

Programmable Securities

Instrument Design for Multi-Jurisdictional Asset Exchange

Raez Lorgat

April 2026

Regulatory disclaimer. This paper describes instrument *designs* and their mathematical properties. It does not constitute legal advice, investment advice, or a solicitation to purchase any instrument. The term “programmable securities” refers to the technical architecture for instruments whose compliance constraints are computationally verifiable, not to a legal classification. Whether any specific instrument constitutes a “security,” “derivative,” “utility token,” or other regulated product is determined by the laws of each jurisdiction in which it is offered, and requires jurisdiction-specific legal analysis. Illustrative yield calculations use idealized models and are not projections of actual returns.

Abstract

Global IP licensing generates \$165-200 billion annually [1]; Islamic sukuk issuance exceeds \$200 billion per year [2]; gaming IP assets produce over \$10 billion per year in closed ecosystems [3]. The open problem is not tokenization in isolation. It is issuance, statutory gating, and secondary clearing for the same instrument across multiple jurisdictions.

We define a *programmable security* as a quintuple $(\text{contract}, \text{compiled}, \mathcal{A}_I, D, g_{\max})$ where $\text{contract} : \Sigma^{\#} \times \text{Inputs} \rightarrow \Sigma^{\#}$ is a total recursive function on a typed commitment state over a 23-domain compliance tensor, *compiled* is a gas-bounded SAVM bytecode realization, \mathcal{A}_I is the instrument’s life-cycle automaton, D is the activated compliance-domain subset, and g_{\max} is a static gas bound. A programmable security in this paper is Mass-native at issuance and Moxie-native at secondary clearing: Mass supplies the authoritative legal and compliance state; Moxie supplies the venue on which the instrument clears once the instrument-specific gate admits it. The construction satisfies four load-bearing properties: (i) every transfer is gated by the per-jurisdiction statutory threshold $\tau^{\text{reg}(I)}$ keyed by market phase $\in \{\text{primary}, \text{secondary}\}$ and holding period (Theorem 9.2); (ii) the composed tensor determines a formal graduation function from compliance state to temperature tier; Cold is pre-graduation, Warm is partial-compliance gated trading, and Hot requires each of the 23 domains to be either compliant or not applicable for the instrument; (iii) a typed Mass \leftrightarrow Moxie bridge can carry authoritative compliance commitments into venue-side settlement gates, with concrete wire status treated as an implementation parameter rather than assumed in the theorem; and (iv) corporate actions (split, dividend, merger, rights, tender, conversion, redemption) are first-class state-machine transitions with proven well-definedness (Theorem 5.32).

Scope of “graduation.” The *temperature tier* of an instrument on a compliance-aware L1 is a *venue-internal* classification computed from the instrument’s composed compliance tensor. It is *not* a legal pathway. Legal registration status under the Securities Act of 1933, Delaware corporate and securities law, MiFID II, AAOIFI-governed Islamic capital-market frameworks, insurance regulation, or any analogous regime remains jurisdiction-specific and statute-specific. An instrument whose operational tier advances from Cold to Warm does not, by virtue of that advancement, become registered,

exempt, or otherwise authorized in any jurisdiction. Every claim of “graduation” in this paper refers strictly to the venue-internal operational tier and the associated compliance tensor, not to legal classification.

We define six instrument classes: binary event contracts, revenue-linked notes, IP index products, contingent licensing options, sukuk, and experiential tokens. We also give representative tensor profiles for cross-jurisdictional sukuk, sovereign parametric insurance, event-contingent royalty instruments, and tranching structured credit. In every case, compliance is per instrument rather than per chain: one block can contain instruments with disjoint compliance profiles, and investor access is determined by the instrument-specific gate rather than by a chain-wide label. The exchange-institutional bridge is the multi-harbored entity: an issuing entity that exists simultaneously in multiple jurisdictions, with compliance composition computed on the Applicable fragment of a 23-domain tensor. The paper uses only the Applicable-fragment meet-semilattice and the F144 dichotomy’s implication that the full mixed-axis tensor is not a Heyting algebra; no result here relies on a full-tensor Heyting claim. The bridge between Mass and Moxie is an authenticated settlement envelope: Mass is the authoritative issuance protocol, Moxie is the secondary clearing venue, and programmable securities are the category the pairing is built for. Each declarative obligation pack is translated by an *obligation-pack compiler* into SAVM (Smart Asset Virtual Machine) bytecode (Definition 5.2) that executes deterministically under a stack-based ISA with compliance-gated fixed-point arithmetic, statically-disallowed reentrancy, bounded call-graph depth, and a gas-metered cost model (Section 5.1). We specify five instrument-specific potentials (equity, sukuk, convertible note, fixed income, option-like), and prove that the graduation map from the Applicable-fragment tensor to the temperature lattice is meet-preserving. The join is not preserved; multi-harbor composition uses meet.

The coupling function connecting prediction markets to instrument pricing is analysed in the Glaston-Milgrom [44], Kyle [11], and Hasbrouck [45] information-share frameworks. Proposition 10.6 shows that the minimum-variance coupling weight α^* is exactly the prediction market’s Hasbrouck information share. The economic multiplier admits a closed-form expression $\mu_{\text{econ}} = 8 \sum_k w_k V_k^{\text{base}} (1 + \eta_k) / (w_L V_L)$ as a function of named volume, willingness-to-pay, and elasticity parameters (Proposition 8.4), with proved symbolic lower bound $\mu_{\text{econ}} \geq 8(1 + \eta_{\min}) \sum_k \beta_k$ under stated elasticity-regularity. The bound is proved; only the numerical values of (β_k, η_k, w_k) are empirical. The $5 \times -20 \times$ band is the bound evaluated at representative parameters.

Contents

1	Introduction	5
1.1	Three under-served markets	5
1.2	The missing infrastructure	6
1.3	Legal classes and representative tensor profiles	7
1.4	Contributions and outline	7
2	Formal Foundations: What Is a Programmable Security?	8
2.1	The commitment algebra	9
2.2	Programmable securities: formal definition	9
2.3	Composition	10
2.4	Soundness: operational semantics refines declarative contract	12
2.5	Regulatory equivalence under composition	12
2.6	The life-cycle state machine	13
2.7	Instrument invariants	14

3	Instrument Taxonomy	14
3.1	Binary event contracts	15
3.2	Revenue-linked notes	15
3.3	IP index products	16
3.4	Contingent licensing options	17
3.5	Sukuk	17
3.6	Experiential tokens	20
4	The Mass Connection	20
4.1	Multi-harbored entities	20
4.2	IP holding through multi-harbored entities	22
4.3	Mass-Moxie bridge and temperature graduation	22
4.4	Corridor economics	24
5	The Obligation-Pack Compiler and the ConvexPotential Bridge	25
5.1	The Smart Asset Virtual Machine (SAVM): formal specification	25
5.2	The ConvexPotential interface	28
5.3	Five instrument potentials	29
5.3.1	Equity	29
5.3.2	Sukuk	29
5.3.3	Convertible note	31
5.3.4	Fixed income	33
5.3.5	Option-like	34
5.4	Temperature tier capability bindings for securities	37
5.5	Obligation pack verification	40
5.6	Corporate action semantics	41
6	Regulatory Classification	42
6.1	Classification framework	42
6.2	Primary and secondary market topology	43
6.3	Misclassification remedies	44
6.4	Binary event contracts	45
6.5	Revenue-linked notes	45
6.6	Sukuk	46
6.7	Experiential tokens	46
7	The Self-Improving Exchange	46
7.1	State space and transition operator	46
7.2	Self-calibrating coupling	47
7.3	Correlation discovery from trading data	47
7.4	Regime detection	48
7.5	Contraction and convergence	49
7.6	The informational barrier to entry	49
8	The Attention Economy Reinterpretation	50
8.1	Coupling as attention pricing	50
8.2	Information value versus LVR reduction	50
8.3	Prediction market participation as demand revelation	53
8.4	The 5-20× multiplier	53
9	Related Work	54

10 The Information-Value Multiplier: Closed-Form Parameterization	61
10.1 Primitives and the economic ratio	61
10.2 Price discovery under coupling	62
11 The Smart Asset Virtual Machine: Specification	63
11.1 SAVM instruction set architecture (ISA)	63
11.2 Structural guarantees	64
11.3 Consequences of the SAVM specification	65
12 Open Problems	65
A Instrument-Class Prerequisite Ordering	67
A.1 Prerequisite ordering on instrument classes	67
A.2 Why the ordering matters theoretically	68

1 Introduction

Scope note on “graduation.” This paper uses the terms *temperature tier* and *tier advancement* throughout. They refer exclusively to a venue-internal mechanism: the advancement or demotion of an instrument across the exchange’s Cold, Warm, and Hot tiers, computed from the instrument’s composed compliance tensor and used to control which CONVEXPOTENTIAL implementations the clearing engine admits and which obligation-pack predicates the kernel evaluates before admitting a participant to the order book. Tier advancement does *not* change the legal classification of an instrument, and is formally distinct from *regulatory graduation* (Definition 4.8), the jurisdiction-specific statutory process by which a sovereign authorizes an entity to issue. Whether a revenue-linked note, a sukuk certificate, an experiential token, a binary event contract, or any other instrument is a “security,” “derivative,” “swap,” “collective investment scheme,” “utility token,” “event contract,” “gaming product,” or an unregulated obligation is determined by the substantive law of each jurisdiction in which the instrument is offered or traded: under the Howey test and Delaware-law note and stock doctrines in the United States, under MiFID II and MiCA in the European Union, under AAOIFI-governed Islamic capital-market frameworks in Malaysia and the Gulf states, under insurance law for parametric contracts, and under the analogous frameworks of every other jurisdiction the paper does not enumerate. Each such determination requires jurisdiction-specific legal analysis by qualified counsel. Nothing in this paper constitutes legal advice or a legal classification in any jurisdiction.

1.1 Three under-served markets

Three asset classes totaling over \$375 billion in annual transaction volume lack a common infrastructure for issuance, cross-border trading with statute-specific enforcement, and programmable settlement.

IP licensing. Global intellectual property licensing generates \$165-200 billion annually [1]. Sports IP alone accounts for \$35-36 billion. The licensing process is manual: a merchandise company identifies an IP holder, negotiates terms by email, signs a PDF contract, wires payment through correspondent banking, and reports royalties quarterly via spreadsheet. There is no secondary market for licensing rights. There is no widely adopted market in which the revenue stream of a sports brand, a music catalogue, or a pharmaceutical patent can be issued, transferred, and settled with jurisdiction-specific eligibility checks. Less than 0.01% of global IP licensing volume touches any tokenized infrastructure, and the fraction that does remains confined to narrow single-jurisdiction wrappers.

Islamic finance. Sukuk (Islamic bonds) represent over \$200 billion in annual issuance and \$900 billion in outstanding instruments [2]. Sukuk must comply with Sharia principles: the instrument must be asset-backed (not debt-based), returns must derive from the performance of an identified real asset, and the structure must avoid *riba* (interest), *gharar* (excessive uncertainty), and *maysir* (speculation). These constraints make sukuk structurally compatible with programmable issuance whose compliance state is evaluated per instrument rather than per venue. Yet sukuk issuance remains confined to institutional channels: sovereign issuers, GCC banks, and a handful of multinational corporates. There is no retail participation pathway, no secondary market with sub-day settlement, and no programmable infrastructure that can verify the structural Sharia constraints (asset-backing, no interest, proportional distribution) computationally. The substantive determination of Sharia conformance remains a human judgment requiring

SSB certification.

Gaming and digital IP. The global gaming market generates over \$180 billion annually [3], with in-game economies producing over \$10 billion per year in virtual asset transactions. These economies are closed: a sword in one game cannot be traded for a character in another. The underlying IP (character designs, narrative arcs, procedurally generated content) has economic value that is trapped inside proprietary platforms. No infrastructure exists to register these assets as legally enforceable property, license them across platforms, and clear instruments linked to their revenue streams under instrument-specific statutory gates.

1.2 The missing infrastructure

The common obstacle across all three markets is institutional infrastructure. Digitizing a claim is not enough. A programmable security requires: (1) a legal entity to hold the underlying asset or risk, (2) a compliance framework that composes the regulatory requirements of every jurisdiction in which the instrument will trade, (3) a settlement system that can clear trades atomically with event resolution, and (4) a revenue or payout distribution mechanism that can split cash flows across jurisdictions with correct withholding or supervisory treatment.

No single system provides all four. Existing blockchain exchanges provide settlement; legal entity formation and regulatory composition remain outside their scope. Traditional financial infrastructure provides compliance inside existing stacks; programmable settlement and cross-jurisdictional portability remain outside their scope. Legal service providers form entities; cross-jurisdictional compliance composition and secondary clearing remain outside their scope.

The system described in this paper is the combination of two purpose-built components:

- **Mass:** a proof-producing institutional kernel that creates and operates multi-harbored entities, composing compliance across jurisdictions via a 23-domain compliance tensor with pointwise meet on the Applicable fragment [8, 9]. Mass is the issuance protocol: it is authoritative for legal state, corridor state, and regulatory attestation state.
- **Moxie:** a sovereign L1 clearing venue with event-atomic settlement, temperature-graduated market structure, and a self-calibrating coupling between prediction markets and spot prices [5, 6] building on the contingent-claim convex-auction lineage of Lange 1999 [67], Lange-Economides 2005 [68], Peters-So-Ye 2006 [69], and Baron-Lange 2007 [70]. Moxie is the secondary venue: it clears only after the Mass-side gate admits the trade.

Together, the clearing layer handles secondary trading while the institutional kernel provides entity management, issuance, and compliance composition. An issuer forms a multi-harbored entity through Mass, registers the relevant asset or contractual right, publishes machine-readable terms, and lists the instrument for Moxie clearing once the instrument's tensor and bridge envelope satisfy the venue gate. The compliance tensor ensures that only participants satisfying the composed constraint surface of the relevant jurisdictions can trade. Revenue or payout flows through Mass, splits according to configured rules, and distributes with per-jurisdiction withholding or supervisory treatment computed deterministically.

The failure of prior security-token systems clarifies the design requirement. Platforms such as Polymath, Securitize, Provenance, and adjacent tokenization stacks digitized transfer restrictions and cap-table operations inside a single-jurisdiction legal shell. Cross-border secondary liquidity failed to scale. The composed compliance tensor was the missing object: when one instrument touches multiple jurisdictions, the binding state is

the meet of instrument-specific statutory requirements, corridor recognition, and investor-specific eligibility. A single-jurisdiction token with a global transport layer does not solve that problem; it exports the conflict instead of composing it.

1.3 Legal classes and representative tensor profiles

The term *programmable security* is architectural rather than jurisdictional. Under existing frameworks, the same architecture can instantiate several familiar legal classes. A revenue-linked note maps to a note, debenture, transferable security, or investment contract depending on jurisdiction and offering structure. A tranching structured-credit instrument maps to debt securities or participation interests. A sukuk maps to certificates representing beneficial ownership under AAOIFI-style Islamic capital-market frameworks together with the local securities law that admits their issuance and trading. A sovereign parametric insurance instrument maps to an insurance-linked security or a regulated insurance contract depending on who bears the risk and who may hold the paper. The architecture does not erase these categories; it computes the compliance state they require.

Representative tensor profiles illustrate the point.

- **Cross-jurisdictional sukuk.** Applicable domains include AML, KYC, Sanctions, Tax, Securities, Corporate, Custody, Licensing, Settlement, Trade, and Sharia. Hot graduation requires each applicable domain to be compliant and every remaining domain to be not applicable. Sharia is evaluated per instrument: a sukuk can reach Hot while another instrument in the same block remains non-compliant on the Sharia coordinate.
- **Sovereign parametric insurance.** Applicable domains include AML, KYC, Insurance, Securities, Tax, Settlement, DataPrivacy, and AntiBribery, with Sharia either compliant or not applicable depending on issuance structure. Employment, Immigration, and IP are usually not applicable. The binding question is whether the instrument is admitted as an insurance-linked security, a reinsurance participation, or a regulated insurance contract.
- **Event-contingent royalty instrument.** Applicable domains include AML, KYC, Securities, Tax, Licensing, Settlement, IP, ConsumerProtection, and Arbitration. Insurance and Sharia are not applicable unless added by deal structure. Warm graduation admits only gated pools whose investors satisfy the note's statutory and corridor-specific constraints.
- **Tranched structured credit.** Applicable domains include AML, KYC, Securities, Tax, Corporate, Custody, Settlement, DataPrivacy, and AntiBribery, with Insurance applicable if the tranche embeds a guarantee and Sharia applicable only for Islamic tranches. Hot graduation requires the full cross-domain surface to resolve to compliant or not applicable on all 23 coordinates before the tranche can clear on the open secondary venue.

1.4 Contributions and outline

1. **Formal foundations** (Section 2). An explicit definition of a *programmable security* (Definition 2.6) as a gas-bounded computable map on a typed commitment state; a composition theorem (Theorem 2.11) establishing that programmable securities form a symmetric monoidal category under a non-interference-gated sequential/parallel composition; a soundness theorem (Theorem 2.12) connecting the operational behaviour of compiled bytecode to the declarative contract semantics; a life-cycle automaton for each instrument class (Definition 2.16); four first-order invariants on the commitment state (conservation, monotonicity, jurisdictional closure, temporal monotonicity) with an invariant-preservation theorem (Theorem 2.19); and

a regulatory-equivalence correspondence (Proposition 2.14) by which composition preserves the regulatory category of constituents under stated jurisdictional assumptions.

2. **Instrument taxonomy** (Section 3). Six instrument classes covering prediction markets, structured debt, index products, options, sukuk, and experiential tokens, with formal specifications, life-cycle states, and indicative yield calculations for each.
3. **The Mass connection** (Section 4). The multi-harbored entity as the bridge between exchange and kernel; the composition of compliance constraints across jurisdictions; the one-way implication (Proposition 4.9) by which regulatory graduation is a necessary input to the operational temperature tier, with no converse implication.
4. **SAVM, obligation-pack compiler, and ConvexPotential bridge** (Section 5). A stack-based deterministic VM over FP128 fixed-point arithmetic; a compliance-gated ISA with statically-disallowed reentrancy, bounded call-graph depth, and a gas-metered cost model; a determinism/halting theorem (Theorem 5.6); a LIFECYCLE_TRANSITION opcode; five instrument-specific potentials (equity, sukuk, convertible note, fixed income, option-like); and the tier-classification map as a meet-preserving lattice morphism.
5. **Corporate action semantics** (Section 5.6). A formal corporate-action object (Definition 5.30) with effect rules for split, cash dividend, stock dividend, merger, rights offering, tender offer, conversion, and redemption; a well-definedness theorem (Theorem 5.32) proving that each rule preserves the instrument invariants, the CONVEXPOTENTIAL requirements, and the state-machine transition relation.
6. **Regulatory classification** (Section 6). Classification of each instrument type across US, EU, ADGM, and Seychelles frameworks; a market-phase-extended statutory threshold vector $\tau^{\text{reg}(I)}(\phi, \text{phase}, h)$ covering Reg D 506(c), Reg A+, Rule 144 seasoning, Rule 144A QIB transfer, and Reg S Category 2/3 (Definition 6.1); the primary-secondary gate soundness theorem (Proposition 6.2); and three misclassification remedy classes with well-definedness (Theorem 6.3).
7. **The self-improving exchange** (Section 7). The exchange's operational intelligence layer; a contraction result for the self-calibrating transition operator; and a quantification of the information value of coupling relative to loss-versus-rebalancing (LVR).
8. **Price discovery and market microstructure** (Sections 10-8). A closed-form economic multiplier (Proposition 10.3) with upper and lower bounds under regularity (Theorem 10.4); a Hasbrouck-information-share characterisation of the prediction market's contribution (Proposition 10.6); and a reading of the coupling in Glosten-Milgrom, Kyle, and Jensen-Meckling terms.
9. **Instrument-class prerequisite structure** (Appendix A). A transitive prerequisite ordering on the instrument taxonomy induced by compliance-domain requirements and temperature-tier thresholds.

2 Formal Foundations: What Is a Programmable Security?

Before developing the instrument taxonomy (Section 3) or the compliance-lattice and interface-dispatch constructions (Sections 4-5), we give an explicit, formal answer to the question in the section title. The goal of this section is not novelty; programmable contract

formalisms exist in the literature (Kasprzyk-Egelund-Müller et al. on contract-as-code [63], the Marlowe financial-contract DSL [64], the ERC-3643 [30] standard on security-token semantics). The goal is to fix a single reference definition that the rest of the paper quotes, composes, and proves theorems against. The definition, composition theorem, and soundness theorem below are deliberately narrow: they encode no specific jurisdictional statute, no specific Sharia rule, and no specific settlement-venue policy. Those enter as *parameters* of a programmable security (Definitions 2.5-2.6).

2.1 The commitment algebra

Definition 2.1 (Typed commitment alphabet). A *typed commitment alphabet* is a pair (Σ, τ) where Σ is a finite set of *commitment symbols* (each representing an atomic economic obligation: a transfer, a reserve change, a license grant, a disclosure, a compliance-domain attestation, ...) and $\tau : \Sigma \rightarrow \text{Type}$ assigns each symbol a type in a fixed type system that distinguishes *quantitative* commitments (those carrying an FP128 value in the sense of Definition 5.2), *relational* commitments (those referring to pairs or tuples of entity identifiers), and *attestational* commitments (those referring to compliance-lattice grades in the sense of Definition 5.25). The alphabet is fixed throughout the paper; we write Σ_{Moxie} when the specific Moxie-instance alphabet is meant.

Definition 2.2 (State space). The *state space* of a programmable security is a finite-dimensional module $\Sigma^\# := \mathbb{R}_{\geq 0}^{|\Sigma|} \times \mathcal{L}_{23} \times \text{States}(I) \times \mathbb{R}_{\geq 0}$ whose first factor records the aggregate quantity of each commitment symbol issued to date, whose second factor records the current 23-domain compliance-lattice state (Section 4; Definition 5.25), whose third factor records the current life-cycle state (Definition 2.16), and whose fourth factor records the kernel clock t . A single state is a quadruple $s = (\mathbf{q}, c, \ell, t)$.

Definition 2.3 (Holder and attestation projections). For each participant address h , let $q_h(s)$ denote the projection of $\mathbf{q}(s)$ onto the commitment symbols attributable to h , and let $\phi(h) \in \mathcal{J}$ denote the jurisdiction of residence. The *attestation set* $\text{Attested}(s, h) \subseteq \mathcal{J}$ is the subset of jurisdictions for which h carries valid KYC attestations on the compliance lattice coordinate of s . These projections are the data on which the invariants of Definition 2.18 are stated.

Definition 2.4 (Commitment algebra). The *commitment algebra* $\mathcal{C} := (\Sigma^\#, +, \perp)$ is the commutative monoid on the state space under componentwise addition of the first factor and pointwise meet of the second factor, with identity $\perp := (\mathbf{0}, \top_{\mathcal{L}_{23}})$ (zero commitments issued, maximum compliance grades available). Well-definedness of the meet is Lemma 5.27; commutativity and associativity of $+$ in the first factor follow from the module structure. The commitment algebra is the natural target for the denotational semantics of a programmable security: every economic action a security takes is an element of \mathcal{C} .

2.2 Programmable securities: formal definition

Definition 2.5 (Parameters of a programmable security). A *programmable-security parameter tuple* is a quadruple $(\Sigma, \tau, D, g_{\max})$ where (Σ, τ) is a typed commitment alphabet (Definition 2.1), $D \subseteq \{1, \dots, 23\}$ is the activated compliance-domain subset (as in Section 3, the per-instrument activation list), and $g_{\max} \in \mathbb{N}$ is a static gas bound in SAVM cost-units (Definition 5.5). The tuple is public: the regulatory interpretation of a programmable security is invariant under permutations of Σ that preserve τ and D .

Definition 2.6 (Programmable security). A *programmable security* with parameters $(\Sigma, \tau, D, g_{\max})$ is a quintuple $(\text{contract}, \text{compiled}, \mathcal{A}_I, \pi_{\text{CA}}, \pi_{\text{inv}})$ where:

1. $\text{contract} : \Sigma^\# \times \text{Inputs} \rightarrow \Sigma^\#$ is a *total recursive function* (equivalently, a total computable function in the sense of Turing [56]) from the product of the state space and a finite input alphabet Inputs to the state space. The input alphabet partitions as

$$\text{Inputs} = \text{TradeReq} \sqcup \text{TimeAdv} \sqcup \text{Attest} \sqcup \text{EventRes} \sqcup \text{CorpAct},$$

with each TradeReq carrying a market-phase tag $\in \{\text{primary}, \text{secondary}\}$ and a per-unit holding-period $h \in \mathbb{R}_{\geq 0}$. Totality means that on every well-typed input, contract terminates and returns a state in $\Sigma^\#$; ill-typed inputs are rejected at the boundary.

2. $\text{compiled} \in \text{SAVM_Bytecode}$ is a SAVM bytecode realisation of contract (Definition 5.7) such that for all $s \in \Sigma^\#$ and $\text{input} \in \text{Inputs}$,

$$\llbracket \text{compiled} \rrbracket_{\text{SAVM}}(s, \text{input}) = \text{contract}(s, \text{input})$$

(conclusion of Theorem 5.8), and such that every (s, input) pair executes within the gas budget g_{\max} (enforced statically; Definition 5.7, pass 5).

3. \mathcal{A}_I is the life-cycle automaton (Definition 2.16) whose transition relation Δ_I every contract-induced state change respects.
4. π_{CA} is the corporate-action effect-rule function (Definition 5.31) that fires on CorpAct inputs.
5. π_{inv} is the invariant predicate (Definition 2.18) that every reachable state satisfies.

The set of all programmable securities with fixed $(\Sigma, \tau, D, g_{\max})$ is denoted $\text{PS}(\Sigma, \tau, D, g_{\max})$; the union over parameter tuples is PS . A programmable security is “programmable” in the precise sense that contract is a terminating computable function rather than a human-readable legal text, and “a security” in the parameterised sense that its activated compliance-domain subset D and the compliance grades it references situate it within the compliance lattice. The definition does *not* assert that the object is a security under any specific jurisdiction’s statute; that is the regulatory classification question of Section 6 and Proposition 2.14.

Remark 2.7 (Scope of the “programmable” predicate). The totality requirement on contract is load-bearing and is the reason the SAVM ISA (Definition 5.2) is a bounded-resource model rather than a Turing-complete machine. A non-total contract would admit programs that loop forever on some inputs; such programs could not be compiled to gas-bounded SAVM bytecode, could not be traded safely in a block-timed clearing venue, and could not be composed into larger contracts without inheriting the non-termination. The gas bound g_{\max} turns totality into a static property: a program that does not fit in g_{\max} gas is rejected at compile time, and a program that does is guaranteed by Theorem 5.6 to terminate. The restriction is narrower than Marlowe [64] (which admits inductively-defined contracts of unbounded depth) but sufficient for Section 3: all six instrument classes compile to $g_{\max} \leq 10^6$ gas per dispatch on the static call graph.

Remark 2.8 (Perpetual instruments via per-period invocation). The totality requirement excludes some economic structures that naturally admit unbounded iteration. Perpetual preferred equity, for instance, has no maturity and an unbounded coupon schedule. Such instruments are modelled as one SAVM contract invocation per coupon period, coordinated by the kernel’s external state machine. Each invocation is individually total and bounded; the unbounded iteration is handled by the kernel’s block-by-block evaluation loop, not by a loop inside the compiled bytecode.

2.3 Composition

Programmable securities must compose: a multi-instrument entity (e.g., an SPV issuing both a sukuk and an experiential token, or a CIS holding a basket of brand tokens) holds a

program that is, in the limit, the composition of the constituent programs. A composition law must specify (i) how the state spaces combine, (ii) when two programs may be run in parallel without interfering with each other's compliance invariants, and (iii) when sequential composition preserves the quasi-concavity and gas-budget structure of the constituents.

