

Quantifying *Gharar* in Islamic Digital Finance

Tariqullah Khan

Chief Executive Officer & Principal Economist, Ventureethica, Toronto
Formerly Professor, Hamad Bin Khalifa University, Qatar; IZU, Türkiye; INCEIF, Malaysia
Division Chief, Islamic Banking and Finance, IsDBI
President, International Association for Islamic Economics
Email: tariqullah.khan@ventureethica.com

ABSTRACT

The new economy is substantively digital and trust based. Prohibition of *Gharar* (contractual uncertainty) in the traditionally economy serves as a guardrail to enhance trust in the markets. This paper shows that *Gharar* in the digital economy has additional features and develops a quantitative framework for assessing *Gharar* in Islamic digital finance. Two structural axes - Specificifiability (S, the epistemic determinacy of a contract's terms) and Deliverability (D, the ontological certainty of its subject matter) - are derived from the classical taxonomy of *Gharar* sources and combined into a coordinate space $(S, D) \in [-1, +1]^2$. Admissibility is determined by axis signs alone: a contract is admissible if and only if $(S, D) \in Q1$ ($S > 0$ and $D > 0$); contracts in Q2, Q3, and Q4 are categorically batil and instantiate the classical taxonomy without remainder. Within Q1, quality is continuous, stratified by $\min(S, D)$ into Robust, Standard, and Marginal bands and measured by the d-gap from the optimum $v^* = (+1, +1)$. A maqasid-grounded Tawhib constraint $\tau \in [0, 1]$ operates as an independent floor alongside the (S, D) admissibility test. The framework is then extended to digital contracts, which exhibit five structurally novel properties - entanglement of the two axes, composability across protocols, governance mutability, ordering risk through adversarial transaction sequencing, and time-dependence through continuous attestation - that the classical taxonomy was never required to address. The paper formalises these properties, decomposes each axis into three sub-dimensions specific to digital contracts, generalises the framework from points to trajectories, specialises the τ -constraint to digital-asset maqasid considerations, and proposes a Substitution-Transformation-Offset (STO) remediation taxonomy with digital instruments. Five case studies illustrate the framework across structurally distinct digital instruments. The paper concludes by identifying the absence of a dedicated multilateral Shari'ah Standard on cryptoassets as itself a finding and positioning the framework as a candidate scaffold for future standards-development work.

ARTICLE HISTORY

Received: May 6, 2026
Accepted: May 7, 2026
Published: May 08, 2026

KEYWORDS

Gharar; Islamic digital finance; smart contracts; stablecoins; DeFi; oracle risk; composability; governance mutability; AAOIFI Shari'ah Standards, Maqasid al-Shari'ah.

JEL CODES

G29; G23; K12; O33; P40; Z12.

HOW TO CITE

Khan, T. (2026). Quantifying *Gharar* in Islamic Digital Finance. *Journal of Economics, Law and Society*, 3(1), 5-29. <https://doi.org/10.70009/jels.2026.3.1.1>

1. INTRODUCTION

The classical jurists who developed the prohibition of *Gharar* were reasoning about contracts they could touch. A buyer could lift a sealed pomegranate, weigh a bundle of wheat, inspect the leather of a coat. The 'urf-tolerated residual uncertainty of mediaeval commerce was bounded by the tactile properties of goods and by the social embeddedness of the parties, and *Gharar* prohibition functioned alongside reputation, kinship, and repeated interaction as one mechanism among several that kept commerce inside the zone where trust was possible.

Digital finance dissolves these presuppositions. The "object" of a tokenised real-world asset is a cryptographic claim whose underlying referent is geographically remote and institutionally

mediated. The “delivery” of a perpetual swap is a cash flow generated by an algorithm whose state depends on the contract’s own trading. The “specification” of a DeFi yield strategy is a Solidity program that may be perfectly published yet effectively opaque to all but a small number of expert auditors, and that may change next week through a governance vote in which neither counterparty participated. The contemporary literature on Islamic digital finance has, in consequence, oscillated between two unsatisfactory poles: binary fatāwā that declare an entire instrument class permitted or forbidden, and informal scoring rubrics that gesture at the issues without committing to falsifiable evaluation. Neither approach scales to the heterogeneity, hybridity, and rapid evolution of the cryptoasset economy. Compounding this, no dedicated multilateral Shari’ah Standard on cryptoassets currently exists; AAOIFI’s issued Standards reach Shari’ah Standard No. 62 on *Sukuk* (exposure draft 2024), and Shari’ah-supervisory practice for digital instruments proceeds by analogical extension of pre-digital Standards (notably No. 18 on *Qabd*, No. 21 on Financial Papers, No. 17 on Investment *Sukuk*, and No. 57 on Gold).

This paper develops the quantitative middle ground. It does so in two stages. The first stage (Sections 2–3) derives a self-contained vector-based framework — the (S, D, τ) framework — from classical *fiqh* sources and contemporary *maqasid* scholarship, applicable to ordinary commercial contracts. The second stage (Sections 4–11) extends that framework to digital contracts, identifying five structurally novel properties that digital instruments exhibit and that the classical taxonomy was never required to address: entanglement of the epistemic and ontological axes, composability across independent protocols, mutability through governance mechanisms, ordering risk through adversarial transaction sequencing, and time-dependence through continuous attestation. Section 12 applies the extended framework to five structurally distinct digital instruments. Section 13 returns to the standards-development gap, positioning the framework as a candidate scaffold for future Shari’ah-supervisory work. Section 14 concludes.

The framework is offered with three commitments. First, doctrinal grounding: each component is anchored to classical *fiqh* sources (al-Sarakhsi, al-Kasani, Ibn Qudamah, Ibn Rushd) and where appropriate to contemporary *maqasid* scholarship (Ibn Ashur, Chapra, Auda). Second, operational specificity: rubrics, thresholds, and aggregation rules are stated in falsifiable form. Third, scientific discipline: Section 11 commits the framework to disconfirmation criteria, in line with the methodology of Dynamic Prescriptive Islamic Economics. The paper does not claim to declare a new field; it offers the foundations of a research programme.

Four contributions distinguish the work. First, it derives a two-axis structure of *Gharar* directly from classical jurisprudence rather than imposing it as a modelling convenience. Second, it formalises the structure as a coordinate system with binary admissibility and continuous quality, separating doctrinal admissibility from gradient quality without conflating them. Third, it identifies five structural transformations that digital contracts impose on the classical architecture and develops the framework extensions those transformations require. Fourth, it specifies an operational protocol — rubrics, scoring procedure, panel adjudication, falsifiability conditions — that is detailed enough to support pilot implementation and standards-development work.

2. LITERATURE REVIEW AND THE STANDARDS GAP

2.1. Classical and modern *Gharar* literature

The classical *fiqh* treatment decomposes *Gharar* into five recognised sources: *al-jahala* (ignorance of attributes), *al-qadr* (indeterminate quantity), *al-ajal* (indeterminate time), *al-wujud* (non-existence of subject matter), and *al-husul* (uncertainty of attainment or delivery). Al-Sarakhsi’s *Mabsut*, al-Kasani’s *Bada’i’ al-Sana’i’*, Ibn Qudamah’s *Mughni*, and Ibn Rushd’s *Bidayat al-Mujtahid* develop these

sources at length. Two structural observations from the classical corpus drive the framework of this paper. First, the five sources cluster on two distinct axes: *al-jahala*, *al-qadr*, and *al-ajal* concern the epistemic determinacy of the contract's terms; *al-wujud* and *al-husul* concern the ontological status of the subject matter. Second, the classical jurists treated these as independent defects: a buyer could be fully informed about an existing object (epistemic clarity, ontological presence), or fully informed about a non-existent object (epistemic clarity, ontological absence), or ignorant about an existing object (epistemic deficit, ontological presence). The two-axis structure is therefore not imposed on the classical material but read out of it.

The modern quantification literature has developed each axis separately. Al-Saati (2003) treats *Gharar* in Knightian-uncertainty terms (epistemic). Al-Suwailem (2006) develops a game-theoretic structure (epistemic with strategic interaction). El-Gamal (2006) gives a transaction-cost interpretation (mostly epistemic). Obaidullah (2002) proposes multi-criterion derivative tests that span both axes but does not integrate them into a unified coordinate. Kamali (2000) examines customary tolerance (*urf*-bounded *Gharar yasir*) without quantifying it. The literature thus contains the two-axis structure in fragments without an integrating coordinate framework.

2.2. Contemporary crypto-fiqh literature and the standards gap

Three strands of contemporary literature engage Islamic digital finance. The first is the fatwa literature on cryptocurrency permissibility. Mufti Taqi Usmani's position — that conventional cryptocurrencies fail the requirements of *mal* and produce excessive *Gharar* — has been influential and is reflected in jurisdictional rulings including those of the Egyptian Dar al-Ifta' and the Turkish Diyanet. Mufti Faraz Adam (2017, 2019) develops a more granular position distinguishing payment, utility, and security tokens. The Shariyah Review Bureau and Amanie Advisors produce product-level reviews without a unifying coordinate framework. Aishath Muneeza and co-authors (Muneeza & Mustapha 2019; Muneeza et al. 2021) examine tokenisation, smart *sukuk*, and Islamic DeFi, with attention to regulatory and Shari'ah-supervisory architecture. Mohamed Cherif El Amri and others examine the *maqasid* implications of decentralised finance.

The second strand is the standards-development landscape. AAOIFI's Shari'ah Standards reach No. 62 on *Sukuk* (exposure draft 2024). The 21st AAOIFI Shari'ah Council Conference (May 2023, Manama) included a roundtable on digital assets in which the Secretary-General of the International Islamic Fiqh Academy participated, and AAOIFI conferences have repeatedly addressed cryptoassets. But — and this is itself a finding for the present paper — no dedicated AAOIFI Shari'ah Standard on cryptoassets has been adopted. Shari'ah-supervisory practice for digital instruments therefore proceeds by analogical extension of pre-digital Standards: Standard No. 18 (Qabd / possession) for tokenised representations of physical assets, No. 21 (Financial Papers) for security tokens, No. 17 (Investment *Sukuk*) for tokenised *sukuk*, No. 57 (Gold) for precious-metal-backed tokens, and the forthcoming No. 62 for on-chain *sukuk*. Recent national frameworks — the Bahrain Central Bank's 2025 stablecoin issuance module and the UAE's Federal Decree-Law No. 6 of 2025 on digital currency — reference AAOIFI alignment generically without specifying a cryptoassets Standard, because there is none to specify. The Higher Shari'ah Authority of the UAE adopts AAOIFI Standards collectively for licensed Islamic institutions and applies them by analogy to digital products.

