# The Safe Asset without the Risk-Free Rate: Spanning, Heterogeneity, and the $\kappa$-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 **$\kappa$-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 $\rho_x$ — and show it recovers **80–98.5%** of the spanning gap at residual floor volatilities of 0.5–1%. *Third*, we ask what $\rho_x$ 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 $(\gamma_A, \gamma_B)$ 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-$\gamma$ agent (Section 3). *Second*, we define the $\kappa$-ladder — $N$ event-intensity claims with per-claim volatility $\sigma_e$ and cross-domain correlation $\rho_x$, 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 $\sigma_{\text{agg}}$. 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 $\kappa$-ladder paying $1 + \varepsilon$, $\varepsilon$ independent of $Y$ with volatility $\sigma_{\text{floor}}$ — 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 $\gamma$.

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

| $\gamma_A$ | $\gamma_B$ | $\sigma_{\text{agg}}$ | 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 $\kappa$-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 $\kappa$-ladder

It does, and it requires no issuer's promise. A *$\kappa$-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 $\sigma_e$ and average cross-domain correlation $\rho_x$,

$$\sigma_p^2(N) \;=\; \frac{\sigma_e^2}{N} + \Big(1 - \frac{1}{N}\Big)\rho_x\sigma_e^2, \qquad \sigma_p(\infty) = \sqrt{\rho_x}\,\sigma_e.$$

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:

| $\sigma_e$ | $\rho_x$ | $\sigma_\pi$=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 $\kappa$-ladder whose real volatility beats a nominal government bond's (inflation volatility $\sigma_\pi$). 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

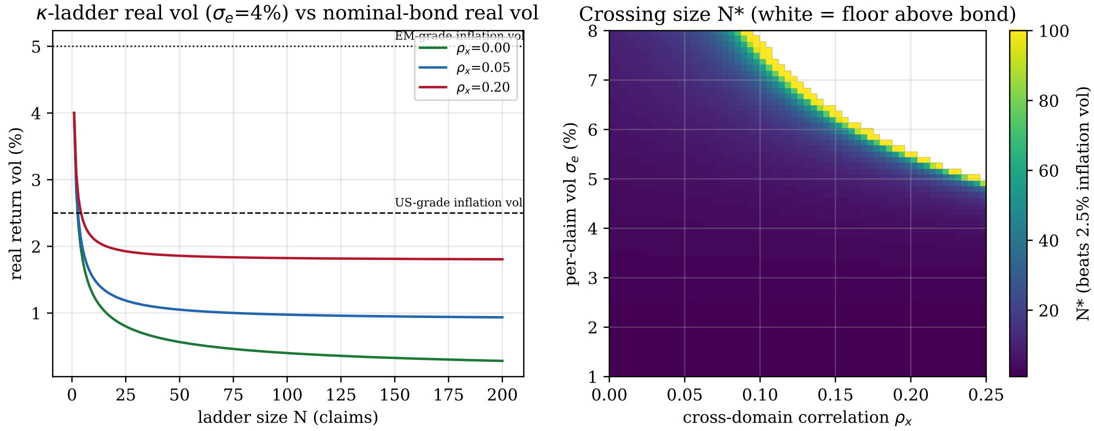


Figure 1: Left: $\kappa$-ladder real volatility against ladder size for three cross-domain correlations ($\sigma_e = 4\%$), with nominal-bond real volatility benchmarks. Right: crossing size $N^*$ over the $(\rho_x, \sigma_e)$ plane against a 2.5% inflation-volatility bond.

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 $\rho_x$? 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 $\rho_w$ and cross-sleeve correlation $\rho_x$, 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 $\rho_x$ 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 ($263 \to 701$ 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 $\rho_x$ 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 ($\rho_x \to 0$) and the conventional failure mode ($\rho \to 1$): removing $\iota$ 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 $\sigma_e$ and $\rho_x$ 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 $\kappa$-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 $\kappa$-Standard. SSRN Working Paper 6890259.

Ahmed, S., & Bhuyan, R. (2026b). Actuarial properties of the $\kappa$-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.