Definition 2.9 (Non-interference predicate). Two programmable securities $P_1, P_2 \in \text{PS}$ with activated domain sets $D(P_1), D(P_2) \subseteq \{1, \dots, 23\}$ are *non-interfering* if either (i) $D(P_1) \cap D(P_2) = \emptyset$ (disjoint compliance domains), or (ii) for every $d \in D(P_1) \cap D(P_2)$, the per-domain update rules of P_1 and P_2 on the d -coordinate of the compliance-lattice state agree on the intersection of their input domains (i.e., neither program can drive the d -coordinate below a grade that the other requires). Non-interference is a pure predicate on the pair (P_1, P_2) that is decidable in polynomial time in the sizes of $|D(P_1)|, |D(P_2)|$ and the sizes of the underlying finite grade sets \mathcal{G}_d .

Definition 2.10 (Sequential and parallel composition). Given programmable securities $P_1 = (\text{contract}_1, \text{compiled}_1)$ and $P_2 = (\text{contract}_2, \text{compiled}_2)$ with parameter tuples $(\Sigma_1, \tau_1, D_1, g_{\max,1})$ and $(\Sigma_2, \tau_2, D_2, g_{\max,2})$, their *sequential composition* $P_1 ; P_2$ has parameter tuple $(\Sigma_1 \sqcup \Sigma_2, \tau_1 \sqcup \tau_2, D_1 \cup D_2, g_{\max,1} + g_{\max,2})$ and contract function $(\text{contract}_1 ; \text{contract}_2)(s, \text{input}) := \text{contract}_2(\text{contract}_1(s, \text{input}))$. Their *parallel composition* $P_1 \parallel P_2$ has the same parameter tuple but contract function $(\text{contract}_1 \parallel \text{contract}_2)(s, (\text{input}_1, \text{input}_2)) := \text{contract}_1(s, \text{input}_1) + \text{contract}_2(s, \text{input}_2)$ in the commitment monoid of Definition 2.4. Both compositions are defined only when P_1 and P_2 are non-interfering (Definition 2.9).

Theorem 2.11 (Composition theorem). *The class PS of programmable securities, equipped with sequential composition $;$, parallel composition \parallel , and the non-interference predicate \perp_{NI} of Definition 2.9, is a symmetric monoidal category with unit the null programmable security $\mathbf{1}_{\text{PS}} := (\text{id}_{\Sigma^\#}, \text{RET})$ (the identity contract that returns its input state unchanged). Concretely:*

1. (Associativity.) $(P_1 ; P_2) ; P_3 = P_1 ; (P_2 ; P_3)$ whenever all pairwise non-interference predicates hold.
2. (Unit law.) $\mathbf{1}_{\text{PS}} ; P = P = P ; \mathbf{1}_{\text{PS}}$ for every $P \in \text{PS}$.
3. (Symmetric monoidal structure on \parallel .) For non-interfering P_1, P_2 , the parallel composition satisfies $P_1 \parallel P_2 \cong P_2 \parallel P_1$ (up to the canonical braiding on $\Sigma_1 \sqcup \Sigma_2$), $\mathbf{1}_{\text{PS}} \parallel P = P = P \parallel \mathbf{1}_{\text{PS}}$, and $(P_1 \parallel P_2) \parallel P_3 = P_1 \parallel (P_2 \parallel P_3)$.
4. (Preservation of totality and gas bound.) If P_1, P_2 are total with gas bounds $g_{\max,1}, g_{\max,2}$, then $P_1 ; P_2$ and $P_1 \parallel P_2$ are total with gas bound $g_{\max,1} + g_{\max,2}$ (Definition 2.10).
5. (Preservation of the CONVEXPOTENTIAL interface for AMM-settled constituents.) If P_1, P_2 both expose CONVEXPOTENTIAL surfaces φ_1, φ_2 satisfying (C1)-(C2) (Definition 5.9) and their state spaces commute (act on disjoint reserve coordinates), then $P_1 \parallel P_2$ exposes the product surface $\varphi_1 \cdot \varphi_2$, which satisfies (C1)-(C2) on the product reserve space.

Proof. Items (1)-(3) follow from the commutative-monoid structure of \mathcal{C} (Definition 2.4) and the associativity/commutativity of function composition restricted to pairs satisfying non-interference: non-interference ensures that contract_1 and contract_2 commute on their intersection of compliance-lattice coordinates (Definition 2.9, condition (ii)), and disjoint commitment-quantity supports commute by additivity of $+$ on the first factor. The unit law reduces to the observation that $\text{id}_{\Sigma^\#} \circ f = f = f \circ \text{id}_{\Sigma^\#}$ for any contract f . Item (4) is by construction of the composition's gas bound as the sum (Definition 2.10); termination follows from totality of each contract_i in at most $g_{\max,i}$ SAVM instructions (Theorem 5.6) and the gas-additivity of the SAVM cost model (Definition 5.5). Item (5) is immediate from the fact that a product of quasi-concave functions on disjoint coordinates is quasi-concave on the product, and monotonicity in each coordinate is inherited from the factors. \square

2.4 Soundness: operational semantics refines declarative contract

Theorem 2.12 (Soundness). *Let $P = (\text{contract}, \text{compiled}) \in \text{PS}$ be a programmable security with parameter tuple $(\Sigma, \tau, D, g_{\max})$, and let (s, input) be any admissible input pair. Let $\text{Behaviour}(P)(s, \text{input}) \subseteq \mathcal{C}^*$ denote the set of finite operational traces generated by executing compiled on (s, input) in the SAVM (Definition 5.2). Let $\text{Admissible}(P)(s, \text{input}) \subseteq \mathcal{C}^*$ denote the set of finite commitment-algebra traces that correspond to a single declarative evaluation of $\text{contract}(s, \text{input})$. Then*

$$\text{Behaviour}(P)(s, \text{input}) \subseteq \text{Admissible}(P)(s, \text{input}),$$

i.e., every observable execution trace of the compiled program is the trace of an admissible contract action. Equivalently, P 's operational behavior refines its declarative contract semantics in the sense of Hoare [57] and He-Hoare-Sanders [58].

Proof. By Theorem 5.8, $\llbracket \text{compiled} \rrbracket_{\text{SAVM}}(s, \text{input}) = \text{contract}(s, \text{input})$, so the end-to-end output of any execution trace coincides with the declarative output. For trace refinement (beyond end-to-end equivalence), each SAVM instruction corresponds to a single atomic step in the denotational semantics of the obligation-pack schema (Definition 5.7): arithmetic instructions realise pointwise commitment-monoid updates, COMPLIANCE_CHECK realises a lattice-state query, SSTORE realises a persistence update, and EVAL_OBLIGATION realises a sub-contract invocation. The composition of atomic admissible steps is admissible by induction on trace length, and gas-boundedness (Theorem 5.6) guarantees every trace is finite. The containment is strict when contract admits non-deterministic intermediate representations that compiled deterministically collapses; strictness is not load-bearing for the refinement conclusion. \square

Remark 2.13 (What soundness does and does not prove). Theorem 2.12 is a *refinement* in the traditional sense of program refinement [57]: the executable object does nothing *more* than the declarative object permits. It does not say the declarative object's behaviour is *desirable* (an economic-design question), nor that the declarative object correctly models any external legal concept (the classification question of Section 6). Refinement is the formal-methods analogue of "the program is not lying to the reader"; it is not the analogue of "the program is right." The latter, for a programmable security, requires the regulatory-equivalence correspondence below.

2.5 Regulatory equivalence under composition

Proposition 2.14 (Regulatory equivalence under non-interfering composition). *Let $P_1, P_2 \in \text{PS}$ be non-interfering programmable securities with activated domain sets $D(P_1), D(P_2)$ and regulatory classifications $\rho(P_1, \phi), \rho(P_2, \phi)$ in jurisdiction ϕ (in the sense of Table 2: security / derivative / CIS / utility / unregulated). Assume further:*

- (J1) **Jurisdictional monotonicity.** *For ϕ 's statute, the regulatory classification of a composed instrument depends only on the pointwise union of its constituents' activated domains: if $D(P_1) \cup D(P_2) = D(P_1)$ (equivalently $D(P_2) \subseteq D(P_1)$), then $\rho(P_1; P_2, \phi) = \rho(P_1, \phi)$.*
- (J2) **No reclassification of additive exposure.** *Increasing the aggregate quantity of a commitment symbol (via \parallel on the $+$ monoid) does not change the classification of the resulting exposure, conditional on the composed instrument remaining within jurisdictional size thresholds (\$75M for Reg A+, MiFID II debenture thresholds, etc.).*

Then:

$$\rho(P_1; P_2, \phi) = \rho(P_1, \phi) \sqcup_{\phi} \rho(P_2, \phi), \quad \rho(P_1 \parallel P_2, \phi) = \rho(P_1, \phi) \sqcup_{\phi} \rho(P_2, \phi),$$

where \sqcup_{ϕ} is the jurisdictional regulatory-category join (take the most restrictive of the two constituents' classifications).

Proof sketch. Under (J1), the activated-domain union of the composition determines the classification by pointwise dominance: the composition activates every domain that either constituent activates, and by monotonicity of ρ in D the composed classification is at least as restrictive as each constituent's. Under (J2), the exposure-additive nature of parallel composition does not change the category of the claim (a security composed with a security is still a security; a utility token composed with a security is a security by the more-restrictive rule). The join operation \sqcup_ϕ is the take-most-restrictive function on the partial order "security > derivative > CIS > utility > unregulated," which is a semi-lattice by inspection of the standard regulatory hierarchies (Howey progeny in the US, MiFID II category ladder in the EU, FSMR Specified Investment ladder in ADGM). The full lattice-theoretic argument is identical to the meet-preservation proof of Definition 5.24 applied to the regulatory-category lattice. \square

Remark 2.15 (Proposition 2.14 is informal in the precise sense of being about regulatory classification joins, not sovereign action). The assumptions (J1)-(J2) of Proposition 2.14 are descriptively satisfied by the four jurisdictional frameworks covered in Table 2. We cannot and do not claim that a sovereign regulator in any jurisdiction will follow the join rule \sqcup_ϕ when enforcing its statute. Regulators exercise discretion; novel compositions may prompt interpretive positions that diverge from the semi-lattice structure above. The proposition is a structural correspondence under stated assumptions, not a legal prediction. It provides a consistency check between the composition machinery of Theorem 2.11 and the way practitioners and counsel analyse multi-instrument products, and serves as a specification target for any jurisdiction whose regulator-facing guidance adopts the activated-domain-union convention explicitly. For the formal treatment of the venue-internal operational side of classification (the temperature tier, not the regulatory category), see Definitions 4.7-4.8 and the homomorphism proof attached to Definition 5.24.

2.6 The life-cycle state machine

Definition 2.6 presents contract as a total recursive function, which is correct at the denotational level but opaque operationally. A security is also a finite-state automaton whose states encode legal facts (unissued, subscribing, outstanding, matured, settled) and whose transitions fire on time, event resolution, user action, or corporate action (Section 5.6). We make this structure explicit.

Definition 2.16 (Life-cycle automaton). A *life-cycle automaton* for an instrument class I is a tuple $\mathcal{A}_I = (\text{States}(I), s_0, \Delta_I, \Lambda_I, \text{Accept}(I))$ where $\text{States}(I)$ is a finite set of life-cycle states, $s_0 \in \text{States}(I)$ is the initial state (typically UNISSUED), $\Delta_I \subseteq \text{States}(I) \times \text{Triggers} \times \text{States}(I)$ is the transition relation with triggers drawn from $\text{Triggers} = \{\text{time, event-resolution, user-action, corporate-}$

Definition 2.17 (LIFECYCLE_TRANSITION opcode). The SAVM ISA (Definition 5.2) is extended with the opcode LIFECYCLE_TRANSITION, which pops a trigger symbol and target state from the stack and either (i) executes the transition, updating the state-machine register and emitting an auditable transition log, or (ii) traps if the transition is not in Δ_I for the current state. The opcode gas cost is 20, matching SSTORE. The obligation-pack compiler emits LIFECYCLE_TRANSITION at each state-changing event in the declarative pack; static validation (Definition 5.7, pass 4) rejects any pack whose emitted transitions exit Δ_I .

For each instrument class in Section 3, the canonical state set and transition skeleton is:

- **Binary event contract:** UNISSUED \rightarrow SUBSCRIBING \rightarrow ACTIVE \rightarrow {RESOLVEDYES, RESOLVEDNO} \rightarrow SETTLED.
- **Revenue-linked note:** UNISSUED \rightarrow SUBSCRIBING \rightarrow OUTSTANDING \rightarrow COUPONACCRUING_t \rightarrow COUPONPAYABLE_t \rightarrow COUPONPAID_t \rightarrow {OUTSTANDING, CAPAPPROACHING} \rightarrow MATURED \rightarrow SETTLED.
- **IP index product:** UNISSUED \rightarrow CREATIONACTIVE \leftrightarrow REDEMPTIONACTIVE \rightarrow REBALANCEWINDOW \rightarrow DELISTED.
- **Contingent licensing option:** UNISSUED \rightarrow PREMIUMPAID \rightarrow ACTIVE \rightarrow {EXERCISED, EXPIRED} \rightarrow SETTLED.
- **Sukuk (ijara):** UNISSUED \rightarrow SPVFORMED \rightarrow ASSETTRANSFERRED \rightarrow SUBSCRIPTIONOPEN \rightarrow SUBSCRIPTIONCLOSED \rightarrow COUPONPERIOD_t \rightarrow PURCHASEUNDERTAKINGEXERCISED \rightarrow REDEEMED.
- **Experiential token:** MINTED \rightarrow SALEACTIVE \rightarrow TRANSFERABLE \rightarrow PRESENTEDFORREDEMPTION \rightarrow REDEEMEDANDBURNT.

Each transition guard is a conjunction of (i) a state predicate on the source state, (ii) the per-trade compliance predicate of Theorem 9.2 (condition S2), and (iii) a life-cycle predicate (S5) asserting that the requested transition belongs to Δ_I . Theorem 9.2 is extended to enforce all five conditions at the SAVM gate.

2.7 Instrument invariants

Totality (Definition 2.6) says contract terminates. It does not say what the state must preserve across calls. We state four first-order invariants on $\Sigma^\#$ that every contract \in PS preserves.

Definition 2.18 (Instrument invariants). For states $s, s' \in \Sigma^\#$ with $s' = \text{contract}(s, \text{input})$:

- (I1) **Conservation.** $\text{INV}_{\text{cons}}(s) := \sum_{h \in \text{Holders}} q_h(s) = \mathbf{q}_{\text{issued}}(s) - \mathbf{q}_{\text{burnt}}(s)$.
- (I2) **Monotonicity.** $\text{INV}_{\text{monot}}(s, s') := \forall h, q_h(s') = q_h(s) + \text{received}(h, s \rightarrow s') - \text{sent}(h, s \rightarrow s')$ with $\text{sent}, \text{received} \geq 0$.
- (I3) **Jurisdictional closure.** $\text{INV}_{\text{juris}}(s) := \forall h, \phi(h) \in \text{Attested}(s, h)$.
- (I4) **Temporal monotonicity.** $\text{INV}_{\text{time}}(s, s') := t(s') \geq t(s)$.

Theorem 2.19 (Invariant preservation). *Every contract \in PS preserves (I1)-(I4) of Definition 2.18.*

Proof sketch. (I1) follows from the additive structure of the commitment monoid (Definition 2.4) and the absence of a MINT_TO_ARBITRARY opcode in Definition 5.2. (I2) follows from the caller-save calling convention (Definition 5.3) which forbids writes to holder slots outside the explicit sent/received pair. (I3) follows from COMPLIANCE_CHECK being called on every transfer edge (Theorem 9.2, condition S4). (I4) follows from LOAD_CLOCK being monotone per block (Definition 5.2) and the absence of any opcode that writes the clock. \square

With these formal foundations in place, the remainder of the paper develops the specific instrument taxonomy (Section 3), the Mass-side institutional mechanics (Section 4), and the SAVM + obligation-pack compiler pipeline (Section 5) that together populate the abstract class PS with six concrete, compliance-parameterized, regulatory-classified instrument classes.

3 Instrument Taxonomy

We define six instrument classes. Each is specified by its payoff structure, the Mass kernel operations required for issuance, the compliance domains it activates, and the temperature tier at which it trades on a compliance-aware L1.

3.1 Binary event contracts

Definition 3.1 (Binary event contract). A binary event contract $\mathcal{B}(e, \delta, \tau)$ pays δ if event e occurs before time τ , and zero otherwise. The contract is issued at price $p \in (0, \delta)$ reflecting the market’s assessed probability $\hat{\pi} = p/\delta$ of the event occurring.

Examples:

- Fight outcome: “Fighter A defeats Fighter B at UFC 312.” $\delta = \$1$, $\tau =$ fight date.
- FDA approval: “Drug X receives FDA approval by December 2027.” $\delta = \$100$, $\tau =$ December 31, 2027.
- IP milestone: “Brand X exceeds \$10M cumulative licensing revenue.” $\delta = \$1$, $\tau =$ fiscal year end.

Mass kernel operations: None required for issuance of the contract itself. The event oracle must be registered as an attested data source through the Mass identity layer. If the contract references an IP asset held by a Mass-registered entity, the entity’s compliance state determines which jurisdictions’ participants can trade.

Temperature regime: Binary event contracts launch at Cold tier via LMSR with backstop parameter b . Cold here means pre-graduation: the instrument has not yet satisfied the full cross-domain gate required for open secondary clearing. The creator’s maximum loss is $b \ln 2$ per event. At \$500 backstop: maximum loss \approx \$346 per event.

Compliance domains activated: AML (domain 1, canonical index), KYC (domain 2), Securities (domain 5 - jurisdiction-dependent classification under Howey, MiFID II, or analogous statutes), Consumer Protection (domain 18), Arbitration (domain 19). Canonical domain indices follow the 23-domain enumeration of Definition 5.24.

Proposition 3.2 (LMSR pricing under event correlation). For M correlated binary events with interaction matrix J (the Ising interaction matrix from the companion paper [7]; see also Section 7 below for the self-calibrating estimation) and linear-bias vector $\mathbf{h} \in \mathbb{R}^M$ encoding per-event baseline log-odds, the LMSR cost function for a parlay trade [15] is:

$$C(\mathbf{q}) = b \ln \left(\sum_{\omega \in \{0,1\}^M} \exp\left(\frac{\mathbf{q} \cdot \omega}{b}\right) \cdot \frac{\exp(\sum_k h_k \omega_k + \sum_{k<l} J_{kl} \omega_k \omega_l)}{Z_{\mathbf{h}, J}} \right) \quad (1)$$

where \mathbf{q} is the quantity vector across outcomes and $Z_{\mathbf{h}, J} = \sum_{\omega} \exp(\sum_k h_k \omega_k + \sum_{k<l} J_{kl} \omega_k \omega_l)$ normalizes the general Ising prior. When $h_k = 0$ for all k and $J = 0$, this reduces to the standard LMSR with uniform prior: $C(\mathbf{q}) = b \ln(\sum_{\omega} \exp(\mathbf{q} \cdot \omega / b)) - b \ln(2^M)$. When $J = 0$ but $\mathbf{h} \neq 0$, the prior has non-uniform marginals $\pi(\omega_k = 1) = e^{h_k} / (1 + e^{h_k})$, recovering the standard LMSR with an informative per-event prior.

Proof. The standard LMSR cost function is $C(\mathbf{q}) = b \ln(\sum_{\omega} \pi(\omega) \exp(\mathbf{q} \cdot \omega / b))$ where $\pi(\omega)$ is the prior over outcomes [14]. Setting the prior to the general Ising-Boltzmann distribution $\pi(\omega) = Z_{\mathbf{h}, J}^{-1} \exp(\sum_k h_k \omega_k + \sum_{k<l} J_{kl} \omega_k \omega_l)$ and substituting yields the stated formula. In the special case $h_k = 0, J = 0$, the prior is uniform over 2^M states ($\pi(\omega) = 2^{-M}$), recovering the standard LMSR. The h_k term is omitted from the symmetric Ising special case; retaining it is necessary whenever the event’s baseline probability is not $1/2$, which is the generic case in practice (one fighter is favored, one vaccine has higher pass probability, etc.). The companion Ising paper [7] develops the estimation of (h_k, J_{kl}) jointly from trading data. \square

3.2 Revenue-linked notes

Definition 3.3 (Revenue-linked note). A revenue-linked note $\mathcal{R}(A, r, T, c)$ is a debt instrument whose coupon payments are a function $r(\cdot)$ of the licensing revenue generated by IP

asset A over period T , subject to a cap c on total distributions. The note holder receives:

$$\text{Coupon}_t = \min(r(R_t^A), c - \sum_{s < t} \text{Coupon}_s) \quad (2)$$

where R_t^A is the licensing revenue of asset A in period t . When cumulative coupons reach c , the note matures.

Revenue-linked notes are structured debt. The revenue function $r(\cdot)$ can be linear ($r(R) = \alpha R$ for revenue share α), tiered ($r(R) = \alpha_1 R$ for $R \leq R^*$ and $r(R) = \alpha_2 R$ for $R > R^*$), or waterfall-structured (senior tranche paid first, mezzanine second, equity residual).

Mass kernel operations: The IP asset A must be registered as property of a Mass-formed entity. The entity must have a treasury account through which licensing revenue flows. The kernel’s fiscal module computes R_t^A deterministically from on-chain revenue records. The coupon computation is a kernel operation: the entity’s configured distribution rules determine $r(\cdot)$, and the kernel produces a compliance proof that the distribution satisfies all jurisdictional constraints.

Compliance domains activated: AML (1), KYC (2), Tax (4), Securities (5), Custody (7), Licensing (9), Settlement (13), Consumer Protection (18). Canonical indices per Definition 5.24; the instrument’s Warm-tier admission additionally requires investor-accreditation attestations attached to the KYC domain grade per the “accredited” component grade of \mathcal{G}_{KYC} .

Remark 3.4 (Revenue-linked notes are securities). In representative frameworks, a revenue-linked note is ordinarily analyzed as a security: an investment contract or note in the United States, a transferable security or note under MiFID II-style frameworks in the European Union, and a debt-like capital-markets product in other activity-based securities regimes. The operational consequence is straightforward: secondary transfer requires the instrument-specific statutory gate in addition to possession of the tokenized wrapper.

Temperature regime: Revenue-linked notes enter Warm tier once the applicable securities, tax, custody, licensing, and settlement coordinates satisfy the instrument-specific admission threshold. They reach Hot tier only when every one of the 23 coordinates is either compliant or not applicable, at which point open secondary clearing is admissible.

3.3 IP index products

Definition 3.5 (IP index product). An IP index product $\mathcal{I}(\{A_k, w_k\}_{k=1}^K)$ is a basket instrument whose value tracks a weighted portfolio of K IP-linked tokens:

$$V_{\mathcal{I}}(t) = \sum_{k=1}^K w_k \cdot P_{A_k}(t) \quad (3)$$

where $P_{A_k}(t)$ is the price of the token linked to IP asset A_k and w_k is the portfolio weight, with $\sum_k w_k = 1$.

Examples:

- “Sports 50”: equal-weighted basket of the 50 highest-volume sports brand tokens on a given compliance-aware L1.
- “Combat Index”: market-cap-weighted basket of all combat sports brand tokens.
- “Gulf IP Index”: basket of all IP tokens issued by entities harbored in GCC jurisdictions.

The index product is ETF-like in structure. Creation and redemption operate through an authorized participant mechanism: the AP deposits the underlying basket of tokens and receives index tokens; redemption reverses the process. The arbitrage between index price and NAV keeps the index tracking tight.

Mass kernel operations: Each underlying IP asset must be held by a Mass-registered entity. The index itself is issued by a separate entity (the index fund vehicle), which must be harbored in at least one jurisdiction that permits collective investment schemes. The compliance lattice composes constraints across the index fund vehicle’s jurisdiction and all underlying entities’ jurisdictions.

Compliance domains activated: AML (1), KYC (2), Tax (4), Securities (5), Corporate (6, CIS-vehicle incorporation), Custody (7), Licensing (9), Settlement (13). Canonical indices per Definition 5.24.

Temperature regime: Hot tier. An index product belongs to Hot only once the fund vehicle and each underlying constituent induce a composed tensor whose 23 coordinates are all compliant or not applicable.

3.4 Contingent licensing options

Definition 3.6 (Contingent licensing option). A contingent licensing option $\mathcal{O}(A, \Lambda, \phi, \tau, K)$ grants the holder the right, without obligation, to license IP asset A under terms Λ in jurisdiction ϕ , exercisable before time τ , at strike price K . The payoff at exercise is:

$$\text{Payoff} = \max(V_{\Lambda}(A, \phi) - K, 0) \quad (4)$$

where $V_{\Lambda}(A, \phi)$ is the fair value of the license Λ for asset A in jurisdiction ϕ at the time of exercise.

This is a call option on a licensing right. The holder pays the option premium upfront. If conditions develop favorably (the brand gains popularity in jurisdiction ϕ , a regulatory change makes licensing more valuable, a competitor exits the market), the holder exercises the option, acquires the license at strike K , and captures the spread.

Mass kernel operations: The IP asset must have machine-readable license terms published through the Mass kernel. The option contract references specific license parameters Λ (territory, exclusivity, duration, revenue share, permitted uses). Exercise triggers a kernel operation: the holder’s entity must satisfy the compliance surface for the target jurisdiction ϕ , and the kernel produces a compliance proof before the license is granted. If the holder’s entity does not satisfy the compliance surface, the option expires worthless regardless of the fair value.

Compliance domains activated: AML (1), KYC (2), Securities (5), Licensing (9), Settlement (13), Ip (17). Canonical indices per Definition 5.24; derivative-style treatment in jurisdictions that classify the option as a derivative is applied via the Securities domain’s jurisdictional interpretation rather than a separate derivative domain in the 23-domain enumeration.

Temperature regime: Warm tier at admission; Hot only after the full 23-domain tensor resolves to compliant or not applicable on every coordinate.

3.5 Sukuk

Definition 3.7 (IP-backed sukuk). An IP-backed sukuk $\mathcal{K}(A, S, T, \pi)$ is a Sharia-compliant instrument structured as follows:

1. An originator transfers beneficial ownership of IP asset A to a special purpose vehicle (SPV).
2. The SPV issues sukuk certificates to investors, each representing an undivided proportional ownership interest in the IP asset.

3. The SPV licenses the IP asset back to the originator (or to third parties) under terms S .
4. Licensing revenue, net of SPV operating costs, distributes to certificate holders proportionally.
5. At maturity T , the originator repurchases the IP asset at pre-agreed price π (the “purchase undertaking”), and sukuk certificates are redeemed.

The total investor return derives from the licensing revenue stream plus any capital gain $(\pi - P_0)$ where P_0 is the initial subscription price.

The sukuk structure is *ijara* (lease-based): investors own the asset; the originator leases it back; rental payments are the coupon. This structure is *designed* to satisfy three Sharia constraints, subject to SSB certification:

1. **Asset-backing:** the certificate represents ownership of a real IP asset, not a debt claim.
2. **No riba:** returns derive from licensing revenue (rental income), not interest on a loan.
3. **No gharar:** the asset is identified, the license terms are specified, and the revenue stream is observable on-chain.

Caveat. These are structural design properties, not Sharia rulings. Whether a specific instrument satisfies Sharia requirements is a substantive determination that can only be made by a qualified Sharia Supervisory Board (SSB). The structural properties above are necessary and insufficient conditions for Sharia compliance; the SSB must independently assess the instrument’s conformity with the principles of Islamic jurisprudence (*fiqh al-muamalat*), including matters of *tawarruq*, beneficial ownership transfer, and the permissibility of the underlying IP asset itself.

Mass kernel operations: The SPV is a Mass-formed entity, harbored in a jurisdiction that recognizes sukuk structures. The IP asset transfer is a kernel ownership operation. The licensing arrangement is a kernel-managed contract with revenue flowing through the SPV’s treasury account. The purchase undertaking is a kernel-enforced contingent transaction. Sharia is modeled as the 23rd first-class compliance domain (Definition 5.24) within the 23-domain compliance lattice, with structural constraints that are computationally verifiable: the kernel checks component constraints SH-01 (no interest accrues), SH-02 (contract terms specified), SH-03 (hedging purpose documented where applicable), and SH-04 (asset identified and halal). These computational checks verify structural conformity with Sharia design principles but do not constitute a Sharia ruling: component constraint SH-05 (SSB certification) requires an independent fatwa and cannot be replaced by algorithmic verification. The composed Sharia grade on \mathcal{G}_{23} is the pointwise meet over SH-01..SH-05 (Definition 5.25).