The third strand is the technical literature on smart-contract verification, oracle reliability, and DeFi risk that has developed largely outside the Islamic finance discourse — work on oracle problems by Juels and others, formal-verification work by CertiK, Runtime Verification, and the Move and Certora prover teams, and empirical literature on DeFi exploits (Werner et al. 2022; Zhou et al. 2023). This literature is empirically detailed but doctrinally agnostic.

2.3. The gap

No existing work offers what the present paper attempts: a quantitative, falsifiable, classically-grounded framework with a continuous quality metric within doctrinally-defined admissibility regions, applicable to digital instruments. The fatwā literature is doctrinally rich but operationally binary. The standards landscape is operationally structured but classificatory rather than coordinate, and lacks a dedicated cryptoassets Standard. The technical literature is empirically detailed but doctrinally agnostic. The framework developed here addresses all three deficits, and Section 13 returns to the standards gap as an opportunity for the framework's prospective integration.

3. THE (S, D, τ) FRAMEWORK

This section develops the framework in self-contained form. The reader is not assumed to have access to prior work in the research programme; the doctrinal derivation, the geometric structure, and the operational thresholds are stated here in full. Sections 4 onward then extend the framework to digital contracts.

3.1. The two axes: doctrinal derivation

The classical taxonomy of *Gharar* sources clusters on two distinct structural axes. *Al-jahala* (ignorance of the attributes of the object), *al-qadr* (indeterminacy of quantity), and *al-qjal* (indeterminacy of time) all concern the buyer's epistemic position relative to the contract: whether the buyer knows what is being agreed to. Al-Sarakhsi treats these as defects of *bayan*, the obligation of clarity that he derives from the prohibition of *bay' al-mulamasah* (the blind sale by touch). Al-Kasani develops the same cluster under *jahalah fahishah* (egregious ignorance) and treats it as a category of *bayan* deficit. The structural commonality is informational: the contract fails because the buyer cannot, at the moment of contracting, determine what the contract obliges.

Al-wujud (non-existence of the subject matter) and *al-husul* (uncertainty of attainment or delivery) cluster on a different axis. Ibn Qudamah's treatment of *bay' al-ma'dum* (the sale of what does not exist) and Ibn Rushd's analysis of *bay' al-Gharar* concern not whether the buyer knows but whether the object can be delivered. The structural commonality here is ontological: the contract fails because the subject matter is not in a state to be transferred. The classical jurists' acceptance of *salam* (forward sale of fungibles) and *istisna'* (manufacturing contract) — both of which involve future objects — establishes that ontological absence is not categorically prohibited; what is prohibited is ontological absence combined with the absence of structural mechanisms (specification at *salam*, milestone delivery at *istisna'*) that secure deliverability.

We therefore name the two axes Specifiability (S) for the epistemic dimension and Deliverability (D) for the ontological dimension. The names follow the underlying jurisprudential function: S measures whether the contract's terms are specifiable at the moment of contracting, D measures whether the contract's subject matter is deliverable as obligated. Each axis takes values in $[-1, +1]$, with +1 corresponding to fully specifiable (or fully deliverable) and -1 corresponding to wholly unspecifiable (or wholly undeliverable). Zero is the boundary of admissibility on each axis.

3.2. Non-substitutability

The two axes are non-substitutable: a strong position on one axis does not offset a weak position on the other. The doctrinal grounding for non-substitutability is the structural separation of *daf' al-mafasid* (harm-prevention) from *jalb al-masalih* (benefit-acquisition) in the *maqasid* framework. Specifiability addresses one structural harm — the buyer not knowing what they are agreeing to — and Deliverability addresses a different structural harm — the buyer not receiving what they were

promised. These are categorically different defects, and the classical jurists' treatment of them as independent grounds for invalidation reflects the recognition that one cannot substitute for the other. A perfect specification of a non-existent object is *batil* (Q4 in the geometry below); perfect knowledge of a non-existent referent does not bring the referent into being. Conversely, an existing object whose terms are unspecifiable is *batil* (Q2); the object's existence does not cure the contract's informational void.

Non-substitutability has an operational consequence that distinguishes the framework from utility-based alternatives: there is no indifference curve between S and D. The admissibility boundary is a right angle at the axes, not a smooth trade-off frontier. Section 3.4 develops the geometry.

The translation from doctrine to coordinate is not arbitrary. The Qur'anic principle of *mizan* (balance, established in *Surat al-Rahman* 55:7–9) requires equilibrium across distinct dimensions of value and is explicit that one dimension cannot offset deficit on another — the prohibition is against *tukhsiru al-mizan*, the diminishing of the balance, in any of its dimensions. *Wasatiyyah*, the principle of balanced positioning developed in *Surat al-Baqarah* 2:143, implies dual-structure evaluation rather than scalar aggregation. The two-axis coordinate space is the natural geometric expression of these principles: the framework does not impose structure on the doctrine but makes explicit the structure already implicit within it.

3.3. Sign-based admissibility

A contract C is represented as a point $(S(C), D(C)) \in [-1, +1]^2$. Permissibility under the *Gharar* prohibition is determined by the conjunction:

$$C \text{ is admissible iff } S(C) > 0 \wedge D(C) > 0$$

The four quadrants of the coordinate space carry distinct juristic interpretations: Q1 ($S > 0, D > 0$) is admissible; Q2 ($S < 0, D > 0$) instantiates the *al-jahala* cluster (epistemic deficit on existing objects); Q3 ($S < 0, D < 0$) is doubly defective; Q4 ($S > 0, D < 0$) instantiates the *al-wujud / al-husul* cluster (ontological deficit despite specification). Contracts in Q2, Q3, and Q4 are categorically *batil* and instantiate the classical taxonomy of *Gharar* sources without remainder.

The admissibility test is binary and lexicographic. A contract whose S-score is +0.95 but whose D-score is -0.01 is not admissible — non-substitutability rules out trade-off. The framework distinguishes admissibility (sign-based, binary) sharply from quality (continuous, gradient), addressed in §3.5. This distinction matters because the contemporary literature has tended to conflate the two: papers that score *Gharar* on a continuous scale and report contracts as “more compliant” or “less compliant” implicitly treat axis values as substitutable, which the doctrinal structure does not permit.

3.4. Geometric representation

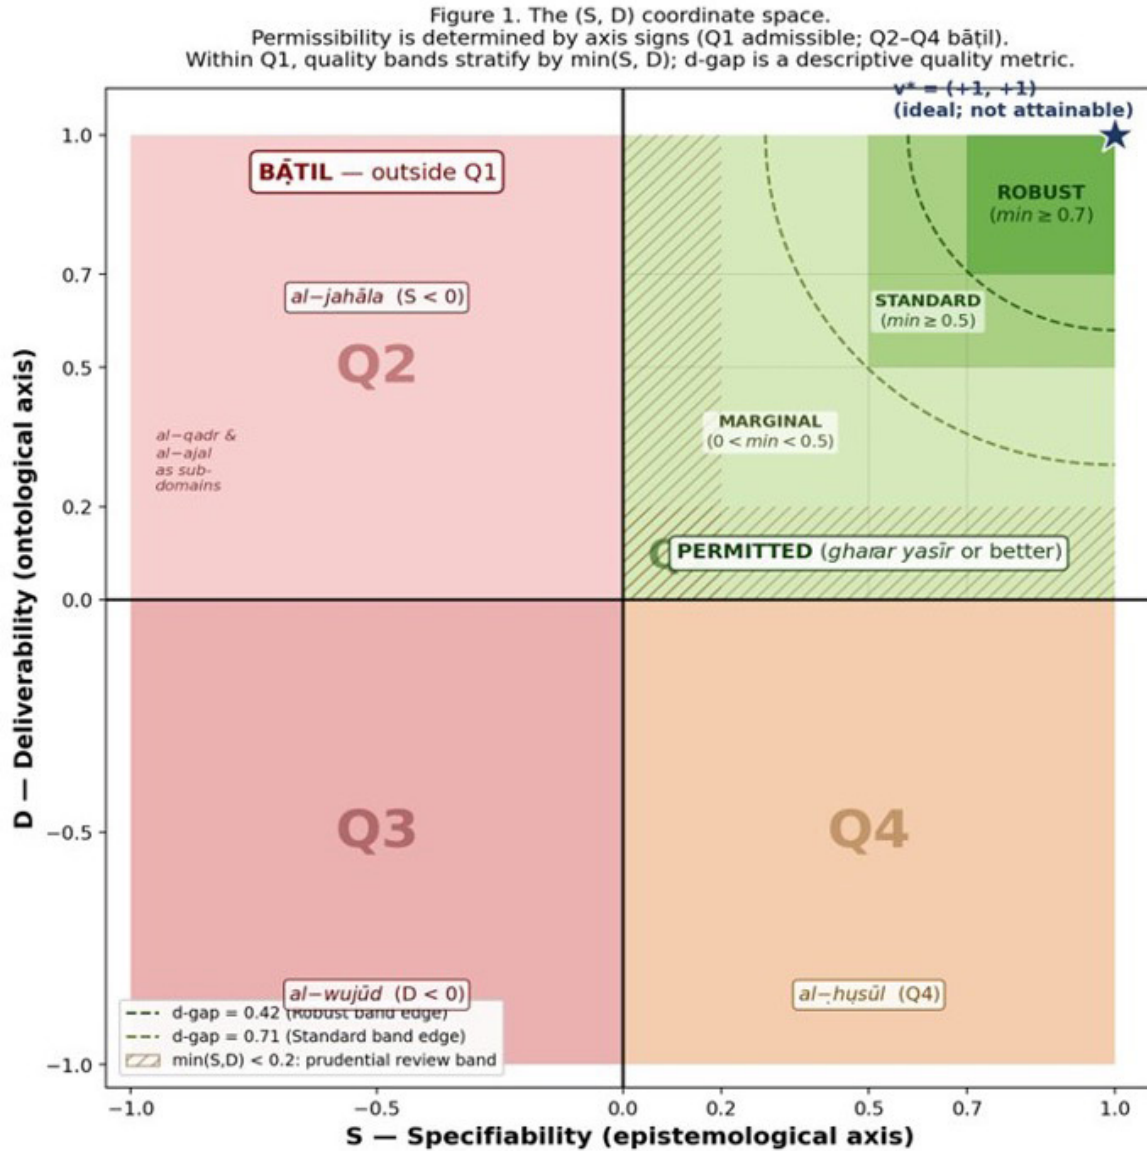
Figure 1 presents the (S, D) coordinate space with its four quadrants, the optimal vector $v^* = (+1, +1)$, the quality bands within Q1, and the prudential review corridor along the axes.

