{-B(\tau)\kappa_t}$, $\tau = T - t$, where A, B solve the Riccati system*

$$B'(\tau) = 1 - \alpha^* B(\tau) - \frac{1}{2}\nu^2 B(\tau)^2, \quad \frac{A'(\tau)}{A(\tau)} = -\alpha^* \bar{\kappa}^* B(\tau), \quad (5)$$

with $B(0) = 0$, $A(0) = 1$, and closed form, writing $\gamma = \sqrt{(\alpha^)^2 + 2\nu^2}$,*

$$B(\tau) = \frac{2(e^{\gamma\tau} - 1)}{(\gamma + \alpha^*)(e^{\gamma\tau} - 1) + 2\gamma}, \quad A(\tau) = \left[\frac{2\gamma e^{(\alpha^* + \gamma)\tau/2}}{(\gamma + \alpha^*)(e^{\gamma\tau} - 1) + 2\gamma} \right]^{2\alpha^*\bar{\kappa}^*/\nu^2}.$$

Proof. Apply the Feynman–Kac theorem to $P(t, T)$; the PDE $\partial_t P + \alpha^*(\bar{\kappa}^* - \kappa)\partial_{\kappa} P + \frac{1}{2}\nu^2 \kappa \partial_{\kappa\kappa} P - \kappa P = 0$ with the affine ansatz separates into (5); the closed form is Cox et al. [1985]. $\square$

The κ -yield curve is $y(\tau) = -\tau^{-1} \log P(t, T)$, upward-sloping for κ_t below its long-run risk-neutral level. This is the continuous-time object behind the empirically observed, monotone, arbitrage-free κ -surface.

Corollary 2 (No direct rate channel). *Because the price (1) at $\iota = 0$ carries no discount factor in r , a κ -priced claim has no direct interest-rate sensitivity within the pricing formula: $\partial f / \partial r = 0$ at every maturity. Its formula sensitivity is to the convergence intensity alone — an event duration $\partial \log P / \partial \kappa = -(1 - \delta)T$ from the affine form — not to the price of time.*

This is the practical content of separation, stated with its scope. It is a property of the pricing formula, not an observed return: no κ -priced instrument traded through the episode that illustrates it. In the 2022 rate shock the ten-year US Treasury yield rose from $\sim 1\%$ to $\sim 4\%$; a conventional, interest-priced bond lost value in proportion to its duration — from about -4.5% at one year to about -45% at thirty years (a single policy decision revaluing the long end by nearly half) — whereas the $\iota = 0$ formula has no rate term to revalue, by Corollary 2. What it retains is the far shallower event duration of the corollary; the realised credit- κ move over the same period was two to three orders of magnitude smaller than the rate move, so the residual formula loss was 12 to 35 times smaller across the curve (Paper 15, the empirical companion, which also states the equilibrium caveat: in a mixed economy a traded κ -claim competes with r -bearing assets, so its market-clearing intensity κ_Q responds to the rate environment even though the formula's direct rate channel is gone; only a closed κ -world removes the indirect channel as well). The separation of interest from the price is real at the level where pricing formulas operate: it is the difference between carrying a built-in duration gradient and not.

8 Corollary: the riba question as a single parameter

Let $\kappa_P = \mathbb{E}_\infty^\mathbb{P}[\kappa] = \bar{\kappa}$ and $\kappa_Q = \mathbb{E}_\infty^\mathbb{Q}[\kappa] = \bar{\kappa}^*$ be the steady-state physical and risk-neutral intensities, and $\Pi = 1 - \kappa_P/\kappa_Q$ the risk-premium share.