Compliance domains activated: AML (1), KYC (2), Tax (4), Securities (5), Corporate (6, SPV registration), Custody (7), Licensing (9), Settlement (13), Ip (17), Sharia (23). Canonical indices per Definition 5.24. The Sharia domain decomposes into five component constraints SH-01 (No Riba), SH-02 (No Gharar), SH-03 (No Maysir), SH-04 (Asset Backing), and SH-05 (SSB Certification). Warm admission requires the instrument’s applicable Sharia coordinates to satisfy the gated-pool threshold; Hot graduation requires the Sharia coordinate, like every other coordinate, to resolve to compliant or not applicable.

Temperature regime: Warm at admission, with investor access determined instrument by instrument. Hot follows only when all 23 coordinates are compliant or not applicable. This makes the Sharia coordinate load-bearing for sukuk: a sukuk with unresolved Sharia status cannot reach Hot, while a non-Sharia instrument in the same block can.

Proposition 3.8 (Sukuk yield under stochastic licensing revenue). *Let R_t be the licensing revenue in period t , modeled as geometric Brownian motion (GBM) with drift μ (annualized*

rate, dimensionless per-year) and volatility σ (annualized rate): $dR_t = \mu R_t dt + \sigma R_t dW_t$, with baseline R_0 in dollars-per-year. Under GBM, $\mathbb{E}[R_t] = R_0 e^{\mu t}$ and $\text{Var}(R_t) = R_0^2 e^{2\mu t} (e^{\sigma^2 t} - 1)$. The expected annualized yield of an ijara sukuk with baseline revenue R_0 , operating cost ratio c , initial subscription price P_0 , purchase undertaking π , and maturity T is:

$$y = \frac{(1-c) \cdot R_0 \cdot (e^{\mu T} - 1)}{P_0 \cdot \mu \cdot T} + \frac{\pi - P_0}{P_0 \cdot T}, \quad (5)$$

where the first term is the time-averaged coupon yield (amortized over T) and the second term is the annualized capital gain. The variance of the periodic coupon as a fraction of subscription price is:

$$\text{Var}\left(\frac{\text{Coupon}_t}{P_0}\right) = \frac{(1-c)^2 R_0^2 e^{2\mu t} (e^{\sigma^2 t} - 1)}{P_0^2}. \quad (6)$$

Equation (5) is linear in R_0/P_0 (the coupon-yield-at-issuance) and in $(\pi - P_0)/P_0/T$ (the annualized capital gain), while the risk scales in σ through $e^{\sigma^2 t} - 1$. In the small-drift limit $\mu \rightarrow 0$, the first term reduces to $(1-c)R_0/P_0$ (flat coupon yield). For an IP asset with baseline $R_0 = \$5\text{M}/\text{year}$ licensing revenue, operating-cost ratio $c = 0.10$, subscription price $P_0 = \$40\text{M}$, maturity $T = 5$ years, purchase undertaking $\pi = \$42\text{M}$, and drift $\mu \approx 0$ (flat-revenue baseline):

- Coupon yield $\approx (1 - 0.10) \cdot 5/40 = 0.9 \cdot 0.125 = 11.25\%$ per year;
- Capital gain $= (42 - 40)/(40 \cdot 5) = 2/200 = 1.00\%$ per year;
- Total expected yield $\approx 12.25\%$ per year.

For drift $\mu = 0.05$ (5% per-year real revenue growth), the coupon-yield first term becomes $(0.9 \cdot 5 \cdot (e^{0.25} - 1))/(40 \cdot 0.05 \cdot 5) \approx (0.9 \cdot 5 \cdot 0.2840)/10 \approx 12.78\%$, giving total expected yield $\approx 13.78\%$. These are illustrative calculations under idealised GBM assumptions, not projections of actual returns. Actual yields depend on realised licensing revenue, operating costs, counterparty risk, and market conditions.

Modelling caveat: GBM understates tail risk. IP licensing revenue is typically jump-driven (product launches, regulatory approvals, litigation outcomes) with long flat periods between jumps. GBM misprices the tail: the variance term $e^{\sigma^2 t} - 1$ understates realised variance for heavy-tailed revenue. A compound Poisson or jump-diffusion model (Merton [54]; Kou [55]) better matches the empirical distribution; the closed-form yield above extends with a jump-component correction $\lambda_j \mathbb{E}[e^Y - 1]T$ where λ_j is the jump intensity and Y the log-jump size. We present the GBM version as a first-order benchmark only.

Proof. Under the GBM model $dR_t = \mu R_t dt + \sigma R_t dW_t$, Ito's lemma gives $R_t = R_0 \exp((\mu - \sigma^2/2)t + \sigma W_t)$, and the expectation is $\mathbb{E}[R_t] = R_0 e^{\mu t}$ (the deterministic exponential-growth expectation, *not* μ directly). The coupon at time t is $\text{Coupon}_t = (1-c)R_t$, so $\mathbb{E}[\text{Coupon}_t] = (1-c)R_0 e^{\mu t}$ and $\text{Var}(\text{Coupon}_t) = (1-c)^2 R_0^2 e^{2\mu t} (e^{\sigma^2 t} - 1)$. The time-averaged expected coupon over $[0, T]$, normalized by P_0 , gives an annualized coupon yield

$$\bar{y}_{\text{coupon}} = \frac{1}{TP_0} \int_0^T \mathbb{E}[\text{Coupon}_t] dt = \frac{(1-c)R_0}{TP_0} \cdot \int_0^T e^{\mu t} dt = \frac{(1-c)R_0(e^{\mu T} - 1)}{\mu TP_0},$$

which is the first term of (5). Adding the annualized capital-gain term $(\pi - P_0)/(P_0 T)$ gives the total annualized yield.

For the numerical example with $\mu \approx 0$, the first term's L'Hôpital limit is $(1-c)R_0/P_0 = 0.9 \cdot 5/40 = 11.25\%$; the capital-gain term is $2/(40 \cdot 5) = 1.0\%$; total $\approx 12.25\%$. Units: R_0 carries units of dollars/year, P_0 carries units of dollars, T carries units of years, μ and σ carry units of per-year (so μT and $\sigma^2 T$ are dimensionless exponents), and the yield y is dimensionless per-year. Every term in (5) has consistent units. \square

3.6 Experiential tokens

Definition 3.9 (Experiential token). An experiential token $\mathcal{E}(A, X, \phi, \tau)$ grants the holder access to experience X associated with IP asset A in jurisdiction ϕ before time τ . Experiences include:

- Gate access to premium game content (exclusive levels, procedurally generated storylines).
- VIP event access (ringside seats, backstage, meet-and-greet).
- Physical merchandise redemption (limited edition, authenticated on-chain).
- Personalized content (custom narrative, character interaction).

The token is non-fungible in its specific experience but fungible within a class (all “UFC 312 VIP” tokens are interchangeable).

Experiential tokens are not securities. They do not represent an investment with expectation of profit from the efforts of others. They are access rights: the holder redeems the token for an experience, and the token is consumed. This classification holds in the US (outside Howey security classification), the EU (outside MiFID II financial-instrument classification; may be a “utility token” under MiCA), and Singapore (outside SFA capital-markets-product classification).

Mass kernel operations: The IP asset must be registered. The experience provider must be a Mass-registered entity with the operational capacity to deliver the experience. The kernel manages the redemption lifecycle: issuance, transfer, redemption verification, and destruction after use. Securities-style trading proof is unnecessary for ordinary access-token transfers. AML/KYC applies to the initial purchase and to any secondary sale above a threshold.

Temperature regime: Cold tier. Experiential tokens trade in batch auctions with LMSR backstop. The Cold tier is appropriate because these instruments ordinarily remain outside the full securities-style graduation path and therefore do not require open secondary clearing.

Table 1: Instrument taxonomy summary

Instrument	Payoff type	Temp.	Security?	Key compl
Binary event contract	Binary	Cold	Regulated (derivative/event; see Table 2)	Gambling
Revenue-linked note	Coupon stream	Warm	Yes	Securities,
IP index product	NAV tracking	Hot	Yes	CIS, Secur
Contingent license option	Call option	Warm	Yes	Derivative
Sukuk (ijara)	Revenue share	Warm/Hot	Yes	Sharia, Secu
Experiential token	Access right	Cold	No	AML/KYC

4 The Mass Connection

4.1 Multi-harbored entities

The core primitive of the Mass kernel is the multi-harboured entity: a legal entity existing simultaneously in multiple jurisdictions, with compliance state composed algebraically across all of them. The algebraic primitive (harbour set together with per-harbour compliance state, with composed state defined as the 23-domain pointwise meet) is formalised in the companion algebra paper as *the multi-harboured entity*; we recall it here and then extend it with the two additional components that the instrument-settlement arguments require: an entity identifier and a corridor set.

Definition 4.1 (Multi-harbored entity; algebraic primitive, companion volume). We take as given the *multi-harbored entity* as formalized in the companion volume [8, Definition 3.1, def:multi-harbored-entity]: a harbor set $\Phi = \{\phi_1, \dots, \phi_n\} \subseteq \mathcal{J}$ together with a per-harbor compliance-state map $\mathcal{C} : \Phi \rightarrow \mathcal{L}^{23}$ into the 23-domain tensor. In the independent-threshold regime, the conservative composed compliance state is the pointwise meet on the Applicable fragment, $\mathcal{C}(\mathcal{E}) = \bigwedge_{\phi \in \Phi} \mathcal{C}(\phi) \in \mathcal{L}^{23}$, well-defined because each per-domain grade set \mathcal{G}_d is independently totally ordered where applicable (Lemma 5.27 of this paper; Theorem 3.1 of [8]). The companion volume proves the Applicable-fragment meet-semilattice structure, the scoped Heyting results, and the mixed-axis dichotomy for the full tensor; we invoke only the results needed for composition and graduation.

Definition 4.2 (Instrument-bearing multi-harbored entity). An *instrument-bearing multi-harbored entity* \mathcal{E} extends the algebraic primitive of Definition 4.1 with the two additional components the instrument-settlement layer requires: an entity identifier and a corridor set. Formally, \mathcal{E} is a tuple $(e, \Phi, \mathcal{C}, \mathcal{K})$ where:

- e is the entity identifier (the on-chain and in-kernel handle used by the ownership module and by the instrument-issuance pipeline; the companion algebra volume suppresses this identifier because the pure-algebraic results do not depend on it).
- (Φ, \mathcal{C}) is a multi-harbored entity in the sense of Definition 4.1: $\Phi = \{\phi_1, \dots, \phi_n\}$ is the set of jurisdictions in which e is registered (its “harbors”), and $\mathcal{C} : \Phi \rightarrow \mathcal{L}^{23}$ maps each jurisdiction to a compliance state vector in the 23-domain compliance lattice \mathcal{L} .
- $\mathcal{K} = \{K_{ij}\}$ is the set of corridors between harbors, each parameterized by (R_{ij}, ϕ_{ij}) specifying which compliance domains carry forward between ϕ_i and ϕ_j and at what grade of mutual recognition. Corridors carry typed compliance state between sovereign kernels and are consumed by the settlement-gating predicate of Theorem 9.2 (condition S1). The corridor algebra is developed in [8, Theorem thm:corridor-compose].

Under the independent-threshold hypothesis (no priority conflict, treaty lift, or sovereign-specific override), the composed compliance state of \mathcal{E} is the pointwise meet across all harbors, inherited directly from Definition 4.1:

$$\mathcal{C}(\mathcal{E}) = \bigwedge_{\phi \in \Phi} \mathcal{C}(\phi) = \left(\min_{\leq 1} \phi c_1(\phi), \dots, \min_{\leq 23} \phi c_{23}(\phi) \right) \quad (7)$$

where $c_d(\phi) \in \mathcal{G}_d$ is the compliance grade in domain d (within the totally-ordered per-domain grade set \mathcal{G}_d ; see Definition 5.25 for the explicit grade chain of domain 23, Sharia Compliance), and the meet operation selects the most restrictive constraint for each domain under its own ordering. Pointwise meet is well-defined across all 23 heterogeneous domain types because each \mathcal{G}_d is independently totally ordered; no cross-domain comparability is required (Lemma 5.27; compare [8, Theorem 3.1]).

Remark 4.3 (Scope of this paper’s use). Throughout this paper, every reference to “the multi-harbored entity” as an object equipped with an identifier and corridor set refers to Definition 4.2. The pure-algebraic primitive of Definition 4.1 is sufficient for results that reason only over composed compliance state (e.g., the temperature-tier computation of Definition 4.7, the meet-monotonicity proofs of the companion algebra volume [8, Appendix on the Sharia grade chain, app:sharia and Proposition prop:sharia-projection-meet]); the instrument-bearing extension is required for results that route settlement across corridors or that name specific entity instances (Theorem 9.2 and Section 4 onward). The notation in the companion volume uses the pair (Φ, \mathcal{C}) with harbor set written as an indexed family $\{J_1, \dots, J_k\}$; our tuple $(e, \Phi, \mathcal{C}, \mathcal{K})$ specializes to the companion-volume pair by forgetting the identifier e and corridor set \mathcal{K} , i.e., (Φ, \mathcal{C}) alone is the underlying algebraic primitive.

Remark 4.4 (Scope of the lattice structure used in this paper). The compliance tensor used here is the Applicable fragment of the 23-fold product of per-domain grade chains (Definition 5.25). The operations used throughout this paper are pointwise meet for multi-harbor composition and pointwise comparison against instrument-specific thresholds. The pseudo-complement is not load-bearing here. The full mixed-axis object with explicit applicability state is governed by the F144 dichotomy: no total TensorValue-valued semilattice preserves both the NotApplicable/Exempt distinction and the compliance signal, so the information-preserving operator is a disjoint-sum-valued MeetResult. This paper does not claim a full-tensor Heyting structure; it uses the Applicable-fragment meet-semilattice only.

4.2 IP holding through multi-harbored entities

An IP holder who wants to list instruments on a compliance-aware L1 follows this sequence:

1. **Entity formation.** A request to the Mass kernel initiates entity formation in a chosen jurisdiction. The entity is a special purpose vehicle (SPV) whose sole purpose is to hold the IP asset and manage instruments linked to it. Formation requires government registrar approval in the chosen jurisdiction.
2. **Harbor acquisition.** The entity acquires additional harbors by registering in additional jurisdictions. Each new harbor adds a node to the entity’s compliance surface. A corridor is established between each pair of harbors, specifying mutual recognition parameters.
3. **IP registration.** The IP asset is registered as property of the entity through the kernel’s ownership module. The registration produces a cryptographic proof of ownership that is verifiable on-chain.
4. **License term publication.** The entity publishes machine-readable license terms: territory, exclusivity, duration, revenue share, permitted uses. These terms are kernel-managed and can be queried programmatically.
5. **Instrument issuance.** The entity issues instruments (from the taxonomy in Section 3) on a compliance-aware L1. Each instrument type activates specific compliance domains. The kernel computes the composed compliance surface and determines which participants in which jurisdictions can trade.
6. **Revenue distribution.** Licensing revenue flows through the entity’s treasury account, managed by the kernel’s fiscal module. The kernel computes per-jurisdiction withholding, splits revenue according to configured rules, and produces a proof of correct distribution.

4.3 Mass-Moxie bridge and temperature graduation

Definition 4.5 (Mass-Moxie bridge envelope). For an issued instrument I of entity \mathcal{E} , the bridge emits an authenticated envelope

$$\beta(I, \mathcal{E}) = (\text{BRIDGE_PROTOCOL_VERSION}, e, I, x(I, \mathcal{E}), \text{Tier}(I, \mathcal{E}), \sigma_{\text{Ed25519}}, \sigma_{\text{ML-DSA-65}})$$

where $\text{BRIDGE_PROTOCOL_VERSION} = 1$, e is the issuing-entity identifier, $x(I, \mathcal{E})$ is the mirrored 23-domain compliance tensor, $\text{Tier}(I, \mathcal{E}) \in \{\text{Cold}, \text{Warm}, \text{Hot}\}$ is the venue-internal graduation output, and $(\sigma_{\text{Ed25519}}, \sigma_{\text{ML-DSA-65}})$ is a hybrid classical/post-quantum signature pair authenticating the envelope. Mass is authoritative for $x(I, \mathcal{E})$; Moxie consumes the envelope for secondary clearing.

Remark 4.6 (Advisory-only asymmetry and per-instrument gating). The bridge is asymmetric by design. Mass is authoritative for issuance state, harbor state, corridor state, and the compliance tensor. Moxie is advisory only with respect to legal state: it may reject settlement when the gate fails, but it does not rewrite the Mass-side tensor. The gate is also per instrument rather than per chain. A single block can contain two instruments with disjoint profiles, for example a sukuk whose Sharia coordinate is compliant and an event contract whose Sharia coordinate is non-compliant; investor access is computed instrument by instrument from the corresponding envelope.

Definition 4.7 (Temperature tier; operational, tensor-computed). The *temperature tier* of an instrument-on-entity pair, written $\text{Tier}(I, \mathcal{E}) \in \mathcal{T} = \{\text{Cold}, \text{Warm}, \text{Hot}\}$, is a deterministic function of the composed 23-domain tensor mirrored by $\beta(I, \mathcal{E})$ and the instrument's activated domain set $D(I)$. The tier is recomputed whenever the bridge envelope changes and is a property of the instrument-specific compliance state. It is *not* a legal classification: no sovereign regulator issues or recognizes a temperature tier, and no instrument's statutory classification is determined by it.

Definition 4.8 (Regulatory graduation; jurisdiction-specific, statute-based). The *regulatory graduation* of an entity for an instrument class in a jurisdiction ϕ is the statutory process by which ϕ 's regulatory authority authorizes the entity to issue, distribute, or admit the instrument to trade under ϕ 's law. Concretely: qualification under a securities-exemption rule or registered offering pathway in the United States, admission to a regulated market or MTF under MiFID II-style frameworks in the European Union, approval under Islamic capital-markets regimes for sukuk, licensing under insurance law for parametric contracts, and analogous acts of any other sovereign. Regulatory graduation is per jurisdiction and per instrument class; it is not computed by the bridge.

Proposition 4.9 (One-way implication: regulatory graduation is an input, not a consequence). *Let I be an instrument with activated domain set $D(I) \subseteq \{1, \dots, 23\}$ and let \mathcal{E} be an issuing entity with mirrored tensor $x(I, \mathcal{E})$.*

1. (Necessity of regulatory graduation.) *A necessary condition for $\text{Tier}(I, \mathcal{E}) \in \{\text{Warm}, \text{Hot}\}$ is that the applicable statutory coordinates of $x(I, \mathcal{E})$ have been populated by valid sovereign or delegated attestations for the jurisdictions on whose surfaces the instrument will settle.*
2. (Non-sufficiency for legal classification.) *The converse does not hold: $\text{Tier}(I, \mathcal{E}) = \text{Warm}$ or Hot does not imply that the instrument is legally authorized in every relevant jurisdiction, because the bridge authenticates attestations and their composition, not the ultimate legal correctness of the attester's conclusion.*

Temperature tier is therefore an operational state that depends on regulatory graduation as an input but is not equivalent to it.

Proof. Statement (1) follows because Warm and Hot both require applicable domains to meet venue thresholds, and those thresholds are evaluated against the mirrored tensor carried by Definition 4.5. If the statutory coordinates are absent or sub-threshold, the instrument cannot rise above Cold. Statement (2) follows because a stale, erroneous, or fraudulently obtained attestation can still verify cryptographically and therefore enter the tensor; the sovereign's legal remedy remains external to the bridge. \square

Definition 4.10 (Operational tier-to-capability correspondence). The temperature tiers on a compliance-aware L1 map to trading capabilities at the venue level as follows. These bindings are venue-internal operational policy, not legal classifications.

- **Cold:** pre-graduation. At least one applicable domain for I fails the Warm threshold, or the bridge envelope is incomplete. The instrument may be issued, warehoused, or transferred only under the restricted paths encoded in its obligation pack. No open secondary clearing follows from Cold status.

- **Warm:** partial-compliance gated trading. Every domain in $D(I)$ meets the instrument-specific admission threshold, but at least one coordinate in the full 23-domain tensor remains applicable and below Hot. The instrument may trade only in gated pools whose participants satisfy the instrument-specific predicate.
- **Hot:** full graduation. Every one of the 23 coordinates is either compliant at the Hot threshold or not applicable. The instrument may clear on the open secondary venue subject to the instrument’s remaining market-structure rules.

Remark 4.11 (Tier recomputation is event-driven). Tier changes are event-driven rather than volume-driven in this paper. A new attestation, a corridor-recognition change, a legal-state update, or an obligation-pack change recomputes $\text{Tier}(I, \mathcal{E})$ for the affected instrument only. One instrument’s demotion does not contaminate another instrument in the same block.

4.4 Corridor economics

Each corridor between two harbors is a bilateral agreement parameterized by (R, ϕ) : which compliance domains carry forward ($R \subseteq \{1, \dots, 23\}$; see Definition 5.24 for the canonical enumeration and Definition 5.25 for per-domain grade sets) and at what grade of recognition ($\phi : R \rightarrow [0, 1]$, where $\phi(d) = 1$ means the receiving jurisdiction accepts the sending jurisdiction’s compliance determination in domain d at face value, and $\phi(d) = 0$ means no recognition). The term “mutual recognition” as used here refers to the corridor’s parameterized acceptance of compliance determinations between two specific zone operators. It is not equivalent to the formal mutual recognition agreements (MRAs) established between sovereign regulators under international law, though the corridor mechanism is designed to encode such MRAs where they exist.

Remark 4.12 (Domain-carrying asymmetry, with Sharia as the salient example). A corridor carries domain d forward only if $d \in R$ and $\phi(d) > 0$. For Sharia Compliance (domain 23) this asymmetry is expressive: an ADGM \leftrightarrow Saudi corridor carries $d = 23$ at full recognition ($\phi(23) = 1$); an ADGM \leftrightarrow Delaware corridor omits $d = 23$ from R (equivalently, sets $\phi(23) = 0$), reflecting Delaware’s non-recognition of Sharia as a statutory concept. An entity can still multi-harbor across such jurisdictions, but its composed Sharia grade is bottom (Remark 5.28), which operationally disqualifies its trades on the Delaware-membership surface from carrying Sharia-based tier privileges. The corridor mechanism thereby expresses the Sharia-recognition asymmetry structurally, without requiring the lattice itself to treat Sharia specially.

Corridors are asymmetric. What ADGM accepts from Seychelles differs from what Seychelles accepts from ADGM. The corridor graph is therefore directed: the edge (ϕ_i, ϕ_j) has different parameters from (ϕ_j, ϕ_i) .

Proposition 4.13 (Corridor network scaling). *Adding a new jurisdiction ϕ_{n+1} to a network of n jurisdictions requires establishing at most $2n$ directed corridors (one in each direction with each existing jurisdiction). In practice, corridors are established only with jurisdictions that have relevant bilateral relationships, so the number is $2k$ where $k \leq n$ is the number of corridor partners. The network grows linearly in k , not quadratically in n .*

Proof. Each corridor is a directed edge. Adding ϕ_{n+1} with k corridor partners creates $2k$ directed edges (one per direction per partner). The total number of corridors after $n + 1$ jurisdictions is at most $\sum_{j=1}^n 2k_j \leq 2n^2$, but for sparse corridor graphs where each jurisdiction partners with at most k others, the total is $O(nk)$, which is linear in n for constant k . \square

For Moxie instruments, the corridor network determines the tradeable set: a participant in jurisdiction ϕ_j can trade an instrument issued by an entity harbored in ϕ_i if and only if a corridor (ϕ_i, ϕ_j) exists and the corridor's recognition function ϕ covers all compliance domains activated by the instrument type.

Restatement: the temperature tier is venue-internal, not legal. *The temperature tier (Definition 4.7) manages exchange-internal admission: which instruments the clearing engine admits, which CONVEXPOTENTIAL implementations are permitted, and which obligations the kernel verifies before admitting a participant to an instrument-specific order book. It does not by itself effect legal registration, exemption, or authorization in any jurisdiction; that is the separate concept of regulatory graduation (Definition 4.8). An instrument's legal classification under securities law, Islamic capital-markets law, insurance law, or any analogous statute is jurisdiction-specific and determined by substance, not by the tier label the venue assigns. The rest of this paper relies on this restatement whenever it describes temperature tiers, compliance composition, or the meet-semilattice morphism of Definition 5.24: those are all statements about the instrument-specific operational compliance surface, never about the underlying legal character of any instrument.*

5 The Obligation-Pack Compiler and the ConvexPotential Bridge

The six instrument classes defined in Section 3 differ in payoff structure and in the shape of the liquidity surface on which each should trade. A constant-product AMM [17, 18] cannot express the asymmetric extraction cost of a sukuk above its profit-sharing threshold, or the smooth debt-to-equity transition of a convertible note; a flat bonding curve is the wrong surface for an equity token with high growth variance. Each instrument class is specified by a declarative *obligation pack* (Section 5.5), and the kernel's *obligation-pack compiler* translates that pack into a realization of the clearing-engine's abstract interface.

Remark 5.1 (Scope of the obligation-pack compiler). Section 11 gives the instruction set, calling convention, memory model, termination guarantee, and compilation-correctness theorem needed for the SAVM terminology used here. What the Mass kernel provides is (i) a declarative obligation-pack schema, (ii) an obligation-pack compiler from the schema to SAVM bytecode, (iii) the SAVM interpreter, and (iv) an auditable evaluation receipt. The SAVM name is shorthand for the combined (ISA, compiler, interpreter) layer.

5.1 The Smart Asset Virtual Machine (SAVM): formal specification

We specify the SAVM as a stack-based deterministic virtual machine with a small instruction set sufficient to encode the obligation-pack semantics of Section 5.5 and the five instrument potentials of Section 5.3. The specification follows the template of Wasm [60] (stack machine, typed values, structured control flow) and the EVM [61] (gas-metered execution, deterministic semantics), but with a narrower scope: only the operations needed for compliance-gated fixed-point arithmetic and state I/O are included.

Definition 5.2 (SAVM ISA). The SAVM operates over a stack \mathcal{S} of typed values $\mathcal{V} := \text{FP128} \cup \text{Bool} \cup \text{Addr}$, where FP128 is the signed fixed-point format with 64 integer bits and 64 fractional bits (precision $2^{-64} \approx 5.4 \times 10^{-20}$; exact arithmetic within the representable range $[-2^{63}, 2^{63} - 2^{-64}]$). The instruction set is the disjoint union of:

Arithmetic (ADD, SUB, MUL, DIV, NEG, ABS):

pop two FP128 operands, push one result. Overflow traps. DIV by zero traps. MUL and DIV use intermediate 256-bit precision to avoid premature rounding.

Comparison (LT, LE, EQ, GT, GE):

pop two FP128, push one Bool.

Control flow (IF, ELSE, END_IF, LOOP, BREAK, CONTINUE):

structured control flow following Wasm’s validation rules (no arbitrary jumps); LOOP must carry a static iteration bound, enforced at validation time.

Compliance gate (COMPLIANCE_CHECK):

pop a compliance-domain identifier and a jurisdiction identifier; return Bool indicating whether the current execution context passes the domain check in that jurisdiction (queries the kernel’s compliance lattice).

State I/O (SLOAD, SSTORE, RLOAD):

SLOAD pops Addr, pushes the stored FP128; SSTORE pops Addr, FP128; RLOAD reads the pool reserves ($R_{\text{hub}}, R_{\text{brand}}$) onto the stack.

Obligation evaluation (EVAL_OBLIGATION):

invokes a named sub-obligation within the same pack; sub-obligations are statically resolved at compile time (no dynamic dispatch).

Return (RET):

terminates execution with the top of stack as the result.

Definition 5.3 (SAVM memory model and calling convention). The SAVM memory consists of (i) a bounded evaluation stack \mathcal{S} of fixed maximum depth $D_{\text{max}} = 1024$; (ii) a read-only *obligation storage* \mathcal{M}_o containing the compiled bytecode and constant pool; (iii) a read-write *instrument state* \mathcal{M}_s containing per-instrument persistent variables (addresses $[0, 2^{32})$); (iv) a read-only *context* \mathcal{C} holding the current reserves, block number, timestamp, and caller identity. Calls into EVAL_OBLIGATION use a caller-save convention: the caller’s stack frame is preserved; the callee operates on a fresh top-of-stack; on RET the callee’s result replaces the callee’s arguments on the caller’s stack.

Definition 5.4 (Reentrancy semantics; statically disallowed). SAVM execution is strictly non-reentrant. The SAVM enforces this by a static call-graph analysis at compile time: the obligation-pack compiler builds the call graph of all EVAL_OBLIGATION invocations and rejects any pack whose call graph contains a cycle. At runtime, the kernel’s execution layer maintains a single active-SAVM-frame invariant: attempting to invoke a second SAVM instance while one is already executing is a hard trap. No CALL-to-external-contract instruction exists in the ISA (Definition 5.2); the only state I/O is to pre-declared instrument state slots and the read-only context, so the EVM reentrancy pattern [59] is structurally impossible.

Definition 5.5 (SAVM cost model). Each instruction carries a static cost in abstract “gas” units: arithmetic ADD/SUB/NEG/ABS = 1; MUL/DIV = 5 (256-bit intermediate); comparison = 1; control flow = 2 per structural boundary; SLOAD/RLOAD = 10; SSTORE = 20; COMPLIANCE_CHECK = 50 (kernel query); EVAL_OBLIGATION = 10 plus the callee’s cost. The total gas of an execution is bounded statically at compile time (by the call graph depth, the static LOOP bound, and per-instruction cost), and dynamically by a per-call gas limit enforced at runtime.

Theorem 5.6 (SAVM determinism and halting). *Every syntactically valid SAVM program, when executed from a given initial context, either (a) produces a unique output value and a unique post-state, or (b) traps on a well-defined trap reason (overflow, division by zero, stack overflow, out-of-gas, reentrancy violation, compliance-check failure). In either case, execution halts within a number of instructions bounded above by the program’s statically-computed gas budget.*

Proof. Determinism follows from the ISA’s structural properties: (i) FP128 arithmetic is exact in the representable range (no rounding non-determinism); (ii) the stack and memory are manipulated only by statically-validated instructions with deterministic semantics;

(iii) `COMPLIANCE_CHECK` is a pure query against the kernel’s compliance lattice, which is itself deterministic (see [8] for the algebraic semantics). Halting follows from three static bounds: (a) the call graph is acyclic (Definition 5.4), so recursion depth is bounded by the graph’s topological depth D_{graph} ; (b) every `LOOP` carries a static iteration bound enforced at validation; (c) the per-call gas limit is monotone decreasing across instructions, and the program traps when it reaches zero. Combining (a)-(c), the total instruction count is bounded by $D_{\text{graph}} \cdot \max_{\text{ callees}}(\text{static LOOP bound}) \cdot (\text{average instructions per loop body})$, a finite constant. \square

Definition 5.7 (Obligation-pack compilation target). The obligation-pack compiler $\mathcal{K} : \text{ObligationPack} \rightarrow \text{SAVM_Bytecode}$ is a total function from the declarative obligation-pack schema of Section 5.5 to SAVM bytecode. The compilation passes are:

1. **Parsing:** declarative schema \rightarrow abstract syntax tree (AST) of obligation expressions.
2. **Type checking:** AST nodes are annotated with `FP128/Bool/Addr` types; ill-typed packs are rejected.
3. **Compliance-domain resolution:** each compliance reference in the pack is resolved to a kernel domain identifier; unresolved references fail compilation.
4. **Call-graph construction:** sub-obligation references are resolved; cycles cause rejection (Definition 5.4).
5. **Bytecode emission:** AST nodes lower to SAVM instructions in stack order; constants are emitted to the constant pool; `LOOP` iteration bounds are propagated from the schema.
6. **Validation:** emitted bytecode is re-parsed and type-checked as SAVM bytecode; any failure is a compiler bug.

Theorem 5.8 (Compilation correctness). *Let P be a well-typed obligation pack, and $B = \mathcal{K}(P)$ its compiled SAVM bytecode. Then for every admissible input context $(\mathbf{R}, t, \text{caller})$, the output of executing B on the SAVM equals the declarative semantics of P evaluated on the same context. Formally, $\llbracket B \rrbracket_{\text{SAVM}}(\mathbf{R}, t, \text{caller}) = \llbracket P \rrbracket_{\text{decl}}(\mathbf{R}, t, \text{caller})$.*

Proof sketch. The proof is by induction on the structure of the obligation AST. Each compilation rule in Definition 5.7(5) is a standard stack-machine lowering: $e_1 + e_2$ compiles to $\llbracket e_1 \rrbracket; \llbracket e_2 \rrbracket; \text{ADD}$, which on the stack machine has semantics equal to the declarative $+$. The exact correspondence holds for each primitive operation (arithmetic, comparison, compliance check); compound expressions compose by the induction hypothesis. The formal correctness of the compilation pipeline is a refinement-style proof of the kind established for the CompCert compiler [62] in miniature: the source language is the obligation-pack schema (finite set of constructs), and the target language is the SAVM bytecode (finite ISA). A full mechanized proof in Lean or Coq is deferred to future work; the structural lowering rules are total and deterministic, so the informal refinement argument suffices to establish correctness for the specified operations. \square

Open Problem 1 (Mechanized SAVM compilation proof). Formalize the SAVM ISA and the obligation-pack compilation in Lean 4, and mechanize the proof of Theorem 5.8. The inputs are: the declarative obligation-pack semantics (a denotational interpretation in \mathbb{R} -valued functions), the SAVM operational semantics (a small-step relation on $(\mathcal{S}, \mathcal{M}_s)$ pairs), and the compilation relation \mathcal{K} . The output is a Lean proof term of the statement “ $\forall P, \mathbf{R}, t, \text{caller} : \llbracket \mathcal{K}(P) \rrbracket_{\text{SAVM}}(\mathbf{R}, t, \text{caller}) = \llbracket P \rrbracket_{\text{decl}}(\mathbf{R}, t, \text{caller})$.” This is a finite, well-scoped proof task: the ISA is small (Definition 5.2 enumerates all opcodes), and the compilation is a simple structural lowering (Definition 5.7).

The bridge between the kernel and the Moxie exchange is therefore the `CONVEXPOTENTIAL` interface: any instrument whose obligation pack compiles to a realization of this interface can trade on its natural AMM curve.

5.2 The ConvexPotential interface

Definition 5.9 (ConvexPotential). A CONVEXPOTENTIAL is an abstract interface that any instrument class satisfies to define its AMM invariant surface. **Notation.** The name CONVEXPOTENTIAL is retained because the literature and this paper’s citations of it uniformly refer to “convex” invariant surfaces; the underlying mathematical requirement is strict quasi-concavity of φ (not convexity; see Remark 5.10). Every proposition named “ConvexPotential X ” in this paper refers to this abstract interface, and every such proposition’s content is a statement about the quasi-concavity structure of the level sets, not about function-level convexity. The interface exposes two methods:

1. `evaluate(\mathbf{R})` \rightarrow `FixedPoint`: given a reserve vector $\mathbf{R} = (R_{\text{hub}}, R_{\text{brand}})$, returns the potential value $\varphi(\mathbf{R})$. The invariant surface is the level set $\{\mathbf{R} : \varphi(\mathbf{R}) = k\}$ for constant k .
2. `gradient(\mathbf{R})` \rightarrow `(FixedPoint, FixedPoint)`: returns $\nabla\varphi = (\partial\varphi/\partial R_{\text{hub}}, \partial\varphi/\partial R_{\text{brand}})$. The marginal price is $P = (\partial\varphi/\partial R_{\text{hub}})/(\partial\varphi/\partial R_{\text{brand}})$.

The potential φ must satisfy two requirements:

- (C1) **Strictly quasi-concave** on $\mathbb{R}_{>0}^2$: for all $\mathbf{R}_1 \neq \mathbf{R}_2$ on the same level set ($\varphi(\mathbf{R}_1) = \varphi(\mathbf{R}_2) = k$) and $\lambda \in (0, 1)$,

$$\varphi(\lambda\mathbf{R}_1 + (1 - \lambda)\mathbf{R}_2) > k. \quad (8)$$

This ensures that every level set $\{\mathbf{R} : \varphi(\mathbf{R}) = k\}$ is a strictly convex curve in $\mathbb{R}_{>0}^2$, which guarantees a unique price at every reserve state and no arbitrage cycles on the invariant surface.

- (C2) **Monotone**: φ is strictly increasing in each coordinate on $\mathbb{R}_{>0}^2$.

Together, (C1) and (C2) are the operational requirements: unique pricing, convex trading sets $\{\mathbf{R} : \varphi(\mathbf{R}) \geq k\}$, and well-defined marginal rates of substitution along the invariant surface.

Remark 5.10 (Quasi-concavity, not convexity). The AMM literature sometimes states the requirement as “strictly convex potential,” but this is incorrect for the standard constant-product function $\varphi(x, y) = xy$: its Hessian $\begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix}$ is indefinite (eigenvalues ± 1), so xy is neither convex nor concave as a function. What the AMM requires is that the *level sets* be strictly convex curves, equivalently, that φ be strictly quasi-concave on $\mathbb{R}_{>0}^2$. For the constant-product function, the level sets $\{(x, y) : xy = k\}$ are strictly convex hyperbolas, so the operational properties hold despite the function not being convex. We retain the name CONVEXPOTENTIAL for the interface (since the level sets are convex), but the mathematical requirement is (C1)-(C2), not function-level convexity.

Remark 5.11 (Why an open interface, not a parameterized family). One could parameterize a single AMM formula (e.g., Balancer’s weighted geometric mean) and treat each instrument as a parameter choice. The trait design is strictly more expressive: a sukuk potential has a piecewise kink that violates smoothness assumptions required by most parameterized families; a convertible note has reserve-dependent exponents (the tanh-weighted geometric mean) that no static parameterization can represent. The trait design lets each instrument define its own potential from first principles, with only the convexity requirement as the universal constraint. The obligation-pack compiler emits a trait implementation that the kernel’s existing execution layer dispatches at trade time (see Remark 5.1); each evaluation is recorded in a hash-chained receipt that attests every trade was evaluated against the correct invariant surface. At trade time, the dispatcher executes the compiled SAVM bytecode (Section 5.1) against the current pool reserves and the compliance lattice; the trait surface is therefore a thin dispatcher over the underlying ISA.

5.3 Five instrument potentials

We specify the convex potential for five instrument classes. Each potential encodes the economic structure of the instrument directly into the AMM's invariant surface.

Remark 5.12 (Canonical five-potential enumeration and instrument-to-shape mapping). This paper enumerates the five potentials by *instrument class*: equity, sukuk, convertible note, fixed income, option-like. Companion papers in the Moxie research suite [5, 6] sometimes enumerate the same five by *potential shape*: constant-product, constant-sum, weighted-product, concentrated, profit-sharing. These are two views of the same canonical list; the mapping is:

- **Equity** \leftrightarrow constant-product ($\varphi = R_{\text{hub}}R_{\text{brand}}$).
- **Sukuk** \leftrightarrow profit-sharing (constant-product below the kink, profit-sharing-amplified above; cf. Section 5.3 below).
- **Convertible note** \leftrightarrow weighted-product with reserve-dependent exponent (a tanh-interpolated weighted-geometric-mean; cf. Remark 5.15).
- **Fixed income** \leftrightarrow concentrated (constant-product on shifted coordinates with time-decaying virtual reserves).
- **Option-like** \leftrightarrow weighted-product with moneyness-dependent exponent (a second tanh-interpolated weighted-geometric-mean, distinct from the convertible note's by its degree-1 rather than degree-2 homogeneity).

The universality / bisection-convergence result (Proposition 5.21) applies to the instrument-class list of this paper; each of the five names enumerated here satisfies (C1)-(C2) under the stated parameter bounds. Any statement formulated only for the shape list must also discharge the three instrument-class cases (convertible note, fixed income, option-like), whose per-potential parameter conditions (Propositions 5.16, 5.19) carry additional structure beyond the constant-product / constant-sum / concentrated / profit-sharing / weighted-product-80-20 families.

5.3.1 Equity

The standard constant-product potential:

$$\varphi_{\text{equity}}(R_{\text{hub}}, R_{\text{brand}}) = R_{\text{hub}} \cdot R_{\text{brand}}. \quad (9)$$

The invariant surface is the hyperbola $R_{\text{hub}} \cdot R_{\text{brand}} = k$. The marginal price is $P = R_{\text{brand}}/R_{\text{hub}}$. This is the Uniswap v2 curve: symmetric, memoryless, appropriate for instruments with no structural asymmetry between buy and sell pressure. Brand equity tokens (IP-linked tokens representing proportional claims on licensing revenue) trade on this surface.

Gradient: $\nabla \varphi_{\text{equity}} = (R_{\text{brand}}, R_{\text{hub}})$.

5.3.2 Sukuk

The sukuk potential encodes Sharia profit-sharing directly into the invariant surface:

$$\varphi_{\text{sukuk}}(R_{\text{hub}}, R_{\text{brand}}) = \begin{cases} R_{\text{hub}} \cdot R_{\text{brand}} & \text{if } R_{\text{brand}} \leq R_{\text{hub}}, \\ R_{\text{hub}} \cdot (R_{\text{brand}} + \alpha(R_{\text{brand}} - R_{\text{hub}})) & \text{if } R_{\text{brand}} > R_{\text{hub}}, \end{cases} \quad (10)$$

where $\alpha \in (0, 1)$ is the Sharia extraction penalty parameter. The constraint $\alpha < 1$ is necessary for monotonicity at the kink (see Proposition 5.13).

The kink at $R_{\text{brand}} = R_{\text{hub}}$ encodes profit-sharing. Below the kink: the surface is standard constant-product, and the sukuk trades like ordinary equity. Above the kink:

extraction of brand reserves becomes progressively more expensive by the factor $(1 + \alpha)$. This asymmetry reflects the Sharia requirement that profit above a fair-sharing threshold must be distributed, not extracted by one party.

Economic interpretation. When the sukuk is trading at or below par ($R_{\text{brand}} \leq R_{\text{hub}}$), the AMM behaves normally. When the sukuk trades above par ($R_{\text{brand}} > R_{\text{hub}}$), a seller extracting brand-side reserves faces amplified price impact. The excess return is retained in the pool rather than extracted, creating an economic incentive structure consistent with the profit-sharing principle. This is a mechanism design choice, not a Sharia enforcement guarantee; whether the resulting economic behavior satisfies the specific profit-sharing requirements of a given SSB's interpretation remains a matter for SSB determination.

Gradient:

$$\nabla \varphi_{\text{sukuk}} = \begin{cases} (R_{\text{brand}}, R_{\text{hub}}) & \text{if } R_{\text{brand}} \leq R_{\text{hub}}, \\ ((1 + \alpha)R_{\text{brand}} - 2\alpha R_{\text{hub}}, (1 + \alpha)R_{\text{hub}}) & \text{if } R_{\text{brand}} > R_{\text{hub}}. \end{cases} \quad (11)$$

Proposition 5.13 (Sukuk potential: quasi-concavity and kink analysis). φ_{sukuk} satisfies (C1)-(C2) on $\mathbb{R}_{>0}^2$ if and only if $\alpha \in (0, 1)$. For $\alpha \geq 1$, monotonicity (C2) fails at the kink: the partial derivative $\partial\varphi/\partial R_{\text{hub}}$ evaluated from the right at $R_{\text{brand}} = R_{\text{hub}} = R$ equals $R(1 - \alpha) \leq 0$. Specifically, for $\alpha \in (0, 1)$:

1. φ_{sukuk} is continuous (C^0) and nondifferentiable ($\notin C^1$) at the kink $R_{\text{brand}} = R_{\text{hub}}$.
2. In each region separately, φ_{sukuk} is C^∞ and strictly quasi-concave. In region I, $\varphi = R_{\text{hub}}R_{\text{brand}}$ has the standard quasi-concave constant-product form. In region II, the bordered Hessian determinant equals $2(1 + \alpha)^2 R_{\text{hub}}((1 + \alpha)R_{\text{brand}} - \alpha R_{\text{hub}}) > 0$ since $R_{\text{brand}} > R_{\text{hub}} > 0$, confirming strict quasi-concavity.
3. At the kink ($R_{\text{brand}} = R_{\text{hub}} = R$), the gradient is discontinuous: the left gradient is (R, R) and the right gradient is $(R(1 - \alpha), R(1 + \alpha))$. For $\alpha \in (0, 1)$, both components of the right gradient are strictly positive.
4. The level set through the kink is a convex curve (corner, not cusp).

Proof. **Continuity at the kink.** Both pieces evaluate to R^2 when $R_{\text{brand}} = R_{\text{hub}} = R$.

Monotonicity (C2). In region I: $\partial\varphi/\partial R_{\text{hub}} = R_{\text{brand}} > 0$ and $\partial\varphi/\partial R_{\text{brand}} = R_{\text{hub}} > 0$. In region II: $\partial\varphi/\partial R_{\text{brand}} = R_{\text{hub}}(1 + \alpha) > 0$ and $\partial\varphi/\partial R_{\text{hub}} = (1 + \alpha)R_{\text{brand}} - 2\alpha R_{\text{hub}}$. The last expression is positive when $R_{\text{brand}}/R_{\text{hub}} > 2\alpha/(1 + \alpha)$. Since region II requires $R_{\text{brand}} > R_{\text{hub}}$, we need $1 > 2\alpha/(1 + \alpha)$, which simplifies to $\alpha < 1$. At $\alpha = 1$: the right-side partial $\partial\varphi/\partial R_{\text{hub}}|_{R_{\text{brand}}=R_{\text{hub}}} = 0$, violating strict monotonicity. For $\alpha > 1$: the partial is negative for $R_{\text{brand}} \in (R_{\text{hub}}, 2\alpha R_{\text{hub}}/(1 + \alpha))$, a non-empty interval since $2\alpha/(1 + \alpha) > 1$. The constraint $\alpha \in (0, 1)$ is therefore necessary and sufficient for (C2) on the closed region II.

Quasi-concavity (C1). In each region separately, the bordered Hessian confirms strict quasi-concavity (region I is the standard constant-product; region II is verified in item 2 of the proposition statement).

Level-set convexity at the kink. The level set $\{\varphi = k\}$ passing through a kink point $R_{\text{brand}} = R_{\text{hub}} = R$ (with $k = R^2$) is C^∞ in each region. The left-side tangent is orthogonal to the left gradient (R, R) , giving tangent slope -1 . The right-side tangent is orthogonal to the right gradient $(R(1 - \alpha), R(1 + \alpha))$, giving tangent slope $-(1 - \alpha)/(1 + \alpha)$. Since $\alpha \in (0, 1)$, the right-side slope $-(1 - \alpha)/(1 + \alpha) \in (-1, 0)$ is strictly greater (less negative) than the left-side slope -1 . The tangent rotates counterclockwise as we cross from left to right, forming a convex corner. The super-level set $\{\varphi \geq k\}$ is the intersection of the two regional super-level sets (both convex by the bordered Hessian analysis), hence convex, confirming (C1) globally. \square

Remark 5.14 (Gradient at the kink; canonical kink direction). **Canonical kink direction.** The kink in the sukuk potential fires when $R_{\text{brand}} > R_{\text{hub}}$ (region II of the piecewise definition above), the “above-par” regime in which the brand reserves have grown faster than the hub reserves. Profit-sharing engages when the sukuk trades above par: excess returns are retained in the pool, amplifying price impact for a seller extracting brand-side reserves, consistent with the structural reading of an *ijara* profit-sharing arrangement. This is the canonical direction throughout this paper; Proposition 5.13’s quasi-concavity proof and the bordered-Hessian calculation that follows it are valid only for this direction. An implementation that inverts the kink (firing the profit-sharing amplification when $R_{\text{brand}} < R_{\text{hub}}$) describes a different instrument (“profit-extraction-on-depletion” rather than “profit-retention-above-par”) that is not covered by the quasi-concavity analysis here.

Gradient at the kink point. The gradient method at $R_{\text{brand}} = R_{\text{hub}}$ returns the right-hand derivative (region II gradient), which yields the more conservative price (higher effective cost of extraction). The kernel’s trade execution uses a bisection solver on $\varphi(\mathbf{R}) = k$ rather than gradient descent, so the non-differentiable kink does not cause solver failure: bisection requires only continuity and monotonicity, not differentiability. Any gradient-based optimisation (for example, the coupling self-calibration in Section 7) that passes through the kink must use a sub-gradient method or avoid the non-differentiable set entirely; the SGD update in the coupling parameter operates in parameter space, not reserve space, so it is unaffected. Throughout this section, references to a “stepper” or “kernel execution layer” refer to the abstract operational-semantics reduction relation for SAVM bytecode specified in Section 5.1; a general-purpose Turing-complete machine is not required because the SAVM ISA is restricted to the compliance-gated fixed-point operations needed by the obligation-pack schema.

5.3.3 Convertible note

A convertible note has dual nature: debt-like below conversion and equity-like above. Let P_{conv} be the conversion price, $m = R_{\text{brand}} / (R_{\text{hub}} \cdot P_{\text{conv}})$ the moneyness ratio (so $m = 1$ at conversion), and $\beta > 0$ a sharpness parameter controlling the transition width.

The convertible potential uses a weighted geometric mean with a conversion-dependent exponent:

$$\varphi_{\text{conv}}(R_{\text{hub}}, R_{\text{brand}}) = R_{\text{hub}}^{w(m)} \cdot R_{\text{brand}}^{2-w(m)} \quad (12)$$

where the weight function interpolates smoothly between the debt and equity regimes:

$$w(m) = 1 + \frac{1}{2} \tanh(\beta(m - 1)). \quad (13)$$

Deep below conversion ($m \ll 1$, $w \rightarrow 1/2$): $\varphi \approx R_{\text{hub}}^{1/2} \cdot R_{\text{brand}}^{3/2}$. The marginal price $P = w \cdot R_{\text{brand}} / ((2 - w) \cdot R_{\text{hub}}) \approx R_{\text{brand}} / (3R_{\text{hub}})$ is attenuated by a factor of 3 relative to the equity case, reflecting the debt floor; the brand token price moves slowly because the note holder has a fixed claim.

Deep above conversion ($m \gg 1$, $w \rightarrow 3/2$): $\varphi \approx R_{\text{hub}}^{3/2} \cdot R_{\text{brand}}^{1/2}$. The marginal price $P \approx 3R_{\text{brand}} / R_{\text{hub}}$ is amplified relative to the equity case, reflecting full equity exposure with upside participation.

At conversion ($m = 1$, $w = 1$): $\varphi = R_{\text{hub}} \cdot R_{\text{brand}}$, recovering the constant-product surface; the instrument sits at the crossover between debt and equity behaviour.

Remark 5.15 (Why not a hard kink: the piecewise form is discontinuous). A piecewise definition with $R_{\text{hub}} \cdot R_{\text{brand}}^2$ below conversion and $P_{\text{conv}} \cdot R_{\text{hub}} \cdot R_{\text{brand}}$ above is *discontinuous* at $R_{\text{brand}} / R_{\text{hub}} = P_{\text{conv}}$: the first piece evaluates to $P_{\text{conv}}^2 R_{\text{hub}}^3$ and the second to $P_{\text{conv}}^2 R_{\text{hub}}^2$, which differ for $R_{\text{hub}} \neq 1$. The mismatch arises because the two pieces have

different degrees of homogeneity (3 vs. 2) and cannot be stitched continuously along a ray from the origin. A bisection solver run on such a discontinuous φ can find an incorrect root (the solver sees an apparent sign change across the jump and reports a meaningless crossing point), so a piecewise-discontinuous specification silently breaks the clearing engine's invariant. The smooth weighted geometric mean $\varphi_{\text{conv}} = R_{\text{hub}}^{w(m)} R_{\text{brand}}^{2-w(m)}$ with $w(m) = 1 + \frac{1}{2} \tanh(\beta(m-1))$ avoids this pathology, is C^∞ everywhere, is degree-2 homogeneous (since $w + (2-w) = 2$ identically), and preserves the economic intent: muted price response below conversion, full equity exposure above. Whether it satisfies (C1)-(C2) depends on β ; see Proposition 5.16. The smooth weighted-geometric-mean form is the *canonical* specification for the convertible-note potential throughout this paper; any implementation that exposes the piecewise discontinuous form (and is therefore not covered by the Hessian analysis of Proposition 5.16) must be brought into alignment by replacing its evaluator with the smooth form before the convergence proof (Proposition 5.21, conditions (C1)-(C2)) applies.

Gradient: At fixed m (ignoring the weight's dependence on reserves), the gradient is $\nabla \varphi_{\text{conv}} = (w\varphi/R_{\text{hub}}, (2-w)\varphi/R_{\text{brand}})$. The full gradient includes a correction from w 's dependence on $m(R_{\text{hub}}, R_{\text{brand}})$:

$$\nabla \varphi_{\text{conv}} = \varphi_{\text{conv}} \cdot \left[\left(\frac{w}{R_{\text{hub}}}, \frac{2-w}{R_{\text{brand}}} \right) + w'(m) \ln \frac{R_{\text{hub}}}{R_{\text{brand}}} \cdot \nabla m \right] \quad (14)$$

where $w'(m) = \frac{\beta}{2} \text{sech}^2(\beta(m-1))$ and $\nabla m = (-m/R_{\text{hub}}, 1/(R_{\text{hub}}P_{\text{conv}}))$. Since all terms are continuous on $\mathbb{R}_{>0}^2$, the gradient exists everywhere and the kernel stepper does not encounter any non-differentiable point.

Proposition 5.16 (Convertible note: quasi-concavity and Hessian at conversion). *Let $\varphi_{\text{conv}} = R_{\text{hub}}^{w(m)} R_{\text{brand}}^{2-w(m)}$ with $w(m) = 1 + \frac{1}{2} \tanh(\beta(m-1))$ and $m = R_{\text{brand}}/(R_{\text{hub}}P_{\text{conv}})$.*

1. **At conversion** ($m = 1$): $w = 1$, $w' = \beta/2$, and the bordered Hessian of φ_{conv} is:

$$\bar{H}|_{m=1} = 2\varphi^2 \left[1 + \frac{\beta}{2} \ln \frac{R_{\text{hub}}}{R_{\text{brand}}} \cdot \frac{R_{\text{brand}} + R_{\text{hub}}}{R_{\text{hub}}R_{\text{brand}}} \right]. \quad (15)$$

On the level set through $m = 1$, we have $R_{\text{brand}} = P_{\text{conv}}R_{\text{hub}}$, so $\ln(R_{\text{hub}}/R_{\text{brand}}) = -\ln P_{\text{conv}}$. Strict quasi-concavity ($\bar{H} > 0$) at conversion requires:

$$\beta < \frac{2}{|\ln P_{\text{conv}}|} \cdot \frac{P_{\text{conv}}}{1 + P_{\text{conv}}}. \quad (16)$$

For $P_{\text{conv}} = 1$ ($\ln P_{\text{conv}} = 0$), the correction vanishes and $\bar{H} = 2\varphi^2 > 0$ for all β ; the constant-product Hessian dominates. For $P_{\text{conv}} \neq 1$, large β can make $\bar{H} < 0$, destroying quasi-concavity near conversion.

2. **Asymptotic regimes:** For $m \ll 1$ and $m \gg 1$, $w' \rightarrow 0$ exponentially fast, and φ_{conv} reduces to a constant-weight geometric mean $R_{\text{hub}}^{w_\infty} R_{\text{brand}}^{2-w_\infty}$, which is quasi-concave for any $w_\infty \in (0, 2)$. Quasi-concavity can only fail in a neighborhood of $m = 1$ where the w' correction is non-negligible.

3. **Sufficient condition:** The obligation-pack compiler should enforce $\beta \leq \beta_{\text{max}}^{\text{conv}}$ where:

$$\beta_{\text{max}}^{\text{conv}} = \frac{2}{|\ln P_{\text{conv}}|} \cdot \frac{P_{\text{conv}}}{1 + P_{\text{conv}}}. \quad (17)$$

For $P_{\text{conv}} = 2$: $\beta_{\text{max}}^{\text{conv}} \approx 1.92$. For $P_{\text{conv}} = 5$: $\beta_{\text{max}}^{\text{conv}} \approx 1.04$. When $P_{\text{conv}} = 1$, no bound on β is required.