3.5. Quality within Q1: the d-gap and the min-constraint

Within Q1 — the admissible region — quality is continuous. Two metrics jointly characterise quality. The first is the d-gap, the Euclidean distance from the optimum:

$$\text{d-gap}(C) = \sqrt{[1 - S(C)]^2 + [1 - D(C)]^2}$$

Figure 1: The (S, D) coordinate space.



Note(s): Permissibility is determined by axis signs (Q1 admissible; Q2-Q4 *batil*). Within Q1, quality bands stratify by $\min(S, D)$; the d-gap measures distance from v^* and operates as a descriptive quality metric, not an admissibility threshold. The hatched corridor at $\min(S, D) < 0.2$ marks the prudential review band requiring *sadd al-dhara'i* procedural caution.

The d-gap is a descriptive quality metric, not an admissibility threshold. A contract with smaller d-gap is closer to the ideal of full specification and full deliverability. The second metric is the binding axis:

$$m(C) = \min(S(C), D(C))$$

$m(C)$ captures the operative consequence of non-substitutability: a contract's effective quality is bounded by its weaker axis. Three quality bands stratify Q1:

- **Robust:** $m(C) \geq 0.7$ and $d\text{-gap} \leq 0.42$. Both axes are strongly positive; the contract is well within the classical *Gharar yasir* tolerance.
- **Standard:** $0.5 \leq m(C) < 0.7$ and $d\text{-gap} \leq 0.71$. Both axes are positive with operational margin; ordinary commercial practice typically falls here.
- **Marginal:** $0 < m(C) < 0.5$. Admissible but with limited margin; institutional documentation requirements are heightened.

A prudential near-axis review band is defined at $m(C) < 0.2$: contracts in this band are technically admissible (Q1) but require panel review under the principle of *sadd al-dhara'i* — the closing of pretexts to harm — because their proximity to the boundary makes them susceptible to drift across the admissibility line. The 0.2 threshold corresponds approximately to the classical *Gharar mutawasit* (intermediate *Gharar*) demarcation in al-Suwailem's typology and is consistent with Ibn Rushd's treatment of *Gharar* that “approaches the impermissible without crossing it.”

3.6. The Tayyib constraint τ as an independent maqasid floor

The (S, D) framework addresses *Gharar* — one of the structural prohibitions in *mu'awadat*. It does not address whether a *Gharar*-compliant contract is also Shari'ah-aligned in the broader *maqasid* sense. A contract may be admissible (Q1 with strong quality) and yet fail to advance, or actively undermine, the *maqasid al-Shari'ah* — the higher purposes of preserving *din* (religion), *nafs* (life), *'aql* (intellect), *nasl* (lineage), *mal* (wealth), and (in Ibn Ashur's and contemporary *maqasid* scholarship) *'ird* extended to environmental integrity. The framework therefore introduces a *Tayyib* constraint $\tau \in [0, 1]$ as an independent *maqasid* floor, structurally separated from the (S, D) admissibility test:

C is *Tayyib*-compliant iff $(S, D) \in Q1 \wedge \tau(C) \geq \tau_{\min}$

τ is computed as a weighted average of sub-scores on the six *maqasid* dimensions, with weights calibrated by asset class. The *Tayyib* floor τ_{\min} varies by asset class to reflect that different instruments are expected to advance different *maqasid*; a social-finance instrument is held to a higher τ_{\min} than a general trade-finance instrument because its claimed *maqasid* alignment is more central to its purpose.

The doctrinal grounding for τ 's independent operation is the *maqasid* principle that *daf' al-mafasid* (the prevention of harm, which the *Gharar* prohibition serves) is structurally separate from *jalb al-masalih* (the acquisition of benefit, which the *maqasid* serve). Ibn Ashur's *Maqasid al-Shari'ah al-Islamiyyah* develops this separation explicitly; Chapra's *Morality and Justice in Islamic Economics* applies it to financial instruments; Auda's *Maqasid al-Shari'ah as Philosophy of Islamic Law* extends it to systems-level analysis. The conjunction $(S, D) \in Q1 \wedge \tau \geq \tau_{\min}$ blocks both failure modes: “compliant but extractive” contracts (Q1 with low τ) and “wholesome but ambiguous” contracts (high τ outside Q1). Treating τ as a third axis would, by contrast, permit substitution between admissibility and *maqasid* — allowing high τ to compensate for failed admissibility, which the doctrinal structure does not permit.

3.7. STO remediation

Contracts outside Q1 are not necessarily lost. The framework specifies three families of intervention by which a contract can be moved into admissibility (or, within Q1, to higher quality). The taxonomy follows the *maqasid* logic of *daf'*-then-*jalb*: first remove the structural defect, then enhance the benefit.

Substitution replaces an opaque or unspecifiable component of the contract with a transparent or specifiable equivalent, without changing the contract's structural form. Substitution acts primarily on S. Classical example: replacing an indeterminate sale (*jahalah fahishah*) with a sale by sample (*bay' al-namudhaj*). Digital example: open-sourcing previously proprietary smart-contract code (developed in §9).

Transformation changes the contract's juristic type to a Shari'ah-recognised structure that authorises what the original form did not. Transformation acts primarily on D. Classical example: restructuring a forward delivery as *salam*, with strict specification re-baselining D. Digital example: restructuring a custodial token as a *wakala* arrangement with explicit agency relationship.

Offset introduces third-party safeguards that absorb residual uncertainty without altering the bilateral terms. Offset can act on either axis. Classical example: *takaful* coverage for delivery risk in long-distance trade. Digital example: on-chain *takaful* coverage of smart-contract risk.

The τ -weighted choice rule selects among admissible STO pathways: when multiple interventions can move a contract into the target quality band, choose the pathway that maximises post-remediation τ subject to the (S, D) admissibility constraint. STO is mandatory for non-admissible contracts that are to enter Islamic finance and is recommended within Q1 for quality uplift, with the prudential review band ($m < 0.2$) triggering recommendation even where the contract is technically admissible.

3.8. The framework in one view

The framework rests on four propositions. First, dual irreducible axes: every contract is evaluated along Specificifiability (S) and Deliverability (D), which are non-substitutable. Second, admissibility by signs: a contract is admissible if and only if $(S, D) \in Q1$; contracts in Q2, Q3, and Q4 are categorically *batil* and instantiate the classical taxonomy of *Gharar* sources without remainder. Third, quality by distance: within Q1, quality is continuous, stratified by $\min(S, D)$ into Robust, Standard, and Marginal bands and measured by the d-gap from $v^* = (+1, +1)$, with a prudential near-axis review band at $m < 0.2$. Fourth, independent τ floor: *Tayyib* compliance is the conjunction of Q1 admissibility with $\tau \geq \tau_{\min}$; τ is a *maqasid*-grounded constraint, not a third axis. Sections 4 onward extend each of these propositions to accommodate the structural properties of digital contracts.

4. THE STRUCTURAL NOVELTY OF DIGITAL GHARAR

The framework of Section 3 is the baseline for ordinary commercial contracts. The extension to digital contracts is non-trivial. Five structural properties of digital contracts have no clean correspondent in the classical taxonomy: the framework can be extended to handle them, but each property defeats a presupposition the classical jurists were entitled to make. The extensions developed in Sections 5–11 are required by these properties; the section closes by addressing the substrate-re-labelling objection (§4.6) before the extensions begin.

Table 1 summarises the five properties against the classical assumptions they defeat. The underlying claim: digital contracts do not merely exhibit more uncertainty than their classical counterparts; they transform the structure of uncertainty itself.

4.1. Entanglement of the epistemic and ontological axes

The classical jurists treated knowledge of the object (*al-jahala*) and existence of the object (*al-wujud*) as independent defects. A buyer could be ignorant of the attributes of an existing camel; a buyer could have full knowledge of an unborn calf that did not yet exist. The framework of §3 preserves this independence: the two axes are scored separately, and the combined assessment is their conjunction.

Many digital contracts violate this independence. Consider a custodial stablecoin in which the only evidence of the underlying reserves is the issuer's attestation. The buyer's epistemic position (S) and the underlying object's existence (D) are no longer independent: the only mechanism by which the buyer can come to know about the reserves is the same mechanism that establishes their existence. A failure of attestation transparency is therefore simultaneously a failure of S and of D. The classical analogy would require a sealed box whose existence could only be confirmed by opening it — and which, when opened, might prove empty. Oracle-dependent contracts, wrapped tokens, and synthetic assets all exhibit the same structure.

Table 1: Five structurally novel properties of digital Gharar

| Property | Classical assumption | Digital reality | § |
|-----------------|--|---|-----|
| Entanglement | S and D are informationally independent; al-jahala and al-wujud fail separately. | S and D are co-determined by a single attestation mechanism; one failure crosses both thresholds. | 4.1 |
| Composability | Contracts are atomic; validity is bilateral and self-contained. | Contracts compose through dependency graphs; validity inherits from the worst component reachable from the user. | 4.2 |
| Mutability | Terms are fixed at issuance; subsequent variation requires a new contract. | Terms evolve through governance over the contract's life; the reachable parameter space, not issuance, bounds S. | 4.3 |
| Ordering risk | Execution price equals contract price for spot exchange. | Execution is subject to MEV; effective price routinely diverges from contract price through adversarial sequencing. | 4.4 |
| Time-dependence | Assessment is one-time at issuance; the contract's coordinate is static. | Assessment must follow a continuous attestation feed; the coordinate is a trajectory through (S, D) space. | 4.5 |

Source: Author's own.

Entangled contracts cannot be scored by independent rubric application to the two axes because the axes' values are not informationally separable. Three operational consequences follow. First, the contract's (S, D) coordinate is computed under a single attestation mechanism, with the mechanism's integrity simultaneously bounding both axes. Second, the prudential review band must be widened: a single failure can move both axes across the admissibility threshold simultaneously. Third, STO remediation must address the entangling mechanism rather than the axis values it produces.