Corollary 3 (Reduction). *Under Theorem 4 with $-\alpha < \lambda < 0$ (so $\alpha^* > 0$ and $\kappa_Q > \kappa_P$), $\Pi = 1 - \bar{\kappa}/\bar{\kappa}^* = |\lambda|/\alpha \in (0, 1)$. After Theorem 2 has removed the predetermined discount at $\iota = 0$, the entire wedge between the priced and the realised intensity is governed by the market price of κ -risk λ , and $\Pi = |\lambda|/\alpha$. The jurisprudential question of whether the residual return is riba in substance is therefore equivalent to a statement about the sign and magnitude of one structural parameter.*

Proof. $\bar{\kappa}^* = \alpha\bar{\kappa}/(\alpha + \lambda)$ gives $\bar{\kappa}/\bar{\kappa}^* = (\alpha + \lambda)/\alpha$, whence $\Pi = 1 - (\alpha + \lambda)/\alpha = -\lambda/\alpha = |\lambda|/\alpha$ for $\lambda < 0$. $\square$

Companion empirical work [Ahmed & Bhuyan, 2026, P15,] estimates Π directly— $\sim 18\%$ for takaful (small $|\lambda|/\alpha$, the priced intensity tracks realised loss) and the great majority for sovereign credit (large $|\lambda|/\alpha$)—so the analytic reduction here and the measured decomposition there describe one quantity.

9 Application: the money base as a stochastic process

For a κ -priced monetary base $M_t = \sum_i N_i P_i(t)$ with $P_i = \mathcal{P}_i e^{-\beta_i \kappa_t^{(i)} T_i}$, $\beta_i = (1 - \delta_i)$, Proposition 2 and Itô's lemma applied to (4) give

$$dM_t = \underbrace{\sum_i N_i P_i \left[\beta_i T_i \alpha^* (\kappa_t^{(i)} - \bar{\kappa}_i^*) + \frac{1}{2} \beta_i^2 T_i^2 \nu_i^2 \kappa_t^{(i)} \right] dt}_{\text{revaluation drift}} - \underbrace{\sum_i N_i P_i \beta_i T_i \nu_i \sqrt{\kappa_t^{(i)}} dW_t^{(i)}}_{\text{diffusion}} + dM_t^{\text{iss}}. \quad (6)$$

The base contracts when intensities rise ($\partial M / \partial \kappa^{(i)} = -N_i P_i \beta_i T_i < 0$): the countercyclical, debt-free money supply of the κ -Standard, here as the drift of (6).

10 Conclusion

The Separation Theorem is the assertion that interest is a removable parameter of no-arbitrage pricing. We have made it precise. Interest enters the perpetual price as a multiplicative discount overlay on a value carried by a Cox stopping time (Theorem 1); for a strictly positive convergence intensity the operator is analytic at the riba-free point, with its pole displaced to $\iota = -\kappa$, so removing interest yields a finite, regular price where the conventional perpetual diverges (Theorem 2); the two parameters are independent axes (Proposition 1); the construction is isomorphic to the Duffie–Singleton claim (Theorem 3); the intensity carries through a measure change that preserves non-negativity (Theorem 4) and an affine term structure (Proposition 2); and the residual, after interest is gone, is governed by a single risk-price parameter whose magnitude the data measures (Corollary 3). The displacement of the pole off the origin by $\kappa > 0$ is the whole of it: r is the price of time, removable; κ is the price of a real event, load-bearing.

References

- Ackerer, D., Hugonnier, J., & Jermann, U. (2025). Perpetual Futures Pricing. *Mathematical Finance*, DOI 10.1111/mafi.70018.
- Cox, J. C., Ingersoll, J. E., & Ross, S. A. (1985). A Theory of the Term Structure of Interest Rates. *Econometrica* 53(2), 385–407.
- Duffie, D., & Singleton, K. J. (1999). Modeling Term Structures of Defaultable Bonds. *Review of Financial Studies* 12(4), 687–720.
- Lando, D. (1998). On Cox Processes and Credit Risky Securities. *Review of Derivatives Research* 2, 99–120.
- Ahmed, S., & Bhuyan, R. (2026). From Perpetual Contracts to Islamic Credit: The Random Stopping Time Equivalence. SSRN 6322858.
- Ahmed, S., & Bhuyan, R. (2026). The κ -Framework Meets the Data.
- Ahmed, S., & Bhuyan, R. (2026). The Riba Question, Reduced: A Risk-Premium Decomposition of the Convergence Intensity.