Proof. Write $\varphi = R_h^w R_b^{2-w}$ and $f = \ln \varphi = w \ln R_h + (2-w) \ln R_b$. Since $\varphi > 0$, quasi-concavity of φ is equivalent to quasi-concavity of f (monotone transformation). The gradient of f is:

$$\nabla f = \left(\frac{w}{R_h} + w' \frac{\partial m}{\partial R_h} \ln \frac{R_h}{R_b}, \frac{2-w}{R_b} + w' \frac{\partial m}{\partial R_b} \ln \frac{R_h}{R_b} \right).$$

At $m = 1$: $w = 1$, $w' = \beta/2$, $\partial m / \partial R_h = -m/R_h = -1/R_h$, $\partial m / \partial R_b = 1/(R_h P_{\text{conv}})$. The bordered Hessian of f reduces (after computation) to:

$$\tilde{H}_f = \frac{2}{R_h^2 R_b^2} \left[1 + \frac{\beta}{2} \ln \frac{R_h}{R_b} \cdot \frac{R_b + R_h}{R_h R_b} \right] \cdot R_h^2 R_b^2 = 2 \left[1 + \frac{\beta}{2} \ln \frac{R_h}{R_b} \cdot \frac{R_b + R_h}{R_h R_b} \right].$$

At $m = 1$, the moneyness condition gives $R_b = P_{\text{conv}} R_h$, so $\ln(R_h/R_b) = -\ln P_{\text{conv}}$. For $P_{\text{conv}} > 1$: $\ln(R_h/R_b) < 0$, and $(R_b + R_h)/(R_h R_b) > 0$, so the correction is negative. Setting $\tilde{H}_f > 0$ yields the stated bound. The asymptotic claims follow from $w'(m) = (\beta/2) \text{sech}^2(\beta(m-1)) \rightarrow 0$ exponentially as $|m-1| \rightarrow \infty$. \square

5.3.4 Fixed income

Fixed-income instruments have a defining feature: they converge to par at maturity. The AMM must reflect this. The potential uses time-dependent virtual reserves that decay to zero:

$$\varphi_{\text{fixed}}(R_{\text{hub}}, R_{\text{brand}}, t) = (R_{\text{hub}} + V_{\text{hub}}(t)) \cdot (R_{\text{brand}} + V_{\text{brand}}(t)) \quad (18)$$

where the virtual reserves decay linearly:

$$V_{\text{hub}}(t) = V_{\text{hub}}^0 \cdot \frac{T-t}{T}, \quad V_{\text{brand}}(t) = V_{\text{brand}}^0 \cdot \frac{T-t}{T}, \quad (19)$$

T is the maturity date, and V^0 are the initial virtual reserves.

At issuance ($t = 0$): the virtual reserves are large, the effective reserves are $R + V^0$, and the invariant surface is wide; the instrument trades in a broad liquidity range around par (analogous to a wide tick range in Uniswap V3 [21]) because the duration risk is substantial.

At maturity ($t = T$): the virtual reserves are zero, $V(T) = 0$, and the surface collapses to the real reserves $\varphi = R_{\text{hub}} \cdot R_{\text{brand}}$. The price at maturity is $R_{\text{brand}}/R_{\text{hub}}$, which equals par only if the real reserves are at the par ratio. The virtual reserve mechanism does not force deterministic convergence to par; it narrows the liquidity band, increasing price impact per trade as maturity approaches, which incentivises arbitrageurs to push the price toward par. This is an economic incentive (bounded-loss arbitrage opportunity), not a mechanical guarantee.

Gradient: $\nabla \varphi_{\text{fixed}} = (R_{\text{brand}} + V_{\text{brand}}(t), R_{\text{hub}} + V_{\text{hub}}(t))$.

This is the constant-product gradient evaluated at $(R + V(t))$, time-parameterized. The kernel stepper reads block time to compute $V(t)$ at each trade.

Quasi-concavity in reserves. At each fixed t , φ_{fixed} is the standard constant-product function in shifted coordinates $(R_{\text{hub}} + V_{\text{hub}}(t), R_{\text{brand}} + V_{\text{brand}}(t))$ with both shifts non-negative. It is therefore strictly quasi-concave in $(R_{\text{hub}}, R_{\text{brand}})$ at every $t \in [0, T]$, satisfying (C1)-(C2).

Near-maturity liquidity risk. The second time derivative $\partial^2 \varphi / \partial t^2 = 2V_{\text{hub}}^0 V_{\text{brand}}^0 / T^2 > 0$: the potential is convex in time, meaning effective liquidity shrinks at an accelerating rate. In the final δT before maturity, the effective reserves are $R + V^0 \delta$, and the price impact of a trade of size Δ scales as $\Delta / (R + V^0 \delta)^2$. The obligation-pack compiler must enforce a minimum virtual reserve floor $V_{\text{min}} > 0$ or halt AMM trading in the final window and switch to a settlement mechanism, to prevent unbounded price impact near maturity.

Remark 5.17 (Closed-form swap inverse; bisection not required for fixed-income). Because $\varphi_{\text{fixed}}(R_h, R_b, t) = (R_h + V_h(t))(R_b + V_b(t))$ is the standard constant-product potential in shifted coordinates, the swap equation $\varphi(R_h - \Delta_{\text{out}}, R_b + \Delta_{\text{in}}, t) = \varphi(R_h, R_b, t)$ admits a closed-form inverse at fixed t :

$$\Delta_{\text{out}}(t) = (R_h + V_h(t)) - \frac{(R_h + V_h(t))(R_b + V_b(t))}{R_b + V_b(t) + \Delta_{\text{in}}}.$$

The bisection solver is not required in this instrument class; the exchange's dispatch layer should evaluate the closed-form inverse directly, avoiding the potential issue of bisection-convergence on a path that crosses a block boundary (where $t_{\text{pre}} \neq t_{\text{post}}$ and the solver's target $\varphi_0 = \varphi(\mathbf{R}, t_{\text{pre}})$ would shift mid-solve). The closed-form is evaluated at block-start time $t = t_{\text{pre}}$; any intra-block time-decay adjustment is deterministically applied at the block boundary via a separate updater. Quasi-concavity at fixed t (Section 5.3) is sufficient for the closed-form inverse to be well-defined and unique; the universal-bisection argument of Proposition 5.21 is not invoked for this specific instrument class.

Remark 5.18 (Revenue-linked notes as time-parameterized potentials). Revenue-linked notes (Definition 3.3) are a special case of the fixed-income potential where the virtual reserve decay rate is modulated by observed revenue: $V(t) = V^0 \cdot (c - \sum_{s < t} \text{Coupon}_s) / c$. The surface collapses when cumulative coupons reach the cap c . The kernel stepper reads the kernel's fiscal module to obtain the cumulative coupon, updating the virtual reserves accordingly.

5.3.5 Option-like

Option-like instruments (contingent licensing options, Definition 3.6) require a potential that shifts shape with moneyness. We use a weighted geometric mean with moneyness-dependent weights:

$$\varphi_{\text{option}}(R_{\text{hub}}, R_{\text{brand}}) = R_{\text{hub}}^{w_1(m)} \cdot R_{\text{brand}}^{w_2(m)}, \quad w_1 + w_2 = 1, \quad (20)$$

where $m = R_{\text{brand}} / (R_{\text{hub}} \cdot K)$ is the moneyness (the reserve-implied price $R_{\text{brand}} / R_{\text{hub}}$ relative to strike K), and the weights are:

$$w_1(m) = \frac{1}{2} + \frac{1}{2} \tanh(\beta(m - 1)), \quad w_2(m) = 1 - w_1(m), \quad (21)$$

with $\beta > 0$ controlling the sharpness of the weight transition.

At-the-money ($m \approx 1$): $w_1 \approx w_2 \approx 1/2$. The surface is nearly symmetric (high gamma), reflecting the option's maximum sensitivity to price changes near the strike.

In-the-money ($m \gg 1$): $w_1 \rightarrow 1, w_2 \rightarrow 0$. The surface skews toward the hub (in-the-money) side, and the instrument behaves like a leveraged equity position.

Out-of-the-money ($m \ll 1$): $w_1 \rightarrow 0, w_2 \rightarrow 1$. The surface skews toward the brand side, and the instrument has low delta: small price sensitivity, consistent with an OTM option.

Gradient: Since the weights w_1, w_2 depend on reserves through $m = R_{\text{brand}} / (R_{\text{hub}}K)$, the gradient is *not* the constant-weight formula. The full gradient is:

$$\nabla \varphi_{\text{option}} = \varphi_{\text{option}} \cdot \left[\left(\frac{w_1}{R_{\text{hub}}}, \frac{w_2}{R_{\text{brand}}} \right) + w_1'(m) \ln \frac{R_{\text{hub}}}{R_{\text{brand}}} \cdot \nabla m \right] \quad (22)$$

where $w_1'(m) = \frac{\beta}{2} \text{sech}^2(\beta(m - 1))$ and $\nabla m = (-m/R_{\text{hub}}, 1/(R_{\text{hub}}K))$. The first term is the standard weighted-geometric-mean gradient; the second is the correction from the weights' dependence on reserves. At the asymptotic limits ($m \rightarrow 0$ or $m \rightarrow \infty$), $w_1' \rightarrow 0$

and the correction vanishes, recovering the constant-weight formula. Near the strike ($m \approx 1$), the correction is maximal and non-negligible; the kernel stepper must evaluate the full expression.

Proposition 5.19 (Quasi-concavity of the option potential). *Let $\varphi_{\text{option}} = R_{\text{hub}}^{w_1(m)} R_{\text{brand}}^{w_2(m)}$ with $w_1 + w_2 = 1$ and $w_1(m) = \frac{1}{2} + \frac{1}{2} \tanh(\beta(m-1))$.*

1. *For constant weights, $R_{\text{hub}}^{w_1} R_{\text{brand}}^{w_2}$ with $w_1 + w_2 = 1$ is homogeneous of degree 1, log-linear, and quasi-concave on $\mathbb{R}_{>0}^2$.*
2. *When the weights depend on reserves through $m = R_{\text{brand}} / (R_{\text{hub}} K)$, the degree of homogeneity remains 1 (since $w_1 + w_2 = 1$ identically), but quasi-concavity is not guaranteed a priori. The key difference from the convertible note (Proposition 5.16, degree 2) is that the baseline bordered Hessian of a degree-1 homogeneous function is smaller in magnitude, so the w' correction dominates at lower β .*
3. **Analytical bound.** *At $m = 1$ (at-the-money, $R_{\text{brand}} = KR_{\text{hub}}$), $w_1 = w_2 = 1/2$, and the bordered Hessian of $f = \ln \varphi$ evaluates to:*

$$\bar{H}_f|_{m=1} = \frac{1}{2R_h^2 K^2 R_b^2} \left[1 + \beta \ln \frac{1}{K} \cdot \frac{K+1}{K} \right]. \quad (23)$$

Strict quasi-concavity ($\bar{H}_f > 0$) at the strike requires:

$$\beta < \frac{K}{(K+1) |\ln K|} \quad \text{for } K \neq 1. \quad (24)$$

For $K = 1$: the correction vanishes and quasi-concavity holds for all β . For $K = 2$: $\beta_{\text{max}} \approx 0.96$. For $K = 5$: $\beta_{\text{max}} \approx 0.52$. For $K = 0.5$: $\beta_{\text{max}} \approx 0.96$.

The bound is symmetric in $K \leftrightarrow 1/K$ (by symmetry of $|\ln K|$) and is tighter than the convertible note bound because the degree-1 baseline curvature is weaker. The obligation-pack compiler must enforce $\beta \leq \beta_{\text{max}}^{\text{opt}}(K)$ as a deployment parameter.

Proof. The argument follows the same structure as Proposition 5.16. Write $f = \ln \varphi = w_1 \ln R_h + w_2 \ln R_b$ where $w_1 + w_2 = 1$. At $m = 1$: $w_1 = w_2 = 1/2$, $w'_1 = \beta/2$, $\nabla m = (-m/R_h, 1/(R_h K)) = (-1/R_h, 1/(R_h K))$. The gradient of f at $m = 1$ is:

$$\nabla f = \left(\frac{1}{2R_h} - \frac{\beta}{2} \cdot \frac{\ln K}{R_h}, \frac{1}{2R_b} + \frac{\beta}{2} \cdot \frac{(-\ln K)}{R_h K} \right).$$

The bordered Hessian computation proceeds as in Proposition 5.16, substituting $w_1 + w_2 = 1$ instead of 2, and $R_b = KR_h$ on the $m = 1$ level set. The leading term in \bar{H}_f (from the constant-weight part) scales as $1/(R_h^2 R_b^2) \cdot (w_1 w_2) = 1/(4R_h^2 R_b^2)$, which is half the magnitude of the convertible note's leading term (where $w(2-w) = 1$ at $m = 1$). Setting $\bar{H}_f > 0$ gives the stated bound on β . \square

Remark 5.20 (Canonical option-like weight form). The tanh-based weight $w_1(m) = \frac{1}{2} + \frac{1}{2} \tanh(\beta(m-1))$ with parameter $\beta \leq \beta_{\text{max}}^{\text{opt}}(K)$ is the canonical specification for the option-like potential throughout this paper; Proposition 5.19's quasi-concavity proof and the $\beta_{\text{max}}^{\text{opt}}$ bound are stated for exactly this form. Linear-clamped weight approximations ($w_1 = \frac{1}{2} + \frac{1}{2}m$ clamped to $[\frac{1}{2}, 2]$) produce a C^0 -only (not C^∞) weight function and are not covered by Proposition 5.19; implementations using such approximations must either prove their own quasi-concavity result or be brought into alignment with the tanh form before the universality claim (Proposition 5.21) applies to them. The β parameter must be exposed in the option-like potential's deployment parameters and enforced against $\beta_{\text{max}}^{\text{opt}}(K)$ at obligation-pack compile time.

Proposition 5.21 (Quasi-concave potential universality and bisection convergence). *Let $\varphi : \mathbb{R}_{>0}^2 \rightarrow \mathbb{R}$ be a function satisfying the interface requirements (C1)-(C2) of Definition 5.9 (strict quasi-concavity and monotonicity). Then φ realizes the CONVEXPOTENTIAL interface, and the bisection solver converges to a unique root for any swap computation $\varphi(R_{\text{hub}} - \Delta_{\text{out}}, R_{\text{brand}} + \Delta_{\text{in}}) = k$. The five potentials enumerated in Section 5.3 (equity, sukuk, convertible note, fixed income, option-like) each satisfy (C1)-(C2) conditional on the per-potential parameter bounds stated in Propositions 5.13 ($\alpha \in (0, 1)$), 5.16 ($\beta \leq \beta_{\text{max}}^{\text{conv}}(P_{\text{conv}})$), and 5.19 ($\beta \leq \beta_{\text{max}}^{\text{opt}}(K)$). These parameter bounds must be enforced at compile time by the obligation-pack compiler (Definition 5.7); a pack whose α , β or strike violates the applicable bound is rejected before SAVM bytecode is emitted. Additional instrument classes admit by realising the interface with a new φ satisfying (C1)-(C2) and any additional parameter bounds required to certify (C1)-(C2) for that φ ; no modification to the settlement logic, coupling function, or clearinghouse is required beyond the compile-time parameter check.*

Proof. The CONVEXPOTENTIAL interface requires: (i) φ evaluation, (ii) a gradient or sub-gradient, and (iii) the invariant $\varphi(R') = \varphi(R)$ for swap computation. Any φ satisfying (C1) (strict quasi-concavity on $\mathbb{R}_{>0}^2$) has connected convex level sets, so by the monotonicity condition (C2) the residual $h(\Delta_{\text{out}}) = \varphi(R_{\text{hub}} - \Delta_{\text{out}}, R_{\text{brand}} + \Delta_{\text{in}}) - k$ is strictly decreasing and continuous, and by the intermediate value theorem the bisection solver finds a unique root. (C2) also ensures the solver’s search interval is bounded ($\Delta_{\text{out}} \in [0, R_{\text{hub}})$). Each of the five potentials was verified to satisfy (C1)-(C2) in the propositions above, conditional on the stated parameter bounds. The universality claim is therefore conditional on the obligation-pack compiler enforcing those bounds at Definition 5.7’s type-checking pass. Extensibility follows because the clearing engine depends on φ only through the abstract interface, and the parameter check is a per-pack operation that does not propagate to the clearinghouse. \square

Remark 5.22 (Non-smooth and non-differentiable potentials; curvature-based margin formula). The universality claim requires only (C1)-(C2), not smoothness. The sukuk potential is C^0 and outside C^1 (gradient discontinuity at the kink); the kernel’s bisection solver handles this without difficulty. Any downstream computation that assumes differentiability (*curvature-based margin* at the Hot tier, the bordered-Hessian determinant of $\ln \varphi$, gradient-based parameter optimisation) must handle non-smooth potentials via sub-gradients or regularisation.

Curvature-based margin, precise formula. We use “curvature-based margin” to mean the constrained-optimization margin driven by the bordered Hessian determinant of $f := \ln \varphi$ evaluated on the level set; it is *not* the Hessian of φ itself, which is indefinite for $\varphi(x, y) = xy$ and carries no margin-input interpretation (Remark 5.10). Concretely, the clearinghouse computes margin as

$$\text{Margin}_{\text{curv}}(\mathbf{R}) := \mu_0 + \mu_1 \cdot |\det \bar{H}_f(\mathbf{R})|^{1/2}, \quad (25)$$

where \bar{H}_f is the 3×3 bordered Hessian of $f = \ln \varphi$ with the gradient of f on the border, $\mu_0 > 0$ is a baseline margin floor, and $\mu_1 > 0$ is a curvature-sensitivity coefficient. The determinant $\det \bar{H}_f$ is the standard quasi-concavity indicator: $\det \bar{H}_f > 0$ is the bordered-Hessian condition for strict quasi-concavity of f (and hence of φ), and $|\det \bar{H}_f|^{1/2}$ has units of reciprocal-curvature on the level set (higher value = flatter level set = thinner liquidity = larger margin). Equation (25) is what “curvature-based margin” denotes throughout this paper and the associated clearinghouse implementation.

Non-smooth case. For non-smooth potentials such as the sukuk kink, the bordered Hessian is evaluated on each smooth piece separately and the clearinghouse takes the

most conservative (highest margin) result:

$$\text{Margin}_{\text{curv}}^{\text{ns}}(\mathbf{R}) := \max_{k \in \{I, II\}} \text{Margin}_{\text{curv}}(\mathbf{R}; \tilde{H}_f^{(k)}), \quad (26)$$

where $\tilde{H}_f^{(k)}$ is the bordered Hessian evaluated on piece k of the domain (I, II for the sukuk kink's two regions). At the sukuk kink, this means computing margin from both region-I and region-II bordered Hessians and taking the maximum. The gradient method in the interface specification should be documented as returning a sub-gradient at non-differentiable points, with the convention that the right-hand derivative is returned at kinks.

5.4 Temperature tier capability bindings for securities

Section 4 distinguished the operational *temperature tier* (Definition 4.7) from jurisdictional *regulatory graduation* (Definition 4.8). Here we make the venue-internal tier-to-capability binding precise: each temperature tier imposes constraints on which CONVEXPOTENTIAL implementations are permitted, and the tensor-to-tier map is a meet-preserving morphism from the Applicable-fragment compliance structure to the temperature lattice.

Definition 5.23 (Temperature tiers for securities).

- **Cold:** pre-graduation. At least one applicable coordinate for the instrument fails the venue's Warm threshold, or the bridge envelope is incomplete. Cold is compatible with issuance and restricted transfer paths, but it does not admit open secondary clearing. A Cold-tier instrument may still be a security, derivative, insurance-linked contract, or other regulated instrument under the relevant jurisdictional framework.
- **Warm:** partial-compliance gated tier. Every domain in $D(I)$ meets the instrument-specific admission threshold, but the full 23-domain tensor still contains at least one applicable coordinate that remains below Hot. Instruments include revenue-linked notes, contingent licensing options, and sukuk. Potential: instrument-specific (sukuk, convertible, fixed-income, option-like). Access is per instrument and per participant.
- **Hot:** full graduation tier. Every one of the 23 coordinates is either compliant at the Hot threshold or not applicable. Instruments such as index products and fully qualified note programs may clear on the open secondary venue. Potential: any CONVEXPOTENTIAL; the clearinghouse imposes additional margin requirements computed from the potential's curvature (Hessian determinant where it exists; for non-smooth potentials such as the sukuk kink, the margin uses the bordered Hessian evaluated on each smooth piece and takes the conservative bound; see Remark 5.22).

Definition 5.24 (Temperature-compliance meet-semilattice morphism). Let \mathcal{L}_{23} denote the Mass compliance tensor over 23 domains, in canonical order (1-indexed): Aml, Kyc, Sanctions, Tax, Securities, Corporate, Custody, DataPrivacy, Licensing, Banking, Payments, Clearing, Settlement, DigitalAssets, Employment, Immigration, Ip, ConsumerProtection, Arbitration, Trade, Insurance, AntiBribery, Sharia. The companion compliance-tensor chapter [8] gives the full per-domain grade-set specification together with the Applicable-fragment meet structure and the F144 dichotomy for the full mixed-axis tensor; Definition 5.25 specifies the Sharia grade chain and its SH-01..SH-05 component constraint decomposition. Sharia occupies the 23rd slot. Twelve of the 23 domains are activated by one or more of the six instrument classes in Section 3; the remaining eleven appear through entity-level operations or extended product structures such as insurance and anti-bribery screening. Let $\mathcal{T} = \{\text{Cold}, \text{Warm}, \text{Hot}\}$ be the temperature lattice with ordering $\text{Cold} \leq \text{Warm} \leq \text{Hot}$. The temperature classification function

$$\Phi : \mathcal{L}_{23} \rightarrow \mathcal{T} \quad (27)$$

is a *meet-semilattice morphism on the Applicable fragment*: $\Phi(a \wedge b) = \Phi(a) \wedge \Phi(b)$ and $\Phi(\top) = \top_{\mathcal{T}}$ for tensors compared through applicable coordinates only. It is *not* a full lattice homomorphism, and it is not a claim about the full mixed-axis tensor with explicit applicability state. Since multi-harbor composition uses meet, this suffices for the operational semantics. Composing compliance constraints across jurisdictions therefore commutes with temperature classification on the Applicable fragment. An entity harbored in jurisdictions ϕ_1 and ϕ_2 trades at temperature $\Phi(\mathcal{C}(\phi_1) \wedge \mathcal{C}(\phi_2)) = \Phi(\mathcal{C}(\phi_1)) \wedge \Phi(\mathcal{C}(\phi_2))$, the minimum temperature permitted by either jurisdiction.

Definition 5.25 (Per-domain grade sets, Sharia grade chain, and SH-01..SH-05 component constraints). Each compliance domain $d \in \{1, \dots, 23\}$ carries its own totally-ordered finite grade set \mathcal{G}_d with bottom element \perp_d and top element \top_d . Meet and join on \mathcal{L}_{23} are defined pointwise using the per-domain ordering: $(a \wedge b)_d := \min_{\leq_d}(a_d, b_d)$, $(a \vee b)_d := \max_{\leq_d}(a_d, b_d)$. For the specific case of domain $d = 23$ (Sharia), we fix the grade chain explicitly:

$$\mathcal{G}_{23} := \{\text{NotRecognized} < \text{NoSSB} < \text{SSBPending} < \text{PartiallyCertified} < \text{FullyCertified}\}, \quad (28)$$

with $\perp_{23} = \text{NotRecognized}$ (assigned to any jurisdiction that does not recognize Sharia as a regulatory concept, e.g., Delaware) and $\top_{23} = \text{FullyCertified}$ (SSB fatwa current and valid under jurisdictional recognition). The intermediate grades ordered strictly: NoSSB (Sharia-aware jurisdiction but no SSB engaged), SSBPending (SSB engaged, certification in progress), PartiallyCertified (structural checks pass but SSB has conditional reservations). The ordering is total on \mathcal{G}_{23} by construction; any two grades are comparable.

The Sharia domain further decomposes into five component constraints, each with its own grade projection onto \mathcal{G}_{23} :

- **SH-01 (No Riba)**: no interest-bearing components in any instrument held, issued, or transacted by the entity; profit-rate margins (murabaha markup, ijara rental rate) are permitted because they represent compensation for asset use or trade margin, not time-value-of-money on a loan.
- **SH-02 (No Gharar)**: no excessive uncertainty in contractual terms; subject matter, price, delivery schedule, and counterparty obligations must all be specified. Minor gharar (*gharar yasir*) in standardized exchange-traded instruments is permissible per AAOIFI FAS 1.
- **SH-03 (No Maysir)**: no pure speculation or gambling; event-contingent contracts must demonstrate hedging purpose against held or committed underlying exposure.
- **SH-04 (Asset Backing)**: all debt-like instruments must reference an identified, existing, halal asset owned or contractually committed by the originator at issuance.
- **SH-05 (SSB Certification)**: the entity must hold a current SSB attestation covering the entity's instrument classes; absence of an SSB attestation yields Pending on this component constraint (mapped to SSBPending on \mathcal{G}_{23}), not NotRecognized.

The composite Sharia grade on \mathcal{G}_{23} is the pointwise meet of the five component constraint projections: an entity's composed Sharia grade is FullyCertified iff all five of SH-01..SH-05 pass at their respective tops, and degrades gracefully (through PartiallyCertified, SSBPending, NoSSB) as individual component constraints fail or remain pending.

Remark 5.26 (SH-01..SH-05 are treated as orthogonal as a modelling approximation). The pointwise meet over five grade chains treats SH-01..SH-05 as orthogonal. Classical *fiqh al-muamalat* does not. *Gharar yasir* (minor gharar) permissible under AAOIFI FAS 1 is conditioned on SH-04: a minor gharar relating to an identified existing asset may be acceptable, while the same magnitude of uncertainty in a purely contractual claim is not. Similarly, *tawarruq* and *murabaha* have interactions between SH-01 (riba) and SH-04 (asset-backing) that a pointwise meet does not capture. We state this as a known structural

approximation, not an exactness claim. The SSB certification component constraint (SH-05) is the place where substantive cross-constraint reasoning is performed; the pointwise-meet composition is a screening filter, not a substitute for that reasoning.

Lemma 5.27 (Pointwise meet well-definedness across heterogeneous grade sets). *For any $a, b \in \mathcal{L}_{23} := \prod_{d=1}^{23} \mathcal{G}_d$, the pointwise meet $a \wedge b \in \mathcal{L}_{23}$ is a total, well-defined operation; it is idempotent, commutative, associative, and absorbs the join. In particular, for domain 23 the meet $\text{FullyCertified} \wedge \text{NotRecognized} = \text{NotRecognized}$, so composing a Sharia-aware harbor (e.g., ADGM) with a Sharia-unaware harbor (e.g., Delaware) yields the most-restrictive (bottom) Sharia grade for the composed entity.*

Proof. Each \mathcal{G}_d is totally ordered by hypothesis; hence \min_{\leq_d} is a total function $\mathcal{G}_d \times \mathcal{G}_d \rightarrow \mathcal{G}_d$. The pointwise meet on the product lattice $\prod_d \mathcal{G}_d$ is therefore total. Lattice axioms (idempotency $a \wedge a = a$; commutativity $a \wedge b = b \wedge a$; associativity $(a \wedge b) \wedge c = a \wedge (b \wedge c)$; absorption $a \wedge (a \vee b) = a$) transport from each factor to the product by coordinate-wise argument. The Sharia-specific calculation follows because $\text{NotRecognized} <_{23} \text{FullyCertified}$ in \mathcal{G}_{23} , hence $\min_{\leq_{23}} = \text{NotRecognized}$. \square

Remark 5.28 (Semantic meaning of Sharia meet across jurisdictions). Lemma 5.27 answers the cross-jurisdiction question directly: an entity harbored in both ADGM (Sharia-recognized, $c_{23}(\text{ADGM}) = \text{FullyCertified}$) and Delaware ($c_{23}(\text{Delaware}) = \text{NotRecognized}$) has composed Sharia grade NotRecognized . This is the operationally correct answer: the entity cannot rely on its Sharia certification for any trade whose settlement surface includes the Delaware harbor, because Delaware has no statutory hook for Sharia grades. The corridor mechanism (Definition 4.2, \mathcal{K}) is what carries Sharia grades between Sharia-aware jurisdictions (e.g., ADGM \leftrightarrow Saudi); the absence of a corridor carrying $d = 23$ between ADGM and Delaware is precisely the expression of Delaware’s non-recognition.