4.2. Composability across independent protocols

Classical contracts are atomic. A sale of wheat is a sale of wheat; its validity does not depend on the buyer's simultaneous lease of an ox. The framework of §3 inherits this atomicity. DeFi contracts are composable. A user's position in a yield-aggregator vault may consist of a deposit in the vault, automated routing of that deposit into a lending protocol, borrowing against that deposit through a different lending protocol, swapping the borrowed asset on a decentralised exchange, and depositing the swapped asset in a fourth protocol. Each step is governed by a separate smart contract written by a different team under different security assumptions and governance regimes. The user's effective position is the composition. A failure at any link breaks the chain.

Composability creates a new class of *Gharar* that the classical taxonomy does not name. Each individual protocol may be in Q1 by its own (S, D) score, but the composed position can exhibit aggregate properties that none of the components do. The framework accommodates composability by treating the composed position as a separate object of assessment with its own (S, D) coordinate, computed by the minimum-aggregation rule over the dependency graph (§7.2): a single Q4 component anywhere in the composition collapses the composition to *batil*.

4.3. Mutability through governance mechanisms

Classical contracts are static at issuance. Many digital contracts are mutable: a DeFi lending protocol's parameters can be modified by token-weighted governance votes; a "stablecoin" contract can have its peg mechanism altered by issuer fiat; a smart contract's upgrade pattern means the code the user agreed to at t may not be the code that executes their position at $t + 1$.

Mutability defeats the static-coordinate assumption. The contract's S-score is no longer a property of the contract at issuance but of the contract's entire reachable state space — the set of configurations it can be moved to through legitimate governance action. The framework accommodates this

by treating S as bounded by the worst-case configuration reachable through governance, not the configuration at issuance. The mutability adjustment is a function of three governance properties: the parameter range subject to change, the threshold required to enact change, and the timelock-and-veto structure that buffers the change.

4.4. Ordering risk through adversarial transaction sequencing

In classical commerce, the price the buyer pays is the price the buyer agreed to. On-chain transactions are subject to maximal extractable value (MEV): adversarial actors observe pending transactions in the public mempool and insert their own transactions before, after, or around them to extract value. The user pays the contractually specified price-per-unit, but the effective price after slippage and adversarial extraction can differ substantially.

MEV-induced ordering risk is a form of *Gharar al-husul* operating not on the existence of the object but on the effective terms of delivery. The contract is fully specified, fully deliverable, and yet routinely produces outcomes that diverge materially from the contracted terms in ways that are systematic rather than accidental. The framework accommodates this by adding a sub-dimension D_{order} to the D rubric (§6.2), with public-mempool execution carrying a deduction relative to MEV-protected venues.

4.5. Time-dependence through continuous attestation

Classical contracts settle at a moment. Many digital contracts have a structurally continuous attestation requirement: a reserve-backed stablecoin's D depends on continuing reserve composition; a bridge's D depends on continuing solvency and validator honesty; a DeFi position's S depends on the protocol's current parameters. Time-dependence is structural rather than incidental: the contract's position in (S, D) space is most accurately represented not as a point but as a trajectory — a function from time to coordinate. Section 7 develops the trajectory generalisation.

4.6. Why these properties are structurally novel, not substrate-relabelled

A reader might object that each property has some classical analogue. Entanglement is reminiscent of the sealed-box trade. Composability resembles chained leverage that Ibn Taymiyyah criticised. Mutability is present whenever a contract has waiver provisions. Ordering risk is structurally similar to *najsh* and *hukrah*. Time-dependence is implicit in any deferred contract. The objection is fair as far as it goes, and §5 below treats open-source verifiability as a partial digital analogue of *bayan* precisely because the underlying coordination problem is continuous across traditions.

But the digital instances differ from classical analogues in three respects. First, scale: where classical instances were edge cases, digital instances are routine and often constitutive — DeFi composition is the mode of existence, not an exception. Second, opacity: classical analogues operated under conditions where parties could in principle inspect; digital analogues often operate under conditions where inspection is technically infeasible without specialised tooling. Third, automation: classical analogues required human intermediation; digital analogues execute autonomously, propagating defects at machine speed. The framework therefore extends rather than relabels.

5. OPEN-SOURCE VERIFIABILITY AS DIGITAL BAYAN: A PARTIAL ANALOGY

The classical principle of *bayan* — the obligation of clarity and disclosure that the jurists treated as constitutive of valid *mu'awadat* — has a contemporary digital analogue: open-source verifiability (OSV). Both traditions address the same coordination problem: how to keep markets inside the zone

where trust is preserved when direct inspection is unavailable. *Bayan* is the classical mechanism; OSV is the digital mechanism. The analogy is structural and the underlying coordination problem is continuous across traditions. But OSV is a partial analogue, not a complete substitute, and the framework must distinguish the layers it addresses.

Open-source code, on its own, raises S only insofar as the relevant audience can read and understand the code. For a Solidity smart contract of moderate complexity, this audience is small. The S-axis distinguishes three layers: transparency (the code is published), intelligibility (the code's behaviour is understood by qualified reviewers), and verified intelligibility (the code's behaviour has been formally established by audit, formal verification, or sustained adversarial testing). Only verified intelligibility produces the epistemic state that *bayan* was meant to secure.

Cryptographic verifiability — Merkle proofs of reserves, zero-knowledge proofs of correct execution — likewise does not, on its own, restore D. Such proofs establish consistency between stated claims and on-chain state but not the truth of claims about off-chain reality. A Merkle proof of reserves shows that claimed balances reconcile with a snapshot of asset addresses; it does not show that the addresses contain real assets, that the assets are unencumbered, or that the issuer has not borrowed them temporarily. The D-axis distinguishes technical verifiability (cryptographic consistency) from ontological deliverability (off-chain reality matches on-chain claims). Only their conjunction produces executability.

OSV suffices to restore S only when the contract's logic is fully specified by published code — no off-chain dependency whose behaviour matters. OSV suffices to restore D only when the subject matter is itself on-chain — a token whose existence is exhaustively determined by chain state. Most digital contracts of practical interest satisfy neither condition. The framework therefore treats OSV as one input to the rubrics rather than a self-sufficient restoration of either axis.

6. EXTENDED RUBRICS FOR DIGITAL CONTRACTS

The S and D axes for digital contracts decompose into three sub-dimensions each. The composite axis is the minimum of the sub-dimensions, reflecting non-substitutability at the sub-dimension level: a contract whose code is open but whose governance can arbitrarily replace the code is not specified.

6.1. The Specifiability rubric for digital contracts

S decomposes into S_{code} (transparency), S_{behav} (intelligibility), and S_{gov} (governance immutability). The composite $S = \min(S_{code}, S_{behav}, S_{gov})$.

A specific governance recommendation follows from the S_{gov} rubric: the *Maqasid* Timelock. Any vote that alters the risk profile of the contract — collateral factors, liquidation thresholds, oracle sources, supported assets — should remain in a pending state for a duration $T \geq 2 \cdot T_{exit}$, where T_{exit} is the time required for a typical user to withdraw 90% of their position through normal redemption channels. The doctrinal grounding is *sadd al-dhara'i*: a timelock shorter than user exit time enables governance changes that users cannot escape, which is structurally a form of imposed risk-shifting analogous to the classical prohibition of *bay' al-mukrah* (compelled sale). A protocol whose timelock falls below this threshold should receive an S_{gov} score no higher than the +0.0 anchor of [Table 2](#); the framework recommends 14-day timelock as a typical operating value for protocols with sub-week exit times.

Table 2: Specificifiability (S) rubric for digital contracts

| Score | S_code (transparency) | S_behav (intelligibility) | S_gov (governance immutability) |
|-------|--|--|---|
| +1.0 | Fully open-source; deployed bytecode reproducibly verified from source. | Audited by ≥ 2 reputable firms; formal verification of security-critical paths. | Critical parameters immutable; no upgrade pathway; no privileged admin keys. |
| +0.5 | Open-source with minor proprietary dependencies; bytecode-source correspondence published. | Single reputable audit; behaviour documented; community review present. | Governance changes require ≥ 7 -day timelock, multi-party veto, on-chain visibility. |
| 0.0 | Source published but no bytecode-source verification. | Code reviewed informally; no professional audit. | Governance can change parameters within disclosed bounded ranges with short timelock. |
| -0.5 | Closed-source with audit summary; or open-source with major closed dependencies. | No professional audit; behaviour determinable only by reverse engineering. | Multi-sig admin can alter critical parameters without timelock. |
| -1.0 | Closed-source proprietary code; no bytecode published. | Behaviour not determinable; no audit; possible undisclosed backdoors. | Single admin key with arbitrary upgrade authority. |

Note: Composite $S = \min(S_code, S_behav, S_gov)$. For composable contracts, S is further reduced to the minimum across all dependencies on which the contract substantively relies (§7.2).

6.2. The Deliverability rubric for digital contracts

D decomposes into D_tech (on-chain technical verifiability), D_onto (off-chain ontological deliverability), and D_order (ordering integrity under adversarial execution). The composite $D = \min(D_tech, D_onto, D_order)$, with D_order applying only in permissionless settlement environments (see [Table 3](#) below).

7. TRAJECTORIES, NOT POINTS

In digital contexts a contract’s coordinate is most accurately represented not as a point but as a trajectory — a function from time to coordinate. The framework’s admissibility, quality, and prudential tests are extended to trajectories along three dimensions: temporal, compositional, and attestation-driven.

Table 3: Deliverability (D) rubric for digital contracts

| Score | D_tech (on-chain verifiability) | D_onto (off-chain referent) | D_order (ordering integrity) |
|-------|--|--|---|
| +1.0 | Fully on-chain settlement; deterministic execution; no oracle dependency. | No off-chain referent or referent fully attested with real-time proof-of-reserves. | Cryptographic ordering guarantee; MEV neutralised at protocol level. |
| +0.5 | On-chain settlement with single reputable independently-verifiable oracle. | Off-chain referent attested \geq monthly by reputable third-party auditor; legal recourse available. | Private mempool or MEV-aware order routing; observable bounded slippage. |
| 0.0 | On-chain settlement with single oracle; periodic but infrequent attestation. | Off-chain referent attested by issuer self-disclosure; legal recourse uncertain. | Public mempool with disclosed slippage tolerance; user has set reasonable bounds. |
| -0.5 | Off-chain settlement contingent on external action; oracle centralised and not independently verifiable. | Off-chain referent’s existence in question; custodian solvency uncertain. | Public mempool with no slippage protection; documented sandwich-attack exposure. |
| -1.0 | Settlement requires trust in single party with no recourse. | No off-chain referent claimed but contract presupposes one; or referent fundamentally non-existent. | Adversarial sequencing structurally guaranteed; user cannot avoid extraction. |

Note: For dynamic-attestation contracts, D is computed against the worst observable state in the attestation window per §7.3.

7.1. Temporal generalisation

Let C be a digital contract instantiated at t_0 with anticipated end at t_T . Its coordinate is $(S(t), D(t))$: $[t_0, t_T] \rightarrow [-1, +1]^2$. Admissibility requires $(S(t), D(t)) \in Q1$ for all t . A trajectory that crosses out of $Q1$ at any time produces a moment of *batil* status that, under stricter readings, may invalidate the contract for that interval or entirely. The framework records this in the admissibility profile $I(t) = 1$ iff $(S(t), D(t)) \in Q1$.

Quality for a trajectory is the time-weighted (or exposure-weighted) average of the d-gap. The within-Q1 quality bands generalise: a trajectory whose worst-state d-gap remains in Robust is uniformly Robust. The prudential near-axis review band is more conservative for trajectories than for static contracts: a trajectory whose (S, D) approaches the axes during any observed interval should trigger panel review, because the brief approach evidences possible worse drift under unobserved conditions.

7.2. Compositional generalisation: dependency-graph aggregation

A composed digital contract is a directed graph $G = (V, E)$ in which vertices are individual protocols and edges are dependency relations. Each vertex has its own (S_v, D_v) trajectory. Under non-substitutability the natural aggregation is the minimum across all vertices reachable from the user's position:

$$S_{\text{composed}}(t) = \min \{ S_v(t) : v \in \text{Reach}(\text{user}, G) \}$$

$$D_{\text{composed}}(t) = \min \{ D_v(t) : v \in \text{Reach}(\text{user}, G) \}$$

Three implications follow. First, a single Q4 vertex anywhere in the user's reachable subgraph collapses the composition to *batil*. Second, the composed coordinate evolves as the dependency graph evolves: an oracle upgrade at any vertex propagates downstream. Third, the user is responsible for understanding the graph: a user who interacts only with the front-line protocol while remaining unaware of its dependencies has not met the *bayan* condition. This places a substantial epistemic burden on retail users and is the framework's strongest argument for institutional intermediation in digital finance.

A practical refinement applies to within-Q1 quality measurement. The minimum-aggregation rule above is the operative rule for admissibility: a single Q4 component renders the composition *batil*, reflecting the doctrinal non-substitutability of §3.2. For quality stratification within Q1, however, pure minimum-aggregation can be overly conservative, in that a composition of one Marginal-band protocol with several Robust-band protocols inherits the worst-case score even though the user's effective exposure to the worst component is small. The framework therefore proposes a dual rule: admissibility uses the minimum (unchanged); within-Q1 quality is computed as the exposure-weighted harmonic mean of constituent d-gaps, floored at the minimum component d-gap. This preserves the binding-axis logic at the admissibility boundary while permitting beneficial composition within Q1 to be recognised. The exposure-weighted harmonic mean is the natural choice because, like the minimum, it is dominated by small values — preserving conservatism — while permitting marginal contribution from secondary exposures.

7.3. Attestation-driven generalisation

For contracts whose D depends on continuing real-world conditions, D must be re-evaluated against a live attestation feed. Three regimes deserve named treatment. Real-time on-chain attestation: D is continuously updated. Periodic off-chain attestation: D is updated at the cadence and bounded between attestations by the most recent value with a deduction proportional to the lag and

the underlying volatility. Attestation-on-event: D is updated only when triggered (e.g., redemption requests) and bounded between events by the worst plausible state.

8. THE TAYYIB CONSTRAINT IN DIGITAL CONTEXTS

Digital assets raise distinctive *maqasid* considerations that the asset-class weight schemata for ordinary commercial contracts do not capture. A separate digital-asset weight schema is required.

8.1. Energy intensity and ‘ird-with-environment

Proof-of-work blockchains, particularly Bitcoin, consume substantial electrical energy with persistently high carbon intensity (Cambridge Bitcoin Electricity Consumption Index). The *maqasid* implication is direct: a contract whose execution depends on proof-of-work settlement inherits the carbon intensity of that layer, and the ‘ird-with-environment weight in the τ assessment falls accordingly. Proof-of-stake settlement (post-Merge Ethereum, Solana, Cosmos chains) reduces this concern by orders of magnitude. Ibn Ashur’s extension of ‘ird to environmental integrity, developed further by Auda, grounds this consideration doctrinally.

8.2. Financial inclusion and mal

Stablecoins have, in jurisdictions with weak banking infrastructure or hyperinflationary currencies, emerged as the de facto savings instrument for unbanked populations – the *mal maqasid* is served. At the same time, retail crypto trading in developed markets has generated substantial losses for individual investors, with documented patterns of speculative addiction. The same instrument class can therefore raise τ in one usage context and lower it in another. The framework requires usage-context disclosure: a stablecoin used for savings in Argentina has different τ than the same stablecoin used as leveraged trading collateral in Singapore.

8.3. Surveillance, dignity, and ‘aql

Centralised stablecoins exhibit comprehensive transactional surveillance – every transaction logged, traceable, subject to ex-post issuer intervention (freezing, blacklisting, reversal). The dignity dimension of ‘ird is implicated when an instrument exposes users to unannounced surveillance or discretionary issuer intervention without due process. Permissionless blockchains exhibit the opposite property – transparency without identity binding – with its own dignity implications and concerns. The ‘aql *maqasid* is implicated by tokenomics that encourage gambling-like engagement patterns. The τ assessment must specify which dimension is operative for the contract’s typical use case.

8.4. Proposed digital-asset weight schema

Table 4: Proposed maqasid weight schema for digital assets.

| Asset sub-class | <i>din</i> | <i>nafs</i> | ‘aql | <i>nasl</i> | <i>mal</i> | ‘ird/env | τ_{\min} |
|--------------------------------|------------|-------------|------|-------------|------------|----------|---------------|
| Custodial stablecoin (savings) | 0.10 | 0.10 | 0.15 | 0.10 | 0.35 | 0.20 | 0.55 |
| DeFi protocol (yield/lending) | 0.10 | 0.10 | 0.20 | 0.05 | 0.30 | 0.25 | 0.60 |
| Tokenised real-world asset | 0.15 | 0.10 | 0.15 | 0.15 | 0.25 | 0.20 | 0.55 |
| On-chain sukuk | 0.20 | 0.10 | 0.10 | 0.15 | 0.20 | 0.25 | 0.65 |
| Native cryptocurrency (PoW) | 0.10 | 0.10 | 0.15 | 0.05 | 0.20 | 0.40 | 0.65 |
| Native cryptocurrency (PoS) | 0.10 | 0.10 | 0.20 | 0.05 | 0.30 | 0.25 | 0.60 |

Note(s): Weights sum to 1.00 across each row. Boards may deviate by ± 0.05 on any single weight with disclosed justification. Empirical refinement of these weights, through structured Shari’ah-board consultation, is part of the research agenda (§13).

8.5. Measurable indicators for τ

To prevent τ scoring from collapsing into unobservable expert discretion, the framework proposes one measurable primary indicator per maqсад (see [Table 5](#)). Boards may select alternative indicators with disclosed justification; the proposal is offered as a default that supports inter-rater reproducibility and ex-post auditability. Each indicator is normalised to [0, 1], multiplied by the asset-class weight from [Table 4](#), and summed to produce τ .

Table 5: Measurable indicators

| Maqсад | Proposed primary indicator | Source / measurement |
|---------------------------|---|---|
| <i>din</i> | Presence of audited Shari'ah governance documentation (binary) | Direct disclosure |
| <i>nafs</i> | User-loss rate from exploits or forced liquidations per 10,000 active addresses | Rekt.news, DeFi-Llama incident registry |
| <i>'aql</i> | Absence of gambling-like mechanics (binary; e.g., variable APY > 30% absent underlying yield) | Tokenomics audit |
| <i>nasl</i> | KYC/AML compliance proportionate to risk (graded 0 to 1) | Regulatory-filing review |
| <i>mal</i> | Share of transactions originating from jurisdictions with <50% bank-account penetration | World Bank Findex × on-chain analytics |
| <i>'ird / environment</i> | Settlement-layer carbon intensity (kWh per transaction, normalised to PoS baseline) | Cambridge CCAF index |

Note(s): The proposal converts τ from a holistic judgment into an auditable score while preserving maqasid flexibility through the weight schema. Empirical anchoring of τ_{\min} — including the percentile-based calibration that some advisors have proposed — requires the structured Shari'ah-board consultation specified in §13.3, and is therefore part of the research agenda rather than a current claim.

9. STO REMEDIATION IN DIGITAL CONTEXTS

The STO taxonomy of §3.7 extends to digital contracts with new instruments specific to the technological substrate.

9.1. Substitution: open-sourcing, formal verification, attestation

Substitution acts primarily on S. Principal digital substitution instruments: open-sourcing previously proprietary code (raising S_{code}); commissioning formal verification of security-critical paths (raising S_{behav}); replacing closed dependencies with audited open-source equivalents; publishing reproducible build pipelines so deployed bytecode matches source; replacing self-attested reserves with third-party attestation (raising D_{onto}).

9.2. Transformation: re-baselining via Islamic legal frame

Transformation acts primarily on D. Principal digital transformations: restructuring a forward-deliverable token contract as on-chain *salam*, with strict specification re-baselining D from negative to positive; restructuring a tokenised future asset as on-chain *istisna'* with milestone-based release; restructuring a synthetic derivative as a *musharakah*-style profit-and-loss-sharing position with real underlying exposure; restructuring a custodial token as a *wakala* arrangement with explicit agency.

9.3. Offset: takaful, multi-sig, circuit breakers, real-time proofs

Offset acts on either axis. Principal digital offset instruments: on-chain *takaful* coverage of smart-contract risk; multi-signature governance with documented signers and time-delayed execution; protocol-level circuit breakers; real-time proof-of-reserves; MEV-aware transaction routing.