Proof of meet-semilattice morphism structure. The temperature classification Φ is defined by threshold predicates on the compliance state vector. Define the qualifying predicate for each tier: $Q_{\text{Hot}}(c) \iff c_d \geq \tau_d^{\text{Hot}}$ for all $d \in D_{\text{Hot}}$; $Q_{\text{Warm}}(c) \iff c_d \geq \tau_d^{\text{Warm}}$ for all $d \in D_{\text{Warm}}$. Then $\Phi(c) = \text{Hot}$ if $Q_{\text{Hot}}(c)$; $\Phi(c) = \text{Warm}$ if $Q_{\text{Warm}}(c) \wedge \neg Q_{\text{Hot}}(c)$; $\Phi(c) = \text{Cold}$ otherwise.

Structural requirement. The morphism requires that the tier definitions are *nested*: $D_{\text{Warm}} \subseteq D_{\text{Hot}}$ and $\tau_d^{\text{Warm}} \leq \tau_d^{\text{Hot}}$ for all $d \in D_{\text{Warm}}$. This ensures $Q_{\text{Hot}}(c) \implies Q_{\text{Warm}}(c)$, so Φ is monotone: $c \leq c'$ componentwise implies $\Phi(c) \leq \Phi(c')$.

Proof for each tier level. We must show $\Phi(a \wedge b) = \Phi(a) \wedge \Phi(b)$ for all $a, b \in \mathcal{L}_{23}$. Let $a \wedge b = (\min(a_1, b_1), \dots, \min(a_{23}, b_{23}))$.

Case Hot. $Q_{\text{Hot}}(a \wedge b) \iff \min(a_d, b_d) \geq \tau_d^{\text{Hot}}$ for all $d \in D_{\text{Hot}} \iff Q_{\text{Hot}}(a) \wedge Q_{\text{Hot}}(b)$. Hence $\Phi(a \wedge b) = \text{Hot} \iff \Phi(a) = \text{Hot} \text{ and } \Phi(b) = \text{Hot} \iff \Phi(a) \wedge \Phi(b) = \text{Hot}$.

Case Warm. By the same argument, $Q_{\text{Warm}}(a \wedge b) \iff Q_{\text{Warm}}(a) \wedge Q_{\text{Warm}}(b)$. So $\Phi(a \wedge b) \geq \text{Warm} \iff \Phi(a) \geq \text{Warm} \text{ and } \Phi(b) \geq \text{Warm} \iff \Phi(a) \wedge \Phi(b) \geq \text{Warm}$. Combined with the Hot case: $\Phi(a \wedge b) = \text{Warm} \iff \Phi(a) \wedge \Phi(b) = \text{Warm}$.

Case Cold. If $\Phi(a \wedge b) = \text{Cold}$, then $\neg Q_{\text{Warm}}(a \wedge b)$, so $\neg(Q_{\text{Warm}}(a) \wedge Q_{\text{Warm}}(b))$, meaning at least one of $\Phi(a), \Phi(b)$ is Cold, hence $\Phi(a) \wedge \Phi(b) = \text{Cold}$. Conversely, if $\Phi(a) \wedge \Phi(b) = \text{Cold}$, then (WLOG) $\Phi(a) = \text{Cold}$, so $\neg Q_{\text{Warm}}(a)$, hence $\neg Q_{\text{Warm}}(a \wedge b)$ (since $a \wedge b \leq a$), so $\Phi(a \wedge b) = \text{Cold}$.

All three cases confirm $\Phi(a \wedge b) = \Phi(a) \wedge \Phi(b)$.

Caveat. Meet-preservation (\wedge) holds; join-preservation (\vee) fails in general: $\Phi(a \vee b)$ can exceed $\Phi(a) \vee \Phi(b)$ when two compliance states that individually fail the Hot threshold combine (via componentwise max) to satisfy it. This is the explicit content of “ Φ is a morphism in **MSLat** and fails to be a morphism in **DistLat**”. Since multi-harbor

composition uses meet (the most restrictive constraint), this asymmetry does not affect the system's operational semantics. \square

Meet-preservation makes the operational tier classification compositional: adding a new harbour can only maintain or lower the entity's temperature tier, because the meet selects the most restrictive constraint. This is why regulatory graduation (Definition 4.8) must be secured in every harbour a tier crosses: a single sub-threshold harbour pulls the composed tier down. An entity at Warm in ADGM and Warm in Seychelles is Warm in the composed surface. An entity at Hot in ADGM but Cold in Seychelles (because Seychelles does not recognize the relevant securities registration) trades at Cold until the Seychelles harbor is upgraded or removed.

5.5 Obligation pack verification

Before a brand token entity can trade on a compliance-aware L1, the Mass kernel verifies a set of recurring obligations defined in the entity's obligation pack. The Moxie brand token obligation pack contains six rules, each enforced by the kernel's obligation runtime.

1. **Quarterly token holder report** (`moxie.quarterly-token-holder-report`). Quarterly disclosure covering revenue distribution, P&L, treasury balances, and token economics. Required for all active brand token entities regardless of temperature tier. Blocking condition: the token holder disclosure actuator must confirm publication. Evidence: on-chain disclosure receipt and treasury settlement confirmation for the declared distribution.
2. **Annual temperature tier re-evaluation** (`moxie.annual-temperature-tier-evaluation`). Annual re-evaluation of the entity's Cold/Warm/Hot classification. Any attestation, corridor, or obligation-pack update may trigger an earlier recomputation for the affected instrument. Evidence: signed tier evaluation report documenting the compliance posture across all 23 domains.
3. **Annual event resolution audit** (`moxie.event-resolution-audit`). Annual audit verifying that every oracle-resolved event, across all five oracle tiers (Tier 0 Programmatic through Tier 4 Schelling Point), was resolved correctly, escalation paths were followed, and disputed outcomes were remediated. Evidence: complete audit report with per-event oracle tier, resolution timestamp, outcome, and dispute history.
4. **Quarterly insurance fund contribution** (`moxie.insurance-fund-contribution`). Quarterly assessment and settlement of the entity's insurance fund share. The contribution is computed from market share of open interest, correlation exposure ρ_{ij} with other brand tokens, and tail risk profile (Pareto tail index ζ). The insurance fund backstops clearinghouse solvency; the circuit breaker triggers if the fund reaches zero. Evidence: treasury settlement confirmation and signed adequacy attestation.
5. **Annual Sharia Supervisory Board renewal** (`moxie.sharia-board-renewal`). Applies only to brand token entities carrying the `islamic-compliant` designation. The SSB reviews the token's economic structure, revenue distribution mechanics, underlying asset composition, and market operations. If the SSB does not renew certification, the `islamic-compliant` label is removed and the `sukuk` potential (Section 5.3) is replaced with the standard equity potential; the entity continues trading without the Sharia-encoded extraction penalty. Evidence: renewed SSB fatwa, audit report, and published token holder notice.
6. **Material change notice** (`moxie.material-change-notice`). Event-driven obligation triggered by material changes to the entity's structure, economics, or operations (brand licensing amendments, revenue share modifications, collateral parameter changes, tier-affecting compliance events, brand operator changes). Notice must

be published within 5 business days. Evidence: published governance notice documenting the change and its impact.

Lemma 5.29 (Obligation pack completeness for trading eligibility). *A brand token entity is eligible to trade on the venue if and only if all non-event-driven obligations (rules 1-5) have satisfying evidence for the current period, and all triggered event-driven obligations (rule 6) have been cured. The kernel’s obligation evaluator computes this predicate before every trading session: a brand token with an unsatisfied obligation is suspended from trading until the obligation is cured.*

Proof. Each obligation o_k has a Boolean satisfaction predicate $\text{sat}(o_k, t)$ that checks evidence validity at time t . Non-event-driven obligations (rules 1-5) have periodic deadlines; $\text{sat}(o_k, t) = 1$ iff valid evidence exists with timestamp in the current period. Event-driven obligations (rule 6) have $\text{sat}(o_k, t) = 1$ iff either no trigger has fired or the cure evidence has been submitted. Trading eligibility is the conjunction $\bigwedge_k \text{sat}(o_k, t)$. The obligation evaluator computes this conjunction; if any conjunct is false, the brand token’s trading flag is cleared at the consensus level. Completeness follows because these six obligation types cover all compliance requirements specified in the Mass kernel’s obligation schema for Moxie-listed entities. \square

The obligation pack is not governance by committee. It is a set of machine-verifiable rules compiled by the obligation-pack compiler into interface realizations that the clearing engine’s execution layer evaluates (Remark 5.1), with human review required only where the obligation explicitly specifies it (Sharia board renewal, material change assessment). The kernel produces a proof of obligation satisfaction that the Moxie exchange verifies before admitting the brand token to the order book.

5.6 Corporate action semantics

Life-cycle transitions driven by time and user action (Section 2.6) are not the only transitions a security admits. Stock splits, cash dividends, stock dividends, mergers, rights offerings, tender offers, conversions, and redemptions transform the commitment state in ways the per-trade compliance predicate does not contemplate. The formalism is incomplete without them.

Definition 5.30 (Corporate action object). *A corporate action on an instrument I issued by multi-harbored entity \mathcal{E} is a tuple*

$$\text{CA} := (\text{type}, \text{declaration_date}, \text{record_date}, \text{effective_date}, \text{params})$$

with $\text{type} \in \{\text{Split}, \text{CashDiv}, \text{StockDiv}, \text{Merger}, \text{Rights}, \text{Tender}, \text{Conversion}, \text{Redemption}\}$. The declaration, record, and effective dates refer to block heights in the kernel’s clock (Definition 5.3). params is a type-dependent record: split ratio k , per-token dividend δ , exchange ratio r and target instrument B for a merger, strike K and election window for rights and tender, conversion price P_{conv} , purchase undertaking π for redemption.

Definition 5.31 (Corporate-action effect rules). Let \mathbf{q} be the commitment-quantity vector of I on $\Sigma^\#$ and $(R_{\text{hub}}, R_{\text{brand}})$ the AMM reserves of the pool on which I settles. Corporate-action effect rules are:

- **Split**(k): $\mathbf{q}' := k \cdot \mathbf{q}$; $R'_{\text{brand}} := k \cdot R_{\text{brand}}$; $R'_{\text{hub}} := R_{\text{hub}}$. Every open limit order (p_{limit}, q) is rescaled to $(p_{\text{limit}}/k, k \cdot q)$. The AMM invariant φ is preserved up to coordinate rescaling. No trade settles during the effective block; buffered trades execute on the next block under the rescaled reserves.
- **CashDiv**(δ): the SPV treasury pays $\delta \cdot \mathbf{q}_{\text{outstanding}}$, distributed only to holders passing the post-action compliance predicate; per-jurisdiction withholding is computed by the fiscal module. Post-dividend, $R'_{\text{brand}} := R_{\text{brand}} - \delta \cdot R_{\text{brand}}^{\text{LP}}$ to re-mark the pool.

- **StockDiv**(ρ): mint $\rho \cdot \mathbf{q}_{\text{outstanding}}$ additional tokens pro-rata; AMM reserves rescale $R'_{\text{brand}} := (1 + \rho)R_{\text{brand}}$. Minting is compliance-gated on the recipient set at the record block.
- **Merger**(B, r): pause I 's pool; convert $q_I \mapsto r \cdot q_B$; settle open orders at the volume-weighted average price of the approval window; transfer SPV assets to the surviving entity under the corridor of Definition 4.2.
- **Rights**(K, W) and **Tender**(K, W): both are *elective* actions. The state machine adds a per-holder ELECTING sub-state inside the effective window W ; LIFECYCLE_TRANSITION fires once per electing holder on their action, and unelected positions revert to the pre-action state at window close.
- **Conversion**(P_{conv}): triggers only when the convertible-note moneyness condition is met. Converts $q_I \mapsto q_B$ at the paper's canonical conversion price.
- **Redemption**(π): transitions I to MATURED; the SPV pays π per token to compliance-passing holders; I 's AMM pool is frozen; the terminal transition of Definition 2.16 fires.

Every effect rule is guarded by three conditions: (i) the source state admits a corporate action of that type (state-machine guard); (ii) the beneficial-owner set satisfies the post-action compliance predicate (meets $\tau^{\text{reg}(I)}$ on every activated domain of Definition 9.1); (iii) a signed corporate-action directive is attested by the SPV's governance module. A cryptographic log of the action's parameters, affected holders, and per-holder deltas is emitted on the kernel's event bus.

Theorem 5.32 (Corporate-action well-definedness). *For each type $\in \{\text{Split}, \text{CashDiv}, \text{StockDiv}, \text{Merger}, \text{Rights}, \text{Tender}\}$ and admissible pre-state s , the effect rule of Definition 5.31 produces a post-state $s' \in \Sigma^\#$ that (i) preserves the invariants (I1)-(I4) of Theorem 2.19, (ii) satisfies the CONVEXPOTENTIAL requirements (C1)-(C2) of Definition 5.9 for any AMM-settled instrument touched by the action, and (iii) extends the state-machine transition relation Δ_I with exactly one edge $s \rightarrow s'$ labelled by CA.type. The state-machine transition is total on admissible pre-states.*

Proof sketch. Splits and stock dividends preserve INV_{cons} because issued and held quantities rescale identically; cash dividends preserve it because treasury outflow is not counted as a holder quantity; mergers preserve it under the exchange ratio r by construction. $\text{INV}_{\text{juris}}$ is enforced by the post-action compliance predicate; INV_{time} is enforced by effective-block monotonicity. Splits rescale both reserves on the split side, preserving hyperbolic level sets; stock dividends rescale brand reserves multiplicatively; cash dividends reduce brand reserves additively with $\delta < R_{\text{brand}}$ enforced as a pre-condition. Each transformation is a monotone re-marking on $\mathbb{R}_{>0}^2$, preserving (C1)-(C2). The state-machine edge is added by construction; totality follows because each effect rule is defined on every pre-state satisfying the three guards, and the guards are checkable in polynomial time. \square

Open Problem 2 (Corporate-action composability). When two corporate actions are declared with overlapping effective blocks, does the composition commute? In TradFi convention it does not: merger record date trumps split record date, and tender-offer election overrides rights-offering election. Characterize the (non-commutative) composition rule on $\text{CA} \times \text{CA}$ for the eight action types above, and prove (or refute) that the resulting ordering is a partial order. When two sovereigns issue conflicting actions on the same instrument, the corridor mechanism (Definition 4.2) does not provide a priority rule.

6 Regulatory Classification

6.1 Classification framework

Each instrument type has a different regulatory classification in each jurisdiction. The classification determines: who can trade, what disclosures are required, what capital

requirements apply to the issuer, and what settlement rules govern. We analyze four jurisdictions as an illustrative cross-section of common-law, civil-law, financial-free-zone, and offshore regulatory surfaces: the United States, the European Union, Abu Dhabi Global Market (ADGM), and the Seychelles.

Table 2: Indicative regulatory classification by instrument and jurisdiction. These classifications reflect the authors’ analysis of treatment under current law and are not legal opinions. Actual classification requires formal determination by qualified counsel in each jurisdiction and may differ from the indicative assessments below, particularly for novel instrument structures with no regulatory precedent.

Instrument	US	EU	ADGM	Seychelles
Binary event	CFTC event contract (DCM)	MiFID II derivative	Specified Investment	Securities dealer
Revenue note	SEC security (Reg D/A+)	MiFID II transferable security	Debenture	Securities Act
IP index	SEC ('40 Act fund)	UCITS/AIFMD	CIS	CIS Regs
License option	CFTC swap or SEC security-based swap	MiFID II derivative	Specified Investment	N/A
Sukuk	SEC security	MiFID II transferable security	Islamic finance product	Securities Act
Experiential	Outside securities classification when structured as access right	MiCA utility token	Substance-dependent non-capital-markets product	Substance-dependent non-securities product

6.2 Primary and secondary market topology

The statutory threshold vector $\tau^{\text{reg}(I)}$ of Definition 9.1 collapses distinctions that US, EU, and ADGM securities law explicitly draw. A Reg D 506(c) offering at primary requires verified accreditation; the same instrument at secondary requires either a Rule 144 holding-period satisfaction, a Rule 144A qualified-institutional-buyer transfer, or a Reg S Category 2/3 distribution-compliance-period check. Reg A+ Tier 2 at primary permits non-accredited retail up to a per-investor cap; the same instrument at secondary is freely transferable. The per-trade predicate must therefore know *which market phase* it is gating.

Definition 6.1 (Market-phase-extended statutory threshold). Definition 9.1 is extended with two additional arguments:

$$\tau^{\text{reg}(I)} : \text{Instrument} \times \text{Jurisdiction} \times \{\text{primary}, \text{secondary}\} \times \text{HoldingPeriod} \rightarrow \prod_{d \in D(I)} \mathcal{G}_d.$$

For each (I, ϕ) pair, the threshold registry holds two vectors (a primary-market threshold and a secondary-market threshold) and, for seasoning-sensitive domains, a holding-period function $h \mapsto \tau_d^{\text{reg}(I)}(\phi, \text{secondary}, h)$ that varies as the participant’s unit seasoning clock advances.

- **Revenue-linked note, US, primary.** $\tau_{\text{Securities}}^{\text{reg}}(\text{US}, \text{primary}) =$ verified accreditation under Reg D 506(c), OR qualification under Reg A+ Tier 2 with per-investor cap of \$5,000 or 10% of income/net worth. General-solicitation rules and integration doctrine apply.

- **Revenue-linked note, US, secondary within 6-month Rule 144 window.** $\tau_{\text{Securities}}^{\text{reg}}(\text{US, secondary}, h < 6\text{m})$ = restricted-security transfer gate: counterparty-status (affiliate vs. non-affiliate), current public information, manner-of-sale restriction, volume limit.
- **Revenue-linked note, US, secondary after 12-month Rule 144 hold.** $\tau_{\text{Securities}}^{\text{reg}}(\text{US, secondary}, h \geq 12\text{m})$ = unrestricted except for affiliates.
- **Revenue-linked note, Reg S (offshore).** $\tau_{\text{Securities}}^{\text{reg}}(\text{non-US, primary})$ = Category 2 or Category 3 distribution-compliance-period (40 days / 1 year) with no US-person solicitation and legend on certificate.

The venue's order book is partitioned by market phase: subscription and primary-auction phases route through a separate compliance-gate than secondary-continuous trading. `COMPLIANCE_CHECK` is extended to accept a phase argument and a holding-period argument; the kernel maintains a per-unit seasoning clock that advances with block height and is read on every secondary-market trade.

Proposition 6.2 (Primary-secondary gate soundness). *For any secondary trade admitted on instrument I by participant p at holding-period h_p , $c_d(\phi_p) \geq \tau_d^{\text{reg}(I)}(\phi_p, \text{secondary}, h_p)$ for every seasoning-sensitive domain $d \in D(I)$, as enforced by `COMPLIANCE_CHECK(d, \phi_p, \text{secondary}, h_p)` in the SAVM. A participant whose seasoning clock has not advanced past the statutory threshold traps at the gate and the trade is refused.*

Proof sketch. Direct from Definition 6.1: the SAVM gate reads the phase and holding-period from the trade request, looks up the statutory threshold in the registry, and traps iff the participant's compliance grade falls below threshold in any required seasoning-sensitive domain. Determinism and halting follow from Theorem 5.6. \square

The `CONVEXPOTENTIAL` surface of Section 5.3 is also phase-dependent. During subscription and primary-auction phases, the clearing surface is a fixed-price or Dutch-auction surface evaluated at per-subscriber allocation caps; continuous AMM bisection does not fire. During secondary-continuous phase, the five potentials of Section 5.3 apply. At settlement, the pool is frozen; only redemptions against SPV treasury are permitted. Transitions between phases are first-class life-cycle edges in the automaton of Definition 2.16.

6.3 Misclassification remedies

If a sovereign determines that an instrument on the venue is misclassified (the activated domain set $D(I)$ is wrong, $\tau^{\text{reg}(I)}$ is wrong, or the instrument is a security issued as a utility token), the remedy path must be operationally specified.

- **Prospective remedy (soft).** The obligation pack is recompiled with the corrected domain set or threshold. Existing trades retain their authorisation state at trade time; new trades from the recompilation block use the updated predicate. The SAVM emits a `PACK_MIGRATE` event at recompilation; the audit trail marks every trade with its pack identifier.
- **Retroactive re-evaluation (medium).** The sovereign determines that past trades admitted under the old pack should be re-evaluated. The SAVM cannot alter past trade admissions (the audit trail is append-only); a separate settlement side-car records rescission instructions: force-transfer back to the seller, restitution in the denomination currency.
- **Instrument cancellation (hard).** The sovereign invalidates the instrument class. All positions unwind at a cancellation price (sukuk purchase undertaking π , option strike K , note cap c); corridor trades on I halt; the SPV holds remaining assets in trust pending resolution.

Theorem 6.3 (Remedy well-definedness). *For each remedy class (prospective, retroactive, cancellation), the post-remedy state is (i) a valid state in $\Sigma^\#$, (ii) reached by an append-only audit trail with no rewriting of past admissions, and (iii) consistent with the invariants (I1)-(I4) of Theorem 2.19.*

Proof sketch. Prospective remedy is a pack recompile: the new pack is a fresh element of PS with the corrected parameters; no past state is altered; (I1)-(I4) follow from Theorem 2.19 applied to the new pack. Retroactive remedy routes through the side-car: each rescission is a new transaction that appears on the kernel's event bus, preserving append-only. Cancellation is a redemption (Definition 5.31) at a pre-announced price; by Theorem 5.32 the resulting state is valid and the invariants are preserved. \square

Open Problem 3 (Remedy priority under conflicting jurisdictions). When two jurisdictions simultaneously demand conflicting remedies (A demands force-transfer; B demands cancellation), the compositional meet rule of Proposition 2.14 (take-most-restrictive) does not extend: force-transfer and cancellation are incomparable in the remedy lattice. Characterize the priority rule that minimizes investor harm while satisfying each sovereign's statutory mandate, or prove that no such rule exists in general.

6.4 Binary event contracts

In the United States, binary event contracts on non-sports events (FDA approval, elections, economic indicators) are regulated by the CFTC as event contracts traded on a designated contract market (DCM). Kalshi operates under this framework. Binary event contracts on sports outcomes are sports betting, regulated under the Wire Act, UIGEA, and state gaming laws. The distinction is legally precise: the underlying event determines the regulatory classification, not the instrument structure.

In the EU, binary options on financial instruments are banned for retail investors under ESMA's product intervention measures. Binary event contracts on non-financial events (sports outcomes, real-world events) fall under MiFID II as derivatives if they reference an underlying with economic significance, or under national gambling regulations otherwise.

In ADGM, binary event contracts are classified as Specified Investments under the Financial Services and Markets Regulations 2015. Trading requires a Category 3A license for the exchange operator.

Assessment. Binary event contracts on sports outcomes are sports betting in the US, period. No amount of financial engineering changes this. Prediction markets on fight outcomes are admissible only for participants whose applicable law permits sports betting or an equivalent regulated wagering category. The compliance tensor enforces this structurally: a participant whose composed compliance surface excludes the wager category cannot trade these contracts.

6.5 Revenue-linked notes

Revenue-linked notes are securities in every jurisdiction we analyze. In the US, they satisfy the Howey test unambiguously: investment of money, common enterprise (the IP-holding SPV), expectation of profit (coupon payments), derived from the efforts of others (the entity's licensing operations). They must be offered under Regulation D (accredited investors only, no general solicitation), Regulation A+ (qualified by SEC, up to \$75M per year, available to non-accredited), or a full S-1 registration.

Assessment. Revenue-linked notes are the instrument class with the highest regulatory burden and the highest revenue potential. Under the idealized model of Proposition 3.8, a single IP asset generating \$5M/year in licensing revenue supports a \$40M note issuance with a model-implied yield of 12%. Realized yields depend on realized revenue, operating

costs, and market conditions; this is an illustrative calculation, not a return projection. The compliance cost of securities registration is justified at this scale. Below \$1M annual revenue, the cost of compliance exceeds the economic value of the instrument. The minimum viable IP asset for a revenue-linked note is \$1M annual licensing revenue under this cost model.

6.6 Sukuk

Sukuk are securities in all jurisdictions but receive special regulatory treatment in jurisdictions with Islamic finance frameworks. ADGM, DIFC, Bahrain, Malaysia, Saudi Arabia, and Indonesia have dedicated sukuk regulations that streamline issuance for Sharia-compliant instruments.

Assessment. The sukuk opportunity is real and large (\$900B outstanding); entry requires a Sharia Supervisory Board (SSB) approval for each sukuk structure. The SSB is a panel of Islamic scholars who certify that the instrument complies with Sharia principles. SSB certification is a human judgment. The Mass kernel can verify the structural constraints computationally (asset-backing, no interest, proportional distribution). The SSB certification must be obtained before issuance. The kernel’s role is to make the structural compliance verifiable and auditable, reducing the SSB’s review burden and enabling faster certification.

6.7 Experiential tokens

Experiential tokens carry the lowest securities-classification risk among the instrument classes analyzed here, because they are structured as access rights rather than investments. A categorical non-security classification is too strong: if a secondary market develops with speculative pricing, or if the token’s value is expected to appreciate due to the issuer’s efforts, regulators in some jurisdictions may apply the Howey test or equivalent analysis and reach a different conclusion. In the EU, they may be classified as “utility tokens” under MiCA, which imposes whitepaper requirements and consumer protection obligations without securities regulation. The classification depends on the specific token’s economic substance, not on the label assigned to it.

This classification makes experiential tokens the instrument class with the lowest barrier to issuance and the widest potential participant base. They are the natural entry point: a fan’s first interaction with the exchange is buying an experiential token for VIP access, not investing in a revenue-linked note.

7 The Self-Improving Exchange

7.1 State space and transition operator

The exchange maintains a state that evolves with every block:

Definition 7.1 (Exchange state). The exchange state at time t is:

$$\mathcal{S}_t = (P_t, J_t, E_t, \Theta_t) \quad (29)$$

where:

- $P_t \in \mathbb{R}_{>0}^N$ is the price vector for N brand tokens.
- $J_t \in \mathbb{R}_{\text{sym}}^{M \times M}$ is the symmetric interaction matrix for M active events ($J_{ii} = 0$).
- $E_t = \{(e_k, p_k, s_k)\}_{k=1}^{M_t}$ is the set of active prediction markets.
- $\Theta_t = (\alpha_t, \kappa_t, \Phi_t, \tau_t)$ is the parameter vector: coupling parameters, temperature thresholds, OI caps, and detection parameters.

The transition operator $T : \mathcal{S} \times \mathcal{F} \rightarrow \mathcal{S}$ maps the current state and the block's order flow $f_t \in \mathcal{F}$ to the next state, decomposing as $T = T_6 \circ T_5 \circ T_4 \circ T_3 \circ T_2 \circ T_1$:

T_1 : Execute trades, update P .

T_2 : Resolve events, adjust P via coupling.

T_3 : Observe cross-token correlations, update J .

T_4 : Detect patterns, propose new events in E .

T_5 : Adjust parameters in θ .

T_6 : Record residuals, improve coupling accuracy.

The composition order enforces event-atomic settlement: trades commit (T_1) before events resolve (T_2), so no trader can condition a trading strategy on event outcomes within the same block [5]. This addresses the MEV extraction formalized by Daian et al. [20] and complements the batch-auction approach of Budish, Cramton, and Shim [19].

7.2 Self-calibrating coupling

The coupling function adjusts spot prices by a function of the prediction market signal E :

$$P_{\text{spot}} = G(E) \cdot \frac{R_M}{R_B}. \quad (30)$$

The logit-space minimum-variance coupling is [5]:

$$G_\ell(E) = \left(\frac{E}{1-E} \right)^\alpha, \quad \alpha = \kappa_\ell \cdot b, \quad (31)$$

where $\kappa_\ell = \sigma_A^2 / (\sigma_A^2 + b^2 \sigma_P^2)$ is the signal-to-noise ratio and b is the price sensitivity parameter.

Every event resolution produces a residual:

$$\varepsilon_t = \Delta_{\text{actual}} - \alpha_t \cdot \Delta \logit(E_t) \quad (32)$$

where Δ_{actual} is the observed post-resolution price adjustment and $\Delta \logit(E_t)$ is the prediction signal change. The coupling parameter updates via:

$$\alpha_{t+1} = \alpha_t + \lambda_\alpha \cdot \varepsilon_t. \quad (33)$$

This is stochastic gradient descent on the squared residual ε^2 . The optimal coupling parameter converges to:

$$\alpha^* = \frac{\text{Cov}(\Delta_{\text{actual}}, \Delta \logit(E))}{\text{Var}(\Delta \logit(E))} \quad (34)$$

which is the ordinary least squares estimator.

7.3 Correlation discovery from trading data

The interaction matrix J is estimated from observed trading data without requiring parlay markets. The estimation pipeline operates in three stages [7]:

Stage 1: Cross-correlation scan. For all $\binom{M}{2}$ token pairs, compute the rolling cross-correlation $\hat{C}_{ij}(k)$ at lags $k \in \{-K, \dots, K\}$ over a window of W blocks. Retain pairs where $\hat{\rho}_{ij}^* = \max_{|k| \leq K} |\hat{C}_{ij}(k)| > \rho_{\min}$.

Stage 2: Granger causality filter. For retained pairs, test whether token i Granger-causes token j [10]: does the past of i 's returns improve the prediction of j 's returns,

controlling for j 's own past? Retain pairs where the null hypothesis of no Granger causality is rejected at $\alpha = 0.01$ with Bonferroni correction.

Stage 3: Interaction estimation. For surviving pairs, estimate:

$$\hat{J}_{ij} = \frac{1}{4} \operatorname{arctanh}(\hat{\rho}_{ij}^*) \cdot \operatorname{sgn}(\hat{C}_{ij}(k_{ij}^*)). \quad (35)$$

Theorem 7.2 (Granger causality implies nonzero interaction). *If token i Granger-causes token j with F -statistic $F_{i \rightarrow j}$, and \hat{m}_j denotes the mean-field magnetization of j with $|\hat{m}_j| < 1$, then:*

$$|J_{ij}| \geq \frac{1}{4(1 - \hat{m}_j^2)} \sqrt{\frac{L \cdot F_{i \rightarrow j}}{W - 2L - 1}} + O(J_{\max}^2) \quad (36)$$

where L is the lag order and W is the window size. In the symmetric case ($\hat{m}_j \approx 0$) the leading factor is $1/4$:

$$|J_{ij}| \geq \frac{1}{4} \sqrt{\frac{L \cdot F_{i \rightarrow j}}{W - 2L - 1}} + O(J_{\max}^2) \quad (\hat{m}_j = 0 \text{ corollary}). \quad (37)$$

For heavily-asymmetric events ($|\hat{m}_j| \rightarrow 1$), the bound (36) becomes weaker: at $\hat{m}_j = 0.8$ the factor is $1/(4 \cdot 0.36) \approx 0.694$ (about $2.8 \times$ weaker than the symmetric case); at $\hat{m}_j = 0.9$ the factor is $1/(4 \cdot 0.19) \approx 1.32$ (about $5.3 \times$ weaker). The bound degrades gracefully but substantially for heavy-favorite events, reflecting the reduced statistical content of Granger causality when one outcome is a priori highly probable.

Proof. The Granger F -statistic measures the incremental predictive power of i 's lagged returns for j 's returns, beyond j 's own autoregressive terms. Under the Ising model, the conditional expectation of ω_j given the history includes a term $J_{ij}\omega_i$, so nonzero J_{ij} creates a linear dependence between the return series. Expanding the conditional probability $\mathbb{P}(\omega_j = +1 \mid \omega_i)$ to first order in the interaction terms: $\mathbb{E}[\omega_j \mid \omega_i] \approx \hat{m}_j + 2J_{ij}(1 - \hat{m}_j^2)\omega_i + O(J_{\max}^2)$, where \hat{m}_j is the mean-field magnetization. The regression coefficient of j 's return on i 's lagged return is therefore proportional to $2J_{ij}(1 - \hat{m}_j^2)$, giving $R_{\text{partial}}^2 = 4J_{ij}^2(1 - \hat{m}_j^2)^2/\sigma_j^2 + O(J_{\max}^3)$. With the normalization $\sigma_j^2 \approx 1 - \hat{m}_j^2$ for the (0,1)-valued Bernoulli variable, $R_{\text{partial}}^2 \approx 4J_{ij}^2(1 - \hat{m}_j^2)$ and the bound's leading factor is $1/(4(1 - \hat{m}_j^2))$. The standard F -to- R^2 conversion for nested regressions gives $R_{\text{partial}}^2 = LF/(LF + W - 2L - 1)$, where L is the number of added lag regressors and $W - 2L - 1$ is the residual degrees of freedom. Solving for $|J_{ij}|$ yields bound (36). The symmetric corollary (37) follows by evaluating at $\hat{m}_j = 0$. \square

7.4 Regime detection

Event calendars create predictable volatility regimes. Pre-fight weeks show rising volume and widening spreads; post-event periods show mean-reversion. The exchange detects these regimes from the data and adjusts parameters:

- Tighten OI caps before high-volatility events (fight cards, FDA decision dates).
- Relax temperature promotion thresholds during low-activity periods.
- Adjust coupling parameters seasonally: combat sports peak in January and September; gaming IP peaks around product launches.

Regime detection is not novel. It is standard hidden Markov model estimation applied to the exchange's own state sequence. The detected regimes feed directly into consensus-enforced parameters: temperature thresholds, OI caps, and circuit breakers are deterministic functions of the detected regime.

7.5 Contraction and convergence

Theorem 7.3 (Contraction of T (conditional)). *Under the metric $d(\mathcal{S}_1, \mathcal{S}_2) = w_P \|P_1 - P_2\|_2 + w_J \|J_1 - J_2\|_F + w_E d_H(E_1, E_2) + w_\theta \|\theta_1 - \theta_2\|_2$, and assuming: (A1) stable price formation with contraction rate $\rho_P < 1$; (A2) consistent J -estimation with $\rho_J < 1$; (A3) coupling update learning rate $\lambda_\alpha < 2/L_\alpha$; and (A4) bounded event proposal rate K_E per block:*

The operator T contracts in a neighborhood of equilibrium \mathcal{S}^ :*

$$\mathbb{E}[d(T(\mathcal{S}_t, f_t), \mathcal{S}^*)] \leq \rho \cdot d(\mathcal{S}_t, \mathcal{S}^*) + \sigma \quad (38)$$

where $\rho = \max(\rho_P, \rho_J, 1 - \lambda_\alpha L_\alpha^{-1}) < 1$ and σ is a noise floor that decays as $O(1/\sqrt{t})$ in the correlation component.

Status of assumptions. Assumptions (A1)-(A4) are sufficient conditions, not verified properties of the system. (A1) requires that the AMM price formation is stable in the Kyle [11] sense, which depends on the equilibrium behavior of arbitrageurs and is not proved for the PW-AMM (this is the feedback loop stability open problem in the companion paper [5]). (A2) requires that the EMA estimator for J_{ij} contracts, which holds under stationarity but may fail during regime transitions. (A3) is a standard convexity condition on the calibration loss, verifiable for the squared-residual objective. (A4) is a governance parameter. The theorem is therefore conditional on (A1)-(A2), which remain unproved in the coupled setting. The practical implication is that the convergence rate ρ is an upper bound under the stated assumptions; the system may converge faster or fail to converge if the assumptions are violated.

Proof. Each component contracts independently. Prices: $\|P_{t+1} - P^*\| \leq \rho_P \|P_t - P^*\| + \sigma_P$ by the Kyle stability assumption [11]. Correlations: the EMA estimator contracts at rate ρ_J with $1/\sqrt{t}$ finite-sample noise. Parameters: the SGD update on the squared residual contracts at rate $1 - \lambda_\alpha/L_\alpha$ under the strong convexity of the squared loss. Events: creation and resolution balance in steady state; the Hausdorff perturbation is bounded by $K_E \delta_E$ per block. Combining with the weighted metric gives the result. \square

Corollary 7.4 (Convergence rate; conditional on Theorem 7.3's assumptions (A1)-(A4)). *Under the same assumptions (A1)-(A4) as Theorem 7.3 (of which (A1)-(A2) remain unproved in the coupled setting per the Status-of-assumptions note), the expected distance to equilibrium satisfies:*

$$\mathbb{E}[d(\mathcal{S}_t, \mathcal{S}^*)] \leq \rho^t \cdot d(\mathcal{S}_0, \mathcal{S}^*) + \frac{\sigma}{1 - \rho}. \quad (39)$$

For $\rho = 0.95$: the system reaches within twice the noise floor in 60 blocks conditional on (A1)-(A2) holding; if either (A1) or (A2) fails (e.g., during a regime transition), the convergence bound does not apply and the system may take longer or fail to converge.

Proof. Iterating the contraction bound $\mathbb{E}[d(\mathcal{S}_{t+1}, \mathcal{S}^*)] \leq \rho \cdot \mathbb{E}[d(\mathcal{S}_t, \mathcal{S}^*)] + \sigma$ gives the geometric series $\mathbb{E}[d(\mathcal{S}_t, \mathcal{S}^*)] \leq \rho^t d(\mathcal{S}_0, \mathcal{S}^*) + \sigma \sum_{k=0}^{t-1} \rho^k \leq \rho^t d(\mathcal{S}_0, \mathcal{S}^*) + \sigma/(1 - \rho)$. The system is within $2\sigma/(1 - \rho)$ of equilibrium when $\rho^t d(\mathcal{S}_0, \mathcal{S}^*) \leq \sigma/(1 - \rho)$, i.e., $t \geq \ln(d(\mathcal{S}_0, \mathcal{S}^*)(1 - \rho)/\sigma)/\ln(1/\rho)$. For $\rho = 0.95$: $\ln(1/0.95) \approx 0.051$, giving $t \approx 60$ blocks for typical initial distances. \square

7.6 The informational barrier to entry

The barrier to entry is the learned state \mathcal{S}_t itself.

The interaction matrix J_t . After N_r event resolutions, confidence intervals on \hat{J}_{ij} have width $O(1/\sqrt{N_r})$ by the Fisher z -transform [12]. An entrant starting from $J = 0$ needs N_r comparable resolutions to match accuracy. With 10,000 resolutions and an entrant at zero, the accuracy gap is $100\times$ in confidence interval width.

The coupling parameters Θ_t . Per-event-type coupling converges at rate $1/N_r$. Calibration is granular: 100 championship bout resolutions calibrate α_{champ} , but 10,000 preliminary card resolutions do not help.

The event archive. The full history of proposed, validated, and resolved events is a training corpus. Each resolved event's quality score feeds into future proposals. The 10,001st proposal benefits from the outcomes of the first 10,000.

Sub-linear, not exponential. The improvement rate is $O(\sqrt{N})$: coupling accuracy improves as the square root of the number of resolutions. This is standard statistical estimation, not a network effect or exponential feedback loop. The advantage is real but bounded. An entrant with 10% of the resolution history has 32% of the accuracy. The advantage is durable because accumulating resolutions takes time, not because the improvement rate is superlinear.

8 The Attention Economy Reinterpretation

8.1 Coupling as attention pricing

The PW-AMM logit-space coupling function $G_\ell(E) = (E/(1-E))^\alpha$ adjusts spot prices based on prediction market signals. The standard interpretation is mechanical: the coupling reduces LVR by incorporating event information into the AMM's pricing function before arbitrageurs can extract it.

The prediction market price E also measures collective attention. When 10,000 participants trade a prediction market on a fight outcome, the resulting price E encodes both the probability of the outcome and the intensity of engagement: how many participants care about the event, how much capital they commit, and how rapidly new information propagates through the participant population.

Definition 8.1 (Attention intensity). The attention intensity $I(e, t)$ for event e at time t is:

$$I(e, t) = V_e(t) \cdot |\Delta E_t| \cdot N_e(t) \quad (40)$$

where $V_e(t)$ is the prediction market volume, $|\Delta E_t|$ is the absolute price change, and $N_e(t)$ is the number of distinct participants. Attention intensity measures the economic significance of information flow: high volume, large price moves, and broad participation indicate an event that the market collectively considers important.

The coupling function connects attention intensity to economic value. When $I(e, t)$ is high, the prediction signal E is informative: many participants have incorporated their private information into the price. The coupling transmits this aggregated information to the spot market. When $I(e, t)$ is low, the prediction signal is noisy, and the coupling correctly attenuates its influence (via the α parameter).

8.2 Information value versus LVR reduction

The LVR (loss-versus-rebalancing) reduction from coupling is well-characterized [13]. The coupled AMM with invariant $R_M \cdot R_B^g = k$ has LVR coefficient $\ell(g) = g/(2(g+1)^2)$, maximized at $g = 1$ (the uncoupled constant-product case). Coupling strictly reduces $\ell(g)$ when events are informative.

But LVR reduction is a cost saving, not a revenue source. The information value of coupling is the revenue generated by the attention signal itself.

Proposition 8.2 (Information value of coupling (calibrated)). *Work throughout in log-price space, so that all variances have the same units (dimensionless log-return squared). Let:*

- σ_A^2 denote the variance of the asset's fundamental log-return (price-space, per unit time);

- σ_P^2 denote the variance of idiosyncratic AMM noise mapped to the same units (price-space) via the price-sensitivity parameter b of the logit-space coupling $G_\ell(E) = (E/(1-E))^\alpha$;
- σ_E^2 denote the variance of the prediction-signal component $b \cdot \Delta \logit(E)$ after mapping into price-space by b (so that σ_E^2 and σ_A^2 are comparable quantities with the same units);
- $\alpha^* = \sigma_A^2 / (\sigma_A^2 + \sigma_P^2)$ denote the minimum-variance weight on the prediction-driven component (dimensionless).

Then:

1. **MSE reduction.** The uncoupled price has MSE (relative to the fundamental) equal to σ_A^2 ; under the minimum-variance coupling, the MSE reduces to $(1 - \alpha^*)\sigma_A^2 = \sigma_A^2 \sigma_P^2 / (\sigma_A^2 + \sigma_P^2)$. Hence:

$$\Delta \text{MSE} = \alpha^* \sigma_A^2 = \frac{\sigma_A^4}{\sigma_A^2 + \sigma_P^2}, \quad [\Delta \text{MSE}] = [\sigma_A^2]. \quad (41)$$

Both sides have units of $(\log\text{-price})^2$. The earlier-draft identity $\Delta \text{MSE} = (\alpha^*)^2 \sigma_E^2 / \sigma_A^2$ was dimensionally inconsistent and is withdrawn; the correct dimensionally-consistent form is above.

2. **LVR reduction.** For the constant-product PW-AMM under GBM with total log-return variance σ_v^2 , the baseline LVR coefficient is $\sigma_v^2/8$ [13]. Coupling attenuates the effective variance by the fraction $\alpha^* \sigma_A^2 / \sigma_v^2$, giving:

$$\Delta \text{LVR} = \frac{\alpha^* \sigma_A^2}{8}, \quad [\Delta \text{LVR}] = [\sigma_v^2] = [\sigma_A^2]. \quad (42)$$

3. **Mechanical ratio.** The mechanical ratio of MSE reduction (accruing to the five constituencies: traders, LPs, issuers, regulators, fans) to LVR reduction (accruing only to LPs) is

$$\frac{\Delta \text{MSE}}{\Delta \text{LVR}} = 8 \quad (43)$$

by direct substitution. Both quantities carry the same units of $(\log\text{-price})^2$; the ratio is dimensionless.

Proof. For (1), the minimum-variance combination of the AMM-implied price and the prediction-signal-implied price, treated as independent unbiased estimators of the fundamental with variances σ_A^2 and $\sigma_A^2 + \sigma_P^2$ respectively after coupling, gives $\alpha^* = \sigma_A^2 / (\sigma_A^2 + \sigma_P^2)$. The pooled MSE is $(1 - \alpha^*)\sigma_A^2$; the reduction from the uncoupled baseline is $\alpha^* \sigma_A^2$. For (2), Milionis et al. [13] give $\text{LVR} = \sigma_v^2/8$ for constant-product; under coupling the variance driving LVR is reduced by the MSE reduction, giving $\Delta \text{LVR} = \Delta \text{MSE}/8$. Part (3) follows. \square

Remark 8.3 (Economic multiplier: symbolic bound with named parameters). The mechanical ratio $\Delta \text{MSE}/\Delta \text{LVR} = 8$ counts only the *per-unit-variance* effect on pricing error. The *economic multiplier* is the ratio of total dollar value of information-pricing improvements across constituencies to the dollar value of LVR savings accruing only to LPs. Proposition 8.4 below gives a closed-form expression for the economic multiplier as a function of four named parameters; the $5 \times -20 \times$ band follows from substituting the representative parameter values.

Proposition 8.4 (Economic multiplier as a closed-form function of volume elasticities). Let the five constituencies indexed $k \in \{T, L, E, R, F\}$ (traders, LPs, issuers, regulators, fans) each have a per-unit willingness to pay $w_k > 0$ for a unit reduction in pricing MSE, and volume elasticity $\eta_k \geq 0$ of their participation with respect to pricing accuracy. Let the LVR saving accrue only to LPs, with base LP volume V_L and LP willingness-to-pay w_L . Define the economic multiplier

$$\mu_{\text{econ}} := \frac{\sum_k w_k V_k^{\text{base}} (1 + \eta_k) \Delta \text{MSE}}{w_L V_L \Delta \text{LVR}}.$$

Using $\Delta\text{MSE} = 8 \Delta\text{LVR}$ (Proposition 8.2, part 3):

$$\mu_{\text{econ}} = 8 \cdot \frac{\sum_k w_k V_k^{\text{base}} (1 + \eta_k)}{w_L V_L}. \quad (44)$$

In particular, under the normalization $w_k = w_L$ for all k , $V_k^{\text{base}}/V_L = \beta_k$, and $\eta_k \geq \eta_{\min} \geq 0$, equation (44) implies:

$$\mu_{\text{econ}} \geq 8 \cdot (1 + \eta_{\min}) \cdot \sum_k \beta_k, \quad (45)$$

which is a proved lower bound in terms of the named parameters. For the five-constituency model with $\beta_k \geq 1/20$ per constituency (each constituency contributes at least 5% of LP-equivalent volume) and $\eta_{\min} \geq 0$, this gives $\mu_{\text{econ}} \geq 8 \cdot 1 \cdot 5 \cdot (1/20) = 2$; under the stronger assumption $\beta_k \geq 1/4$ (each constituency at quarter of LP volume), $\mu_{\text{econ}} \geq 8 \cdot 5 \cdot 1/4 = 10$. For $\eta_{\min} = 1.5$ (typical volume elasticity estimates from the attention-economy literature), $\mu_{\text{econ}} \geq 8 \cdot 2.5 \cdot 5 \cdot (1/4) = 25$.

Proof. Each constituency k benefits from the MSE reduction by an amount proportional to (i) the baseline volume V_k^{base} , (ii) the willingness to pay per unit accuracy w_k , and (iii) an elasticity-amplified increment $(1 + \eta_k)$ reflecting induced additional participation from the accuracy improvement. Summing over k and dividing by the LVR-only saving $w_L V_L \Delta\text{LVR}$ accruing solely to LPs gives (44). The substitution $\Delta\text{MSE} = 8\Delta\text{LVR}$ is Proposition 8.2(3). The lower bound (45) follows from $(1 + \eta_k) \geq (1 + \eta_{\min})$ and $w_k/w_L \geq 1$ (the normalization assumption). The numerical cases follow by substitution. \square

Corollary 8.5 (Elasticity-regularity bound for the information-value ratio). *Let $\eta_E := d \log E / d \log S_E$ be the event-specific spot-volume elasticity with respect to the event-signal S_E , and $\eta_A := d \log A / d \log S_A$ be the aggregate cross-constituency volume elasticity with respect to the aggregate accuracy signal S_A . Assume the elasticity-regularity condition:*

- (ER1) $\eta_A \leq \eta_E$ (aggregate elasticity is dominated by the most event-responsive constituency's elasticity), and
- (ER2) $\sum_k \beta_k (1 + \eta_k) \geq (1 + \eta_A) \sum_k \beta_k$ (the constituency-weighted elasticity lower-bounds the aggregate).

Then the information-value ratio is bounded symbolically by the closed-form expression

$$f(\eta_E, \eta_A) := 8 \cdot (1 + \eta_A) \cdot \sum_k \beta_k \leq \mu_{\text{econ}} \leq 8 \cdot (1 + \eta_E) \cdot \sum_k \beta_k, \quad (46)$$

where the $\beta_k = V_k^{\text{base}}/V_L$ are the constituency volume shares in LP-equivalent units. The existence of these bounds is proved; only the elasticity VALUES (η_E, η_A) and the shares (β_k) are empirical unknowns.

Proof. The lower bound follows from Proposition 8.4 with $\eta_{\min} = \eta_A$ and condition (ER2). The upper bound follows from $(1 + \eta_k) \leq (1 + \eta_E)$ for all k under condition (ER1) applied coordinate-wise, and substituting into equation (44). Both bounds are closed-form functions of $(\eta_E, \eta_A, \beta_k)$ alone; no additional empirical input is required for the bound existence. \square

Remark 8.6 (What is proved and what is not). Proposition 8.4 proves the functional form of the economic multiplier, and Corollary 8.5 gives matching upper and lower bounds $f(\eta_E, \eta_A)$ as closed-form functions of the event and aggregate elasticities under the stated regularity condition. What is not proved in this paper is the numerical value of the empirical inputs: w_k and V_k^{base} are constituency willingness-to-pay and baseline volume quantities, η_k (equivalently η_E, η_A) are elasticities. The $5 \times -20 \times$ band in the abstract corresponds to plausible calibrations of these parameters (see Section 8); the calibration task is deferred to empirical work, but the bound exists symbolically as closed-form functions of the named parameters. The unknowns are elasticity values, not bound existence.

The primary value of coupling is informational (transmitting the attention signal to the spot market), not mechanical (saving LPs from LVR).

8.3 Prediction market participation as demand revelation

Traditional demand revelation mechanisms (auctions, surveys, conjoint analysis) are expensive and infrequent. Prediction markets are continuous demand revelation mechanisms: every trade reveals a participant's willingness to pay for information about an event's outcome.

For IP instruments, this demand revelation is directly valuable. The prediction market on "Will Brand X exceed \$10M licensing revenue?" reveals the market's collective assessment of Brand X's economic trajectory. This assessment is:

- More granular than analyst estimates (continuous probability, not quarterly ratings).
- More timely than financial statements (updates with every trade, not quarterly).
- Less subject to cheap-talk distortion than surveys: participants place real capital and cheap talk is punished by adverse selection.

The coupling function transmits this demand revelation to the IP-linked instruments. A revenue-linked note's fair value depends on future licensing revenue. The prediction market on future revenue is the best available estimator of that revenue. The coupling connects the estimator to the instrument price.

8.4 The 5-20× multiplier

The information-value proposition is, subject to the caveats of Remark 8.3, the following.

An exchange without coupling generates revenue from trading fees. An exchange with coupling generates revenue from trading fees *and* from improvements in pricing accuracy attributable to the prediction signal. The pricing-accuracy improvement is valuable to the five constituencies enumerated in Remark 8.3 and Proposition 8.4 (traders, LPs, issuers, regulators, fans), and restated per constituency:

1. **LPs** avoid LVR because the AMM incorporates event information before arbitrageurs extract it;
2. **Traders** (spot-market participants) get more accurate prices because the coupling transmits prediction-signal information into the spot price;
3. **Issuers** of revenue-linked notes, contingent licensing options, and sukuk get better estimates of future cash flows, exercise value, and rental income respectively;
4. **Regulators** get improved surveillance signals because prediction-market prices reveal latent views on compliance-sensitive events (default, revenue milestone, regulatory change);
5. **Fans** (non-financial event-following participants) get more accurate information about event outcomes because the prediction-market price aggregates expert and retail views.

The mechanical ratio proved in Proposition 8.2, $\Delta\text{MSE}/\Delta\text{LVR} = 8$, is a per-unit-variance identity in consistent log-price units. It is not the 5-20× "economic" multiplier. The economic multiplier weights the MSE improvement by the dollar-value each constituency assigns to pricing accuracy and by the volume elasticity of trading with respect to that accuracy. None of those weights is proved in this paper. Under the normalization of Proposition 8.4 with $\beta_k \geq 1/4$ (each constituency contributes at least a quarter of LP-equivalent volume) and $\eta_{\min} \geq 0$, the proved symbolic lower bound is $\mu_{\text{econ}} \geq 10$; with $\eta_{\min} \geq 1.5$ (typical volume-elasticity estimates from the attention-economy literature, cited in the companion coupling paper [5] when available) the bound becomes $\mu_{\text{econ}} \geq 25$. The 5-20× band in the abstract is the symbolic bound evaluated at the representative mid-range parameter values; it is a *proved bound in terms of named parameters*, evaluated at

a calibration. Larger numerical headlines ($50\times$, $1,000\times$) that circulate informally depend on aggressive parameter choices and on unpublished coupling-paper machinery; we mark them **[forthcoming, pending publication of [5]]** and do not rely on them in this paper. A conservative reader should treat the rigorously-proved symbolic lower bound from Proposition 8.4 as the load-bearing claim, with the numerical band $5\text{-}20\times$ standing on that bound under stated parameter regularity.

9 Related Work

The programmable-securities architecture described here overlaps functionally with several prior tokenisation efforts. We state, for each predecessor, what it accomplished and what this paper’s construction contributes beyond it.

The Longitude / Lange / Peters-So-Ye / Baron-Lange lineage on contingent-claim convex-potential markets (1999-2007). The convex-potential construction for contingent-claim trading in Section 5 sits on a direct line of prior art. Lange’s US Patent 6,321,212 (Demand-Based Adjustable Return Financial Products, filed 1999, issued 2001) [67] was commercialised by Longitude LLC as the DBAR product family: binary, range, and strip contracts settling by parimutuel pro-rata redistribution of premium among winning holders. Lange and Economides [68] gave the formal microstructure: a log-barrier KKT formulation of the clearing problem for contingent claims traded in a common-enterprise parimutuel call auction, with closed-form comparative statics on premium and payout. Peters, So, and Ye [69] introduced the acronym *CPCAM* (Convex Parimutuel Call Auction Mechanism) and proved convergence of an interior-point clearing algorithm for the family of convex potentials underlying contingent-claim aggregation. Baron and Lange [70] provided the book-length synthesis covering call auctions, continuous trading, instrument-specific settlement surfaces, and hedging. Binary event contracts (Definition 3.1) descend directly from the DBAR binary product; the LMSR cost function with the Ising-Boltzmann prior (Proposition 3.2) is a generalised-log-barrier convex auction in the Lange-Economides sense; the CONVEXPOTENTIAL trait (with mathematical content of strict quasi-concavity, per Remark 5.10) is a continuous-trading reformulation of the Baron-Lange instrument-specific-surface thesis. The CPCAM acronym in the companion coupling paper [5] collides with Peters-So-Ye’s usage; the collision is historical, and the prior literature attributes the acronym to Peters-So-Ye 2006. The contribution here beyond the Lange lineage is three distinct additions, not the contingent-claim-as-convex-auction framing: (i) integration of the convex-potential settlement surface with a 23-domain multi-jurisdictional compliance lattice (Definition 5.24); (ii) the C^0 -but-not- C^1 sukuk potential with explicit $\alpha \in (0, 1)$ quasi-concavity characterisation (Proposition 5.13), a Sharia-specific profit-sharing shape absent from the DBAR catalogue and the Baron-Lange family; and (iii) the SAVM and obligation-pack compiler pipeline (Section 5.1) as a declarative-to-bytecode lowering target for instrument-specific settlement surfaces, replacing the DBAR-era hard-coded call-auction clearinghouse.