9.4. The limit principle for digital contracts

STO is effective only where at least one axis is structurally recoverable. Where the deliverability deficit reflects ontological inexistence rather than contingent inaccessibility, no remediation can restore admissibility. Synthetic perpetual swaps with no real-economy referent, algorithmic stablecoins whose peg mechanism is endogenously circular (§12.2), and self-referential token economies whose only “value” derives from the willingness of subsequent participants to enter at higher prices all exhibit ontological inexistence. STO cannot rescue them. The framework declares such contracts irreducibly *batil*.

10. 10PERATIONAL PROTOCOL: HYBRID HUMAN-MACHINE SCORING

Human raters cannot scale to the volume and technical complexity of digital contracts. The protocol must be hybridised: partial automation of S (static analysis) and D (runtime monitoring), with human raters and panel adjudication preserved for τ and contested cases.

10.1. Static analysis layer (S-axis automation)

S_code can be substantially automated. Tools (Slither, Mythril, Securify, Echidna) statically analyse Solidity bytecode for known vulnerabilities, dependency footprints, and access-control structures. S_behav cannot be fully automated — formal-verification tools require human specification of properties — but the presence and quality of formal verification can be detected automatically. S_gov is fully automatable: admin keys, upgrade pathways, timelock parameters, and governance contract addresses are all on-chain.

The framework additionally specifies three Shari’ah-relevant pattern categories that the static analysis layer should flag for human review. First, hidden-*riba* patterns: algorithmic fee or yield structures that compound on principal without corresponding underlying-asset growth, structurally equivalent to interest accrual regardless of how named. Second, *Gharar*-inducing administrative backdoors: functions that allow a single admin (or small multi-sig below the timelock threshold) to alter the contract’s subject matter post-issuance — replacing the underlying token, changing the redemption asset, or modifying the price feed without holder vote. Third, gambling-mechanic patterns: tokenomics that distribute rewards through luck-based mechanisms (random allocations, lottery-structured emissions) rather than productive use, implicating both the *maysir* prohibition and the *‘aql maqsid*. Detection of any of these patterns triggers panel review regardless of the rubric scores produced by the standard tools.

10.2. Runtime monitoring layer (D-axis automation)

D_tech and D_order are observable at runtime. Oracle reliability metrics, bridge solvency (reserves vs. issued tokens), protocol health (TVL, withdrawal patterns), and the user’s execution venue all yield direct readings. D_onto is harder to automate but the cadence and provenance of attestation are observable. The framework therefore proposes a continuous monitoring layer that updates the (S, D) trajectory of each held position, flagging excursions in real time.

10.3. Panel adjudication for contested cases

Three categories remain panel-mandated. First, all τ assessments — the *maqasid* weighting decisions of §8 cannot be automated. Second, cases where automated and human assessments disagree by more than 0.3 on any axis. Third, structurally novel cases where the rubric has no precedent. The panel’s record must include the trajectory observed during the assessment window, not only a snapshot.

10.4. Annual digital integrity heatmap

Institutions report, per quarter and annually, six metrics: (i) holding-time-weighted exposure proportions in each quality band; (ii) worst-state coordinate per protocol; (iii) duration and frequency of prudential-band excursions; (iv) time-integrated d-gap normalised by exposure; (v) proportion of positions with automated/human disagreement beyond 0.3; (vi) realised exploit and slashing losses, separated by the axis of origination. The reporting form is a time-distribution rather than a snapshot.

11. FALSIFIABILITY CONDITIONS

The framework commits to disconfirmation criteria. Six independently testable conditions specify the empirical evidence that would falsify each major commitment.

11.1. F1: predictive failure on exploit and dispute incidence

If digital contracts assessed in the Robust band — $m(C) \geq 0.7$ sustained across the trajectory — produce realised exploit losses, oracle failures, governance attacks, or formal Shari'ah-board adverse opinions at rates statistically indistinguishable from contracts in the Marginal band, the within-Q1 quality stratification is falsified. Test: an exploit-tracked cohort study using the Rekt loss tracker, DeFi-Llama incident registry, and institutional risk-event reports.

11.2. F2: juristic disagreement on classification

If recognised Shari'ah authorities — Higher Shari'ah Authority (UAE) determinations, Shariyah Review Bureau opinions, Amanie Advisors' published reviews, the Shari'ah-supervisory boards of major Islamic-finance institutions with digital exposure, and any future AAOIFI-issued cryptoassets Standard — systematically rule invalid contracts the framework classifies as Q1, or systematically rule valid contracts the framework classifies outside Q1, the juristic alignment is falsified. Test: a corpus study of published rulings matched against framework scores, reported as a confusion matrix and area-under-ROC for the binary admissibility classification.

11.3. F3: STO remediation impotence

If STO interventions predicted to move a contract from outside Q1 into Q1 — open-sourcing, formal verification, on-chain proof-of-reserves, MEV-protected routing, transformation into Islamic legal frame — produce no measurable reduction in exploit losses or adverse Shari'ah opinions when implemented, the prescriptive claim for digital remediation is falsified. Test: prospective treatment-control comparison of structurally similar protocols undergoing or foregoing remediation.

11.4. F4: the structural-novelty claim

If empirical exploit-and-dispute data show that risk in digital contracts is well-explained by static, atomic, two-axis-independent rubrics — i.e., if a static-rubric model predicts exploit incidence as well as the dynamic-trajectory model — the structural-novelty claim of §4 is empirically redundant and the extensions of §§4 and 7 should be retracted. Test: model comparison of static (S, D) coordinates versus regression incorporating trajectory worst-state and composability minimum.

11.5. F5: the entanglement claim

If exploit data show that in identified entangled contracts S-failures and D-failures co-occur at rates indistinguishable from independent contracts — i.e., the joint distribution factors into the

product of marginals — the entanglement claim of §4.1 is empirically void. Test: contingency analysis of S- and D-attributed exploit events in entangled versus non-entangled contracts.

11.6. F6: the OSV qualification

If protocols with full open-source code and full cryptographic verifiability produce dispute and exploit rates indistinguishable from the framework’s prediction for fully *bayan*-compliant classical contracts, the OSV qualification of §5 is empirically unmotivated. Test: comparison of realised event rates in OSV-complete digital contracts versus matched samples of high-*bayan* classical contracts.

11.7. Empirical calibration of the quality bands: a protocol

The band thresholds in §3.5 (Robust ≥ 0.7 , Standard ≥ 0.5 , prudential < 0.2) are derived doctrinally from the classical *Gharar yasir / mutawassit / fahish* demarcations. The framework proposes the following protocol for empirical anchoring, which is part of the research agenda (§13.3) rather than completed work.

Construct a reference cohort of digital contracts active over a multi-year window with complete incident histories from established sources (Rekt.news loss tracker, DeFi-Llama Hack Database, institutional risk-event reports). Each contract is scored ex ante using the §6 rubrics by ≥ 3 independent raters blinded to outcomes; inter-rater reliability is measured by Krippendorff’s α for ordinal data, with low- α contracts excluded from the calibration set. Critical failure is operationally defined as one of: exploit loss exceeding 1% of total value locked, oracle failure causing $> 5\%$ deviation from reference price, governance takeover, or formal Shari’ah adverse opinion from a recognised authority. Logistic regression of failure incidence on $\min(S, D)$ yields fitted inflection points; if the fitted points converge on the doctrinal thresholds (0.7, 0.5, 0.2) within measurement error, the bands are validated and retained. If they diverge systematically, the bands are recalibrated and the divergence is itself recorded as a finding about how digital risk distributes relative to the classical typology. Annual recalibration is recommended once the cohort exceeds a stable size.

The framework commits to publishing the calibration dataset, rater instructions, and statistical methods alongside any results, in line with the F1 falsifiability condition. The protocol is offered specifically because the absence of such anchoring is a legitimate methodological objection to the band thresholds as proposed.

12. FIVE APPLIED CASE STUDIES

The cases that follow are chosen for structural variety. Scores are illustrative and reflect the framework applied at representative observation points.

12.1. Custodial stablecoin (USDC): the entanglement archetype

USD Coin, issued by Circle, is redeemable 1:1 against US dollar reserves attested monthly by a Big-Four auditor. The instrument exhibits entanglement (§4.1) in pure form: the user’s only access to the reserves’ composition is the issuer’s attestation, and that same attestation establishes the reserves’ existence for third parties.

Specifiability: ERC-20 contract is open-source and bytecode-verified, but the centralised admin holds documented freeze-and-blacklist authority ($S_code = +0.7$); behaviour is well-audited, reserve-management process is internally governed and opaque in detail though attested in aggregate ($S_behav = +0.5$); admin can pause and freeze without timelock ($S_gov = +0.5$). Composite $S = +0.5$.

Deliverability: on-chain transfer is deterministic, but admin pause-and-freeze can interrupt redemption ($D_{\text{tech}} = +0.7$); monthly attestation of cash-and-Treasuries reserves with US legal recourse, the March 2023 SVB depeg episode notwithstanding ($D_{\text{onto}} = +0.5$); direct issuer redemption is off-chain and not subject to MEV ($D_{\text{order}} = +1.0$). Composite $D = +0.5$.

At $(+0.5, +0.5)$: Q1, Standard band, $d\text{-gap} \approx 0.71$. The contract is admissible. τ assessment varies by use context (Table 4 row 1): savings/remittance use in jurisdictions with weak banking infrastructure produces $\tau \approx 0.65$ (above $\tau_{\text{min}} = 0.55$); leveraged-trading-collateral use in developed markets produces $\tau \approx 0.40$ (below τ_{min}). USDC is therefore *Tayyib*-compliant for savings use, not for leveraged-trading use — a use-context distinction unavailable in classificatory frameworks. STO recommendation: continuous proof-of-reserves to raise D_{onto} from $+0.5$ to $+1.0$, lifting composite D into the Robust band.

12.2. Algorithmic stablecoin (UST/Terra): the limit-principle case

TerraUSD lost its peg in May 2022, destroying ~USD 40 billion within ten days. UST maintained its dollar peg through a mint-and-burn mechanism with sister token LUNA: one UST always exchangeable for one dollar's worth of LUNA. The peg held when LUNA had market value, and LUNA had market value because UST was useful (and because Anchor offered ~20% UST yield subsidised from Terraform Labs reserves). The system was endogenously circular: no exogenous reserve, no real-economy referent, no non-circular grounding. When confidence weakened, the mint-burn arbitrage produced LUNA hyperinflation and the peg collapsed.

Specifiability: open-source code ($S_{\text{code}} = +1.0$), reasonable documentation ($S_{\text{behav}} = +0.5$), Terraform-controlled governance with limited timelock ($S_{\text{gov}} = 0.0$). Composite $S = 0.0$ — at the boundary. The decisive failure is on D . D_{tech} for the on-chain mint-burn was complete ($+1.0$); D_{order} acceptable. The decisive collapse: D_{onto} . The off-chain referent UST claimed to deliver — a stable dollar — did not exist in any non-circular sense. $D_{\text{onto}} = -1.0$. Composite $D = -1.0$.

Verdict: $(S, D) = (0.0, -1.0)$. Q4. *Batil*. The depeg was not a tail risk the framework failed to predict; properly applied at inception, the framework would have placed UST in Q4. The case demonstrates the limit principle (§9.4): STO cannot rescue UST because the deficit is structural inexistence, not contingent inaccessibility. Substitution cannot conjure reserves where none exist; Transformation into *salam* or *wakala* is not available because there is no underlying asset to specify; Offset through *takaful* only redistributes loss. The contract is irreducibly *batil*. The framework generates two independent grounds for refusal: *batil* under (S, D) , and $\tau \approx 0.20$ (well below any plausible τ_{min}) under *maqasid*.

12.3. Cross-chain bridge: the compositional archetype

Cross-chain bridges have suffered exploit losses exceeding USD 2 billion since 2021. The structural reason is composability: every protocol that uses a bridged token inherits the bridge's S and D , while users typically see only the front-line protocol.

Consider a wrapped-token bridge B with proxy-upgrade pattern ($S_{\text{code}} = +0.7$), audited but with multiple post-deployment patches ($S_{\text{behav}} = +0.5$), permissioned validator set replaceable by multi-sig ($S_{\text{gov}} = 0.0$). Composite $S(B) = 0.0$. $D(B)$: on-chain settlement on both sides ($D_{\text{tech}} = +0.7$), locked underlying with validator-set risk ($D_{\text{onto}} = +0.5$). Composite $D(B) \approx +0.5$. B alone is at $(0.0, +0.5)$ — Marginal Q1, on the boundary.

Compositional aggregation produces the more important verdict. A user holding wrapped tokens on the destination chain, depositing into a lending protocol L with $(S(L), D(L)) = (+0.7, +0.7)$, faces effective coordinates:

$$S_{\text{composed}} = \min(+0.7, 0.0) = 0.0$$

$$D_{\text{composed}} = \min(+0.7, +0.5) = +0.5$$

The bridge's S has dominated the composition. The user, observing only L, would estimate Robust quality. The framework reveals it as boundary Marginal. The Ronin exploit of March 2022 (five of nine validators compromised, ~USD 600 million drained) would have been correctly anticipated for any protocol composing through the bridge: Ronin's validator-set sub-dimension was below 0.0 under [Table 2](#), and any composed position inherited that score. STO recommendations: Substitution to permissionless light-client or ZK-bridge verification (raising S_{gov} from 0.0 to +0.7); Offset through bridge insurance via on-chain *takaful*; Transformation restructuring wrapped-token issuance as *wakala-bil-istithmar*. The framework's recommendation: require either ZK or light-client verification, or refuse the position.

12.4. DeFi lending (Aave): the trajectory case

Aave's parameters — collateral factors, liquidation thresholds, supported assets, oracle sources, fee structures — are governed by AAVE token-holder vote. The protocol's coordinate is a trajectory, not a static point.

Single-point assessment. S: open-source, bytecode-verified, proxy-upgrade under timelock ($S_{\text{code}} = +0.7$); multiple audits with formal verification of selected paths ($S_{\text{behav}} = +0.7$); 7-day timelock and guardian veto, but parameter range is wide ($S_{\text{gov}} = +0.5$). Composite S = +0.5. D: on-chain settlement with critical Chainlink dependency ($D_{\text{tech}} = +0.5$); native-token deposits exhaustively determined by chain state ($D_{\text{onto}} = +1.0$); MEV exposure ($D_{\text{order}} = 0.0$ default, +0.5 with private routing). Composite D for native tokens with private mempool = +0.5.

Trajectory observations. A March 2023 governance vote increased a collateral factor; for the 18 hours before an oracle correction, effective S_{gov} was bounded by the new parameter and D_{tech} by oracle reliability. Composite (S, D) during the interval was approximately (+0.3, 0.0) — Marginal, near the prudential band. The Curve crisis of August 2023 produced a cascade through Aave's CRV-collateralised exposure; trajectory analysis would have flagged the D_{tech} excursion to negative values during the cascade peak.

τ assessment ([Table 4](#) row 2): primary use cases are leveraged trading (low τ) and collateralised stablecoin borrowing for liquidity (moderate τ). The protocol enables interest-bearing positions on both sides — a structural impediment to direct Shari'ah-compliance. $\tau \approx 0.30\text{--}0.45$, below $\tau_{\text{min}} = 0.60$. Verdict: not *batil* under (S, D) but *Tayyib*-non-compliant for institutional Islamic-finance use. STO: a Shari'ah-compliant alternative protocol structure (*musharakah*-style positions replacing interest, *wakala* replacing interest-bearing deposits) constitutes the constructive Transformation pathway.

12.5. On-chain sukuk: the constructive case

The first four cases identify framework verdicts on instruments not designed for compliance. The fifth is the constructive complement: an on-chain *sukuk* specifically structured for clean Q1 admissibility and $\tau \geq \tau_{\text{min}}$. The structure is described in Muneeza et al. (2021) and is operationally feasible; AAOIFI Shari'ah Standard No. 62 (*Sukuk*, exposure draft 2024) provides the doctrinal frame, Standard No. 17 covers the underlying investment-*sukuk* structure, and Standard No. 18 (Qabd) governs constructive possession of the tokenised representation.

A *sukuk al-ijarah* is issued on-chain. The SPV issues tokens representing fractional ownership in a leased commercial property. Lease payments are remitted on-chain monthly by the lessee; the SPV

distributes pro rata to token-holders; principal amortises over the tenor through a defined schedule. Asset valuation is attested quarterly by an independent appraiser, with the hash committed on-chain. Governance is locked at issuance (immutable bytecode); changes require a holder vote with 30-day timelock plus Shari'ah-board and trustee approval.

Specifiability: open-source SPV contract, bytecode-verified, simple distribution logic ($S_{code} = +1.0$); audited and formally verified for distribution and amortisation paths ($S_{behav} = +1.0$); bytecode-locked with strong upgrade barriers ($S_{gov} = +1.0$). Composite $S = +1.0$. Deliverability: on-chain settlement, deterministic, oracle limited to appraisal-hash commitment ($D_{tech} = +1.0$); real underlying property, professionally appraised, registry-recorded, legal recourse ($D_{onto} = +1.0$); scheduled SPV distributions not subject to MEV ($D_{order} = +1.0$). Composite $D = +1.0$.

At $(+1.0, +1.0)$: exactly v^* , $d\text{-gap} = 0$, Robust band. τ assessment (Table 4 row 4): real-economy productive use, lease structure compliant by construction, digital infrastructure raises *mal* through cost reduction, environmental performance can be incorporated. $\tau \approx 0.75$, above $\tau_{min} = 0.65$. STO is not required. The case demonstrates the framework's constructive use: digital infrastructure used to enhance, not replace, the Islamic legal frame.

13. DISCUSSION

13.1. Toward a future AAOIFI Cryptoassets Standard

The standards-development gap identified in §2.2 is, on closer view, an opportunity. AAOIFI's Shari'ah Standards have been issued through a mature deliberative process culminating in operational text adopted by member institutions and recognised by national regulators. The absence of a dedicated cryptoassets Standard does not reflect doctrinal indifference — the 21st AAOIFI Shari'ah Council Conference (2023) addressed digital assets at length, and the AAOIFI Shari'ah Board has discussed cryptoassets in working sessions — but reflects rather the substantial technical and doctrinal complexity of the topic. The (S, D, τ) framework is offered as a candidate scaffold for a future Standard.

Three points of contact merit explicit identification. First, the framework operates within the typological structure that AAOIFI Standards already deploy: Standard No. 21 (Financial Papers) for security tokens, No. 17 (Investment *Sukuk*) and No. 62 (*Sukuk*, exposure draft) for tokenised *sukuk*, No. 18 (*Qabd*) for tokenised possession, and No. 57 (Gold) for precious-metal-backed tokens. The framework provides a continuous quality metric within each typological category and a unified test across them, but does not displace the typological work the Standards perform. Second, the falsifiability conditions of §11 anticipate disagreement with future Shari'ah authority and specify it as a falsifier rather than a vindicator: F2 is satisfied when authority and framework converge, falsified when they diverge systematically. Third, the rubric calibrations of Tables 2–4 are proposed; their precise score-anchoring requires the kind of structured deliberative process AAOIFI has institutionalised. The framework is therefore offered as raw material for that process, not as a substitute for it.

13.2. Alignment with the Basel Committee's prudential treatment of cryptoasset exposures

The framework also has substantive contact with the Basel Committee on Banking Supervision's prudential treatment of cryptoasset exposures (BCBS SCO60, finalised December 2022, revised July 2024, with implementation date 1 January 2026). The Basel framework classifies cryptoassets into Group 1 (tokenised traditional assets and stablecoins meeting redemption-test and stabilisation conditions) and Group 2 (everything else, with substantially higher capital charges and a portfolio cap of 1–2% of Tier 1 capital). Group 1 Classification Condition 3 — risk governance and risk control

— explicitly requires that the cryptoasset and its underlying network be designed and operated to mitigate risks affecting transferability, settlement finality, and redeemability, including operational risk and risk of loss of data. The Basel standard additionally permits supervisors to apply an infrastructure-risk add-on to Group 1 exposures based on observed weaknesses in the underlying infrastructure.

The (S, D) coordinate provides a candidate quantitative basis for both the Classification Condition 3 assessment and the infrastructure-risk add-on calibration. A cryptoasset with high composite S (transparent code, audited behaviour, immutable or strongly-buffered governance) and high composite D (verifiable on-chain settlement, attested off-chain referent, ordering integrity) has, by construction, the operational and governance properties Classification Condition 3 requires. A cryptoasset with marginal or sub-band scores on either axis exhibits exactly the kind of infrastructure weakness the add-on is intended to capture. Islamic banks subject to both Shari’ah-supervisory and Basel-prudential reporting can therefore use the (S, D, τ) framework as a single underlying assessment that supports both regulatory regimes — Shari’ah compliance from the (S, D) admissibility test and τ floor, prudential reporting from the (S, D) infrastructure-risk component. This dual use is one of the framework’s strongest practical motivations for institutional adoption.