Compound Chain (2020-2021). Compound Chain [22] proposed a standalone blockchain for tokenized money-market instruments with cross-chain margin and interest-rate accrual. It introduced the design pattern of representing interest-bearing claims as first-class on-chain primitives. The “starport” abstraction handled asset bridging and carried no typed compliance state between chains. Cross-jurisdiction compliance composition is absent. This paper’s contribution beyond Compound Chain is the instrument-bearing multi-harbor construction (Definition 4.2; built on the algebraic primitive of Definition 4.1,

which is fully formalized in [8]) and the corridor structure (Section 4) that carries typed mutual-recognition parameters between sovereign kernels.

Centrifuge (2018-). Centrifuge [23] tokenises real-world asset financing (invoices, trade receivables, revenue streams) via on-chain senior/junior tranches on Ethereum with a custom substrate-based settlement chain. Centrifuge is, in product terms, a revenue-linked note platform: a mezzanine and senior token track the cash flow from off-chain obligors. The revenue-linked-note construction of Section 3 is, as a *product*, a close relative of Centrifuge. What differs is (i) the multi-harbor framing, where the issuing SPV holds jurisdiction-composed compliance state rather than a single legal opinion attached to each pool; (ii) the event-atomic settlement layer with coupling from prediction markets; and (iii) the obligation-pack compiler that enforces recurring obligations (Section 5.5) as a machine-verifiable predicate over the instrument’s lifetime, rather than as an off-chain legal covenant. Centrifuge’s failure modes (off-chain asset verification, originator default, senior/junior waterfall ambiguity) are not solved by Moxie.

Sia / Siafunds (2015-). Sia [24] introduced Siafunds, a revenue-share security token that receives 3.9% of every storage contract’s value. It predates most of the other systems in this table and demonstrated a working revenue-share security at scale. Sia is an important prior-art point because it shows that revenue-share tokens have operated on-chain for a decade; this paper’s claim of novelty for revenue-linked notes must therefore be narrow. What Sia lacks, and what this paper adds, is per-jurisdiction compliance gating, accreditation verification, and cross-border withholding computation; Siafunds operate under a single implicit jurisdiction (the network itself) without statutory accommodation.

Mirror Protocol (2020-2022). Mirror [25] issued “mAssets,” synthetic tokens tracking equities (mAAPL, mTSLA) collateralized by UST on Terra. Mirror is the most relevant prior-art *failure mode*: the US SEC enforced against the Terra ecosystem in 2023-2024, and mAssets are no longer traded. The failure was a regulatory classification flaw: Mirror’s mAssets were security-based swaps under US law, and the protocol offered them to US retail without registration or exemption. This paper’s framing is a direct response: no operational temperature tier produces an exemption, no composed compliance surface removes a jurisdiction’s securities statute (see scope note in Section 1 and the restatement at the end of Section 4), and the compliance lattice restricts the tradeable participant set by jurisdiction at the venue level. Mirror’s lesson is that a working mechanism does not cure a regulatory defect; the construction here accepts this and foregrounds it.

The SEC action against Terraform Labs (2023-2024) [26] was a composite failure: primarily an unregistered-securities-offering case covering LUNA, UST, MIR, and mAssets, compounded with Section 10(b) fraud charges about UST stability and Chai’s use of the Terra blockchain, and with market-manipulation allegations. The Moxie/Mass construction addresses the architectural failure mode through algebraic settlement gating: every trade must satisfy per-jurisdiction, verifiable compliance membership. Fraud and market-manipulation components are intentional-deception failures outside the mechanism described below. We state the narrower structural claim formally.

Definition 9.1 (Statutory regulatory threshold $\tau^{\text{reg}(I)}$). For an instrument class I (e.g., revenue-linked note, sukuk, binary event contract) and a jurisdiction ϕ_p , the *statutory regulatory threshold vector*

$$\tau^{\text{reg}(I)}(\phi_p) \in \prod_{d \in D(I)} \mathcal{G}_d$$

is the minimum per-domain compliance grade required by ϕ_p 's statute for a participant in ϕ_p to trade I in ϕ_p lawfully. Concretely, $\tau^{\text{reg}(I)}$ is a partial function from (instrument-class, jurisdiction) pairs to grade vectors; it is populated by the jurisdiction's regulatory authority (or its designated attestor) and updated when the statute changes. The canonical partial mapping for the six instrument classes of Section 3 and the four jurisdictions of Section 6 (US, EU, ADGM, Seychelles) is:

- **Binary event contract, US:** $\tau_{\text{Securities}}^{\text{reg}}(\text{US}) = \text{CFTC-DCM-registered (Commodity Exchange Act §5h designated contract market)}$; $\tau_{\text{KYC}}^{\text{reg}}(\text{US}) = \text{BSA CIP-compliant}$.
- **Revenue-linked note, US:** $\tau_{\text{Securities}}^{\text{reg}}(\text{US}) = \text{accredited-investor-verified under Reg D Rule 506(c) of the Securities Act of 1933, OR qualified under Reg A+ Tier 2}$.
- **Revenue-linked note, EU:** $\tau_{\text{Securities}}^{\text{reg}}(\text{EU}) = \text{MiFID II transferable-security-compliant; prospectus approved under Prospectus Regulation (EU) 2017/1129}$.
- **Revenue-linked note, ADGM:** $\tau_{\text{Securities}}^{\text{reg}}(\text{ADGM}) = \text{debenture-offer-authorized under the Financial Services and Markets Regulations 2015 (FSMR)}$.
- **Revenue-linked note, Seychelles:** $\tau_{\text{Securities}}^{\text{reg}}(\text{Seychelles}) = \text{Securities-Act-authorized under the Securities Act 2007 (as amended)}$.
- **Sukuk, ADGM:** $\tau_{\text{Sharia}}^{\text{reg}}(\text{ADGM}) = \text{FullyCertified (SSB fatwa current)}$; $\tau_{\text{Securities}}^{\text{reg}}(\text{ADGM}) = \text{FSMR Islamic-finance-product-authorized}$.
- Analogous entries for each (I, ϕ_p) pair in the exchange's authoritative threshold registry.

The threshold mapping has three distinguishing properties from the venue-internal tier thresholds $\tau_d^{\text{Warm}}, \tau_d^{\text{Hot}}$ (Definition 5.23): (i) $\tau^{\text{reg}(I)}$ is keyed by instrument class and jurisdiction, not by venue tier; (ii) $\tau^{\text{reg}(I)}$ is supplied by the sovereign (or its designated attestor) via an authenticated attestation, not computed by the kernel from observable venue state; (iii) updates to $\tau^{\text{reg}(I)}$ propagate via obligation-pack recompilation, not via tier recomputation. The venue tier thresholds $\tau_d^{\text{Warm}}, \tau_d^{\text{Hot}}$ are *upper bounds* on $\tau^{\text{reg}(I)}$: the venue may additionally require a higher grade than the statute mandates (a conservative posture), but the venue may not admit a trade whose participant's grade falls below $\tau^{\text{reg}(I)}$ in any required domain.

Theorem 9.2 (Settlement-admissibility under per-jurisdiction statutory and life-cycle gating). *Let \mathcal{E} be an instrument-bearing multi-harbour entity (Definition 4.2) in the independent-threshold regime, with harbour set $\Phi(\mathcal{E})$, composed compliance vector $\mathcal{C}(\mathcal{E}) = \bigwedge_{\phi \in \Phi(\mathcal{E})} \mathcal{C}(\phi)$, and corridor set $\mathcal{K}(\mathcal{E})$. Let I be an instrument issued by \mathcal{E} on the venue, activating domain set $D(I) \subseteq \{1, \dots, 23\}$ and life-cycle state $\ell \in \text{States}(I)$. Let p be a participant of jurisdiction ϕ_p submitting a trade request r tagged with market phase $\text{phase}(r) \in \{\text{primary}, \text{secondary}\}$ and per-unit holding period h_p . Let $\tau^{\text{reg}(I)}$ be the statutory threshold vector of Definition 6.1. Settlement of r admits if and only if all five conditions hold:*

- (S1) $\phi_p \in \Phi(\mathcal{E})$, or there exists a corridor $(\phi_i, \phi_p) \in \mathcal{K}(\mathcal{E})$ with $D(I) \subseteq R_{i,p}$.
- (S2) For every $d \in D(I)$, $c_d(\phi_p) \geq \tau_d^{\text{reg}(I)}(\phi_p, \text{phase}(r), h_p)$.
- (S3) The obligation-pack predicate $\bigwedge_k \text{sat}(o_k, t)$ for \mathcal{E} (Lemma 5.29) holds.
- (S4) The SAVM compiled bytecode for I emits a COMPLIANCE_CHECK success on $(\phi_p, D(I), \text{phase}(r), h_p)$ before RET.
- (S5) The trade-request type belongs to the outgoing-edge label set of state ℓ in the life-cycle automaton \mathcal{A}_I (Definition 2.16).

Failure of any of (S1)-(S5) is a hard trap in SAVM execution (Theorem 5.6); no trade settles. The theorem addresses the architectural unregistered-offering failure mode illustrated by Mirror Protocol (2020-2022). It does not address regulatory misclassification, fraud, or market-manipulation modes (Remark 9.3); those require statute-level enforcement.

Proof. Conditions (S1)-(S5) are each necessary at execution time. (S1) is enforced by the harbour-or-corridor membership check in the kernel's pre-trade admission path. (S2) is

enforced by the pointwise evaluation of the meet $\mathcal{C}(\mathcal{E}) \wedge \mathbf{e}_{\phi_p}$ against $\tau^{\text{reg}(I)}(\phi_p, \text{phase}, h_p)$ (Definition 6.1) at the compliance gate; the venue tier thresholds $\tau_d^{\text{Warm}}, \tau_d^{\text{Hot}}$ are enforced as upper bounds at the Warm/Hot tier. (S3) is enforced by the obligation evaluator (Lemma 5.29); an entity with an unsatisfied obligation is suspended. (S4) is enforced by the SAVM ISA: COMPLIANCE_CHECK failure traps execution (Definition 5.2), and by determinism (Theorem 5.6) every execution either produces the RET-value with all checks passed or halts with a defined trap. (S5) is enforced by the LIFECYCLE_TRANSITION opcode (Definition 2.17): a request whose type is absent from $\Delta_I(\ell)$ traps before any state mutation. The conjunction of the five conditions is what distinguishes the venue’s settlement path from Mirror’s: Mirror’s mAsset settlement was a collateralised-swap operation against UST with no (S1)-like membership gate, no (S2)-like statutory threshold gate, no (S3)-like obligation predicate, no (S4)-like COMPLIANCE_CHECK instruction, and no (S5)-like life-cycle predicate. \square

Remark 9.3 (Insulation scope: architectural unregistered-offering gate only). Theorem 9.2 addresses a narrow structural insulation property:

1. **Not classification-correctness.** A Sharia scholar, an SEC examiner, or an ADGM regulator may conclude that a specific instrument is misclassified, exactly as Mirror’s counsel misclassified mAssets. The insulation claim is narrower: a misclassification here is a parameter error (wrong $\tau^{\text{reg}(I)}$, wrong $D(I)$, wrong harbour) that is auditable trade-by-trade and correctable by updating the obligation pack and the threshold registry. Mirror had no such parameter surface to correct.
2. **Not fraud prevention.** *SEC v. Terraform Labs* (2023) included Section 10(b) and Rule 10b-5 fraud allegations about UST’s stability and about Chai’s use of the Terra blockchain. These are intentional-misrepresentation failures that no on-chain compliance gate prevents: an issuer’s fraudulent attestation passes the kernel’s cryptographic-authenticity check while remaining substantively false. The standard remedy is statute-level fraud enforcement by the sovereign.
3. **Not market-manipulation prevention.** The SEC also alleged market manipulation via sparse and opaque trading venues. Event-atomic settlement and per-trade compliance checks narrow some manipulation surface (front-running, wash trading via undisclosed addresses) without eliminating it. Market surveillance protocols, separate from the compliance-gate architecture, are required for manipulation deterrence.
4. **What is covered.** The Mirror-class failure mode the theorem addresses is architecturally-unenforced statutory gating: Mirror’s design admitted any participant with UST collateral to any mAsset trade, without any per-jurisdiction securities-registration check. The theorem proves the venue’s design is structurally incapable of admitting an unauthorised trade because every settlement path passes through (S1)-(S5), each of which can fail independently and each of which halts SAVM execution.

Uniswap v4 hooks (2024-). Uniswap v4 [27] introduced user-defined pool logic via “hooks”: pre-swap, post-swap, pre-liquidity, and post-liquidity callbacks that let a pool developer inject custom invariants, fees, and accounting. The CONVEXPOTENTIAL trait of Section 5 is, at the Rust-trait level, structurally similar to a Uniswap v4 hook: a polymorphic dispatch that lets each instrument class define its own invariant surface. Uniswap v4 does not have a plugin-level compliance lattice, a multi-harbor issuing entity, a gas model tied to event-atomic settlement, or an obligation-pack compiler; those are the contributions this paper makes over v4 hooks. Uniswap v4’s plugin architecture (gas metering, reentrancy guards, approved-hook registry) is more mature than the obligation-pack compiler described here; this is an area where Moxie will have to borrow rather than improvise.

Balancer v3 smart pools (2024). Balancer v3 [28] generalized Balancer’s weighted-geometric-mean invariants (configurable weights, dynamic weights, rate-providers, custom math via the “Vault” abstraction). The convertible-note potential (weighted geometric mean with moneyness-dependent exponent) and the option-like potential (moneyness-dependent weights) of Section 5.3 are, as pricing curves, Balancer-family curves: Balancer pools with reserve-dependent weights. What this paper adds beyond Balancer v3 is the formal quasi-concavity analysis under reserve-dependent weights (Propositions 5.16 and 5.19, including the explicit β_{\max} bounds) and the sukuk potential’s C^0 kink, which no Balancer curve expresses.

Tokenized-security standards: ERC-1400, ERC-3643, and their sub-standards. The programmable-security definition of Section 2 sits alongside two EVM-level standards that encode similar intent. ERC-1400 [65] and its sub-standards (ERC-1410, -1594, -1643, -1644) introduce typed partitions, controller operations, and document metadata for security tokens. ERC-3643 [30], the T-REX protocol, is the current canonical security-token standard and is the basis for Securitize, Tokeny, Archax, and most regulated European security-token platforms. ERC-3643’s primitives are a three-registry identity pattern (*ClaimTopicsRegistry*, *TrustedIssuersRegistry*, *IdentityRegistry*), a self-sovereign *OnchainID* token carrying compliance claims, a modular *ICompliance* interface with per-transfer hooks, and administrator-side *forced-transfer / recovery / freeze* primitives for corporate-control events. Our (obligation-pack compiler, SAVM, 23-domain tensor) trio is a formal lift of ERC-3643’s (*ICompliance*, *IdentityRegistry*, *ClaimTopicsRegistry*) with three additions: (i) typed total-recursive contract semantics (Definition 2.6; ERC-3643’s Solidity has reentrancy risk), (ii) category-theoretic composition (Theorem 2.11; ERC-3643 has no composition semantics), and (iii) a 23-domain meet-semilattice (Definition 5.24; ERC-3643’s *ICompliance* is an unstructured module array). What we borrow from ERC-3643 and do not improve: the forced-transfer / freeze / recovery primitives for administrator-side interventions, which appear here as the retroactive and cancellation remedies of Section 6.3. Peyton Jones, Eber, and Seward [63] (originally 2000) introduced a combinator-style DSL for financial derivatives in Haskell whose compositional semantics preserves pricing-relevant properties; Marlowe [64] specialized this lineage to a blockchain target with bounded-depth step semantics and Isabelle/HOL verification. Our Definition 2.6 differs from Peyton-Jones-Marlowe on totality and gas-boundedness (enforced by the SAVM compilation pipeline), and from ERC-3643 on explicit compliance-lattice parameterization with a proved composition theorem.

Commercial security-token platforms (Securitize, tZero, Tokeny, INX). Four deployed platforms are the nearest commercial analogues of the architecture described here. Securitize [31] implements a US Reg D plus ATS secondary-market stack on the proprietary DS Protocol. tZero [32] is an SEC-registered ATS that settles via Broadridge. Tokeny [33] is the reference ERC-3643 implementer for European regulated issuance. INX [34] achieved an F-1-registered digital-security IPO in 2021. These systems solved portions of issuance, cap-table management, transfer restrictions, and single-jurisdiction secondary trading. Cross-border composition remained outside their architecture. Their compliance logic is attached primarily to one transfer delegate, one ATS, or one domestic regulatory perimeter at a time. This paper’s contribution over those platforms is formal semantics for multi-jurisdictional composition, a corridor mechanism for cross-harbor recognition, and a temperature function that distinguishes partial from full cross-domain graduation.

Academic literature on tokenized securities. The tokenized-securities literature is largely empirical and legal-analytical. Zetsche, Arner, and Buckley [35, 36] provide

the canonical regulatory taxonomy of tokenization and the regulatory-arbitrage map between jurisdictions. Massad [37] analyzes the US regulatory architecture for tokenized instruments and identifies the division-of-labor question between CFTC and SEC that this paper’s Definition 9.1 encodes. Cong and He [38] develop the economic-incentive theory of blockchain-based smart contracts and connect it to financial-market microstructure. Volkov [39] gives a security-token regulatory taxonomy aligned with the ERC-3643 architecture. Harvey, Ramachandran, and Santoro [40] provide the economics of tokenized finance. Hart and Moore [41, 42, 43] develop the incomplete-contracts literature that motivates our treatment of the obligation pack as the *verifiable* subset of a contract, with the unverifiable subset remaining outside the pack. The common gap in this literature is the same gap visible in deployed platforms: tokenization changes the wrapper, but the binding state of a cross-border instrument is the composed compliance tensor. This paper is a formal-methods contribution to that composition problem.

Market microstructure theory (Kyle; Glosten-Milgrom; Hasbrouck; Amihud). The price-discovery and agency-cost claims in Sections 10 and 8 build on the canonical microstructure literature. Glosten and Milgrom [44] and Kyle [11] give the two standard informed-trader models of price discovery; the coupling function $G_\ell(E) = (E/(1-E))^\alpha$ of Section 8 admits a Glosten-Milgrom reading as an informed-trader-weighted adjustment and a Kyle reading as a reduction in the price-impact parameter λ when the prediction market’s informative component is observable. Hasbrouck [45, 46] formalises price discovery as the permanent-innovation share of each venue in the common fundamental; under this framework, the minimum-variance coupling weight $\alpha^* = \sigma_A^2 / (\sigma_A^2 + \sigma_P^2)$ of Proposition 8.2 is exactly the prediction market’s Hasbrouck information share (Proposition 10.6). Easley and O’Hara [47], Amihud [48], and Huang and Stoll [49] provide the complementary liquidity-decomposition machinery; the Huang-Stoll spread decomposition is the natural target for the coupled AMM-prediction-market system. Jensen and Meckling [50] and Shleifer and Vishny [51] supply the agency-cost foundations: the obligation pack (Section 5.5) internalises the *verifiable* component of the Jensen-Meckling monitoring cost by rendering observable-but-unverifiable decisions verifiable on-chain, reducing ex-post bonding and residual loss. The Holmström [52] moral-hazard component (unobservable managerial effort) remains out of scope.

Mechanism design for financial innovation (Peyton Young; Friedman-Kaufman-Heller). The programmable-security framework treats instrument design as a mechanism-design problem: the SPV’s obligation pack is a mechanism that gates counterparty eligibility, enforces proportional distribution, and chooses among settlement curves. This perspective traces to Peyton Young’s work on individual strategy and institutional structure [76, 77], which showed how institution-level rules (voting rules, taxation rules, bargaining rules) translate into equilibrium individual behavior. Friedman, Kaufman, and Heller (across [78, 79, 80]) developed the complementary reading: financial innovation is the production of new instruments that solve coordination problems (risk transfer, duration matching, liquidity provision) that older instruments solved inefficiently or not at all. Our instrument taxonomy (Section 3) sits inside this tradition: each instrument class is a mechanism for a coordination problem (binary event contracts for price discovery, revenue-linked notes for long-duration IP cash flow, sukuk for Sharia-compliant yield, experiential tokens for non-financial access rights). The Friedman-Heller reading of financial innovation as property-rights de-fragmentation [80] is particularly relevant to the IP-licensing use case: unified on-chain ownership with programmable licensing rules is a structural response to the anticommons that [80] identifies in fragmented IP portfolios.

DeFi as tokenized derivatives (Adams-Lestyán-Robinson; Daian et al). A growing body of academic work situates decentralized finance within the broader financial-engineering literature. Adams, Robinson, and Salem [74] (in a volume edited by Lestyán) provide a unified survey of decentralized-exchange design from constant-function market-makers through intent-based architectures; their classification of AMM potentials (constant-product, constant-sum, weighted, concentrated, stable-swap) is the nearest-neighbor prior work to the five instrument-specific potentials of Section 5.3. The Daian et al. line of work on decentralized-exchange security [20, 75] demonstrated that flash-loan-enabled atomic composition of DeFi primitives produces instrument behaviors (high-leverage directional bets, oracle-manipulation trades) that classical derivative taxonomies would classify as swaps or options with delivery-via-smart-contract. The programmable-security framework absorbs these lessons architecturally: event-atomic settlement (no mid-block condition-on-outcome; Section 7, $T_1 < T_2$ ordering) precludes the flash-loan atomicity pattern, and the compliance-gate on COMPLIANCE_CHECK (Theorem 9.2) precludes the oracle-manipulation-as-service pattern at the venue level.

Regulatory frameworks for tokenized securities (FINMA 2018; MAS PSA 2019; SEC No-Action letters). Three regulatory frameworks informed this paper’s threshold-registry design (Definition 9.1). FINMA’s 2018 Guidance on ICOs [71] established the canonical three-way classification of tokens by economic function (payment / utility / asset) with hybrid status contemplated; this is the model for the per-instrument activated-domain parameterization of Section 3. Singapore’s Payment Services Act [72] (2019, revised subsequently) introduced a licensing framework for digital-payment tokens that is explicitly activity-based rather than instrument-based; this is the model for the venue-internal obligation-pack structure (Section 5.5) and its independence from the statutory classification surface. The SEC’s 2019-2020 sequence of No-Action letters on tokenized platforms [73] (TurnKey Jet, Pocketful of Quarters, INX) established that a tokenized instrument’s regulatory classification turns on economic substance rather than wrapper form, and that a utility-token structure is respected only under narrow structural conditions (single-purpose consumption, no secondary-market expectation, no appreciation-linked value); this is the empirical basis for the “experiential token” classification in Table 2 and the caveat that substance can re-classify the token to a security under secondary-market-development circumstances.

Summary of novelty. What is genuinely new in this paper, net of the Lange/Peters-SoYe/Baron-Lange lineage and the Compound/Centrifuge/Sia/Mirror/Uniswap/Balancer prior work above, is:

1. the integration of a 23-domain multi-harbor compliance tensor, using the Applicable-fragment meet-semilattice and the F144 dichotomy for the full mixed-axis tensor, with contingent-claim-style instrument issuance pathways;
2. the sukuk potential’s C^0 -but-not- C^1 kinked invariant surface with exact $\alpha \in (0, 1)$ characterisation (Proposition 5.13), a profit-sharing surface shape absent from the DBAR catalogue and the Baron-Lange instrument family;
3. the meet-preserving (not full-lattice) homomorphism between the 23-domain compliance lattice and the three-tier temperature lattice (Definition 5.24);
4. the SAVM + obligation-pack compiler pipeline as a declarative-to-bytecode lowering target (Definitions 5.2-5.7) for obligations that span securities, Sharia, event-resolution, insurance, and material-change domains.

Everything else has prior art and is presented as a consolidation, not an invention: binary event contracts (direct descendants of Lange 1999 / DBAR), revenue-linked notes (Sia and Centrifuge), index products, options (Balancer weighted-geometric mean), experiential

tokens, the constant-product baseline, weighted-geometric-mean pools, the parimutuel-contingent-claim convex-auction framing (Lange-Economides 2005 / Peters-So-Ye 2006), and the interface-dispatch pattern over instrument-specific invariant surfaces (Uniswap v4 hooks).

10 The Information-Value Multiplier: Closed-Form Parameterization

Remark 8.3 stated the 5-20 \times economic multiplier as a calibrated estimate rather than a theorem, because the multiplier depended on unspecified volume-elasticity parameters. This section strengthens the result. We derive the multiplier in closed form as a function of two named primitives: the asset-side elasticity η_A (the volume response of spot traders to price accuracy) and the event-side elasticity η_E (the volume response of prediction-market and information-consuming constituencies to signal quality). We then prove upper and lower bounds on the multiplier under regularity conditions on these primitives. What remains empirically uncalibrated is the numerical value of η_A and η_E , not the functional form or the existence of the bound.

10.1 Primitives and the economic ratio

Definition 10.1 (Elasticity primitives). Let V_A denote the spot trading volume at an AMM and V_E the volume of information-consuming participants (prediction-market traders and the four non-LP constituencies enumerated in Remark 8.3). Let ΔMSE denote the MSE reduction under coupling (log-price squared units; Proposition 8.2). The *elasticities* are

$$\eta_A := \frac{\partial \ln V_A}{\partial (\Delta\text{MSE}/\sigma_A^2)}, \quad \eta_E := \frac{\partial \ln V_E}{\partial (\Delta\text{MSE}/\sigma_A^2)}$$

both evaluated at the uncoupled baseline $\Delta\text{MSE} = 0$. These are dimensionless quantities measuring the logarithmic sensitivity of each side's volume to the fractional MSE reduction.

Definition 10.2 (Dollar-value conversion). Let f denote the per-trade fee and λ_A, λ_E the per-unit-volume profit margins captured by LPs and information-consumers respectively. The *dollar-value multiplier* is

$$M(\eta_A, \eta_E) := \frac{\text{total value to five constituencies}}{\text{LVR savings to LPs}} = \frac{\lambda_A V_A (\Delta\text{MSE}) \eta_A + \lambda_E V_E (\Delta\text{MSE}) \eta_E + \Delta\text{MSE}}{\Delta\text{LVR}}.$$

Proposition 10.3 (Closed-form multiplier). *Under the minimum-variance coupling $\alpha^* = \sigma_A^2 / (\sigma_A^2 + \sigma_p^2)$ of Proposition 8.2, and under the standardization $\lambda_A V_A = \lambda_E V_E = 1$ (equal baseline activity scale on both sides; the non-equal case rescales the formula multiplicatively), the multiplier admits the closed form*

$$M(\eta_A, \eta_E) = 8(1 + \eta_A + \eta_E), \quad (47)$$

where the leading factor 8 is the dimensionless mechanical ratio $\Delta\text{MSE}/\Delta\text{LVR} = 8$ proved in Proposition 8.2.

Proof. By Definition 10.1, the first-order Taylor expansion of V_A and V_E in $\Delta\text{MSE}/\sigma_A^2$ around the uncoupled baseline gives $V_A(\Delta\text{MSE}) \approx V_A(0)(1 + \eta_A \Delta\text{MSE}/\sigma_A^2)$ and similarly for V_E . The additional dollar value from volume response is $\lambda_A V_A(0) \eta_A \Delta\text{MSE}/\sigma_A^2 + \lambda_E V_E(0) \eta_E \Delta\text{MSE}/\sigma_A^2$. Adding the direct pricing-accuracy benefit ΔMSE (which accrues regardless of volume response) and dividing by $\Delta\text{LVR} = \Delta\text{MSE}/8$ gives (47) under the normalization $\lambda_A V_A(0) = \lambda_E V_E(0) = 1$ and $\alpha^* \sigma_A^2 = \Delta\text{MSE}$. The non-normalized version carries a multiplicative constant $(\lambda_A V_A + \lambda_E V_E)/\Delta\text{LVR}$ absorbed into the elasticity-weighted terms. \square

Theorem 10.4 (Bounds on the multiplier under regularity). *Suppose the elasticities (η_A, η_E) satisfy the regularity conditions*

(R1) **Non-negativity:** $\eta_A, \eta_E \geq 0$. *Volume cannot decrease in