13.3. Limitations

Three limitations warrant acknowledgement. First, rubric calibrations are proposed pending the empirical work specified in §11 (F1, F3) and the structured Shari’ah-board consultation specified above. Second, the framework presupposes raters with both fiqh competence and digital-finance technical competence, currently in short supply; the hybrid protocol of §10 partially addresses this through automation but the residual panel work is non-trivial. Third, the framework operates at the contract level; portfolio-level systemic considerations are addressed only through the institutional heatmap of §10.4 and not at the regulatory-systemic level proper.

13.4. Research agenda

Five lines of work follow. First, empirical calibration of the rubrics through the falsifiability tests of §11 — a multi-year cohort study of digital-asset positions matched against realised exploit and dispute events. Second, formalisation of dependency-graph aggregation (§7.2) for multi-protocol DeFi positions, including cyclical dependencies and the user-attribution problem. Third, structured Shari’ah-board consultation across major Islamic-finance jurisdictions to refine the τ -weight schema of [Table 4](#). Fourth, development of automated tooling for the static-analysis layer (§10.1), in particular Shari’ah-aware static analysers that flag *riba*-equivalent fee schedules and gambling-like incentive structures within smart-contract bytecode. Fifth, design and pilot deployment of the on-chain *sukuk* and Shari’ah-compliant DeFi structures of §§12.4 and 12.5, with documented trajectory data over a multi-year holding period.

14. CONCLUSION

This paper has developed a self-contained quantitative framework for assessing *Gharar* in Islamic digital finance. The framework rests on two structural axes derived from the classical taxonomy: Specifiability (epistemic) and Deliverability (ontological), non-substitutable, with sign-based admissibility, continuous within-Q1 quality stratification through the d-gap and the min-constraint, an independent *Tayyib* floor τ grounded in *maqasid* scholarship, and a Substitution-Transformation-Offset remediation taxonomy. The extension to digital contracts identifies five structurally novel properties — entanglement, composability, mutability, ordering risk, time-dependence — and develops the rubric decomposition, trajectory generalisation, digital τ schema, digital STO instruments, and hybrid operational protocol that those properties require.

Three contributions distinguish the work. First, doctrinal: the framework is anchored to classical fiqh sources (al-Sarakhsi, al-Kasani, Ibn Qudamah, Ibn Rushd) and contemporary *maqasid* scholarship (Ibn Ashur, Chapra, Auda) without collapsing into either juristic absolutism or informal pluralism. Second, operational: rubrics, thresholds, and aggregation rules are stated in falsifiable form, with six independent disconfirmation conditions specified in §11. Third, institutional: the framework is offered as a candidate scaffold for the Shari'ah Standard on cryptoassets that does not yet exist, complementing rather than displacing the typological work of existing AAOIFI Standards.

The cryptoasset economy is not waiting for Islamic finance to be ready. The framework's task is to make Islamic finance ready for it — doctrinally faithful, operationally specific, scientifically disciplined, and open to revision through the empirical and deliberative processes the research agenda specifies.

APPENDIX A. NOTATION SUMMARY

| Symbol | Meaning | § |
|-------------------------|---|----------|
| (S, D) | Specifiability (epistemic) and Deliverability (ontological) axes; coordinate space $[-1, +1]^2$ | 3.1 |
| v^* | Optimum vector (+1, +1); ideal but not generally attainable | 3.4 |
| Q1, Q2, Q3, Q4 | Quadrants; Q1 admissible, Q2/Q3/Q4 batil | 3.3 |
| d-gap(C) | Euclidean distance from v^* : $\sqrt{[(1-S)^2 + (1-D)^2]}$ | 3.5 |
| m(C) | Binding axis: $\min(S, D)$ | 3.5 |
| τ | Tayyib constraint, maqasid-grounded floor in $[0, 1]$ | 3.6 |
| τ_{\min} | Asset-class-specific Tayyib floor | 3.6, 8.4 |
| STO | Substitution / Transformation / Offset remediation taxonomy | 3.7 |
| S_code, S_behav, S_gov | S sub-dimensions: code transparency, behavioural intelligibility, governance immutability | 6.1 |
| D_tech, D_onto, D_order | D sub-dimensions: on-chain technical, off-chain ontological, ordering integrity | 6.2 |
| (S(t), D(t)) | Trajectory representation: coordinate as function of time | 7.1 |
| I(t) | Admissibility profile: $I(t) = 1$ iff $(S(t), D(t)) \in Q1$ | 7.1 |
| $G = (V, E)$ | Dependency graph for composed digital contract | 7.2 |
| F1–F6 | Falsifiability conditions | 11 |

Declarations

The authors have no relevant financial or non-financial interests to disclose.

AI Use Declaration

Anthropic's Claude (claude-opus-4-7) was used as a writing assistant during the preparation of this manuscript. The author conceived the framework, the structural-novelty argument, the trajectory generalisation, the digital τ schema, the digital STO instrumentation, and the hybrid operational protocol. Case studies were specified by the author and developed in collaboration with the AI assistant for prose drafting and rubric application. The AI assistant did not contribute to the doctrinal reasoning or the falsifiability conditions, which reflect the author's independent scholarship. The author has reviewed all AI-assisted text and accepts full responsibility for accuracy and integrity. No AI tool is listed as a co-author.

REFERENCES

- AAOIFI (Accounting and Auditing Organization for Islamic Financial Institutions). (2017). *Shari'ah Standards*. Manama: AAOIFI.
- AAOIFI. (2024). *Exposure Draft of Shari'ah Standard No. 62 on Sukuk*. Manama: AAOIFI.
- Adam, M. F. (2017). Bitcoin: Shari'ah Compliant? *Amanah Finance Consultancy Working Paper*.
- Adam, M. F. (2019). Shari'ah Analysis of Cryptocurrencies and Initial Coin Offerings. *Amanah Finance Consultancy Working Paper*.
- al-Kasani, A.-D. (12th c.). *Bada'i' al-Sana'i' fi Tartib al-Shara'i'*. Beirut: Dar al-Kutub al-'Ilmiyya edition (1986).
- al-Saati, A.-R. (2003). The Permissible Gharar (Risk) in Classical Islamic Jurisprudence. *Journal of King Abdulaziz University: Islamic Economics* 16(2): 3–19.

- al-Sarakhsi, M. (11th c.). *Al-Mabsut*. Beirut: Dar al-Ma'rifa edition (1989).
- al-Suwailem, S. (2000). Towards an Objective Measure of *Gharar* in Exchange. *Islamic Economic Studies*, 7(1&2), 61-102.
- al-Suwailem, S. (2006). *Hedging in Islamic Finance*. Jeddah: Islamic Research and Training Institute, Islamic Development Bank.
- Auda, J. (2008). *Maqasid al-Shari'ah as Philosophy of Islamic Law: A Systems Approach*. London: International Institute of Islamic Thought.
- Basel Committee on Banking Supervision (BCBS). (2022). *Prudential Treatment of Cryptoasset Exposures*. Basel: Bank for International Settlements (SCO60, December 2022, revised July 2024).
- Cambridge Centre for Alternative Finance. Cambridge Bitcoin Electricity Consumption Index. University of Cambridge Judge Business School. <https://ccaf.io/cbnsi/cbeci>
- Chapra, M. U. (1992). *Islam and the Economic Challenge*. Leicester: Islamic Foundation.
- Chapra, M. U. (2008). *The Islamic Vision of Development in the Light of the Maqasid al-Shari'ah*. Jeddah: IRTI.
- El-Gamal, M. A. (2006). *Islamic Finance: Law, Economics, and Practice*. New York: Cambridge University Press.
- Ibn Ashur, M. T. (2006). *Treatise on Maqasid al-Shari'ah*. Translated by M. El-Tahir El-Mesawi. Herndon: International Institute of Islamic Thought.
- Ibn Qudamah, 'A.-A. (12th c.). *Al-Mughni*. Riyadh: Dar 'Ālam al-Kutub edition (1997).
- Ibn Rushd, A.-W. (12th c.). *Bidayat al-Mujtahid wa Nihayat al-Muqtasid*. Beirut: Dar al-Ma'rifa edition (1996).
- Juels, A., A. Kosba, and E. Shi. (2016). "The Ring of Gyges: Investigating the Future of Criminal Smart Contracts." Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security: 283–295.
- Kamali, M. H. (2000). *Islamic Commercial Law: An Analysis of Futures and Options*. Cambridge: Islamic Texts Society.
- Khan, T (2026). Dynamic Prescriptive Economics: General Theory and Practice of Decision under Complexity". *Ventureethica*. https://www.tariqullah.com/dynamic-prescriptive-economics_1
- Muneeza, A., and Z. Mustapha. (2019). "Blockchain and Its Shari'ah-Compliant Structure." In *Halal Cryptocurrency Management*, ed. M. M. Billah. Cham: Palgrave Macmillan.
- Muneeza, A., et al. (2021). Smart Sukuk: A Conceptual Framework for the Application of Smart Contracts to Sukuk. *International Journal of Islamic and Middle Eastern Finance and Management* 14(5): 924–946.
- Nakamoto, S. (2008). "Bitcoin: A Peer-to-Peer Electronic Cash System." *Bitcoin White Paper*.
- Obaidullah, M. (2002). Islamic Risk Management: Towards Greater Ethics and Efficiency. *International Journal of Islamic Financial Services* 3(4): 1–18.
- Shariyah Review Bureau. (2018). *The Shari'ah Factor in Cryptocurrencies and Tokens*. Manama: SRB.
- Usmani, M. T. (2018). *Decisive Verdicts on the Cryptocurrency Question*. Karachi: Maktabah Ma'arif al-Qur'an.
- Werner, S. M., D. Perez, L. Gudgeon, A. Klages-Mundt, D. Harz, and W. Knottenbelt. (2022). "SoK: Decentralized Finance (DeFi)." arXiv:2101.08778.
- Zhou, L., et al. (2023). "SoK: Decentralized Finance (DeFi) Attacks." 2023 IEEE Symposium on Security and Privacy: 2444–2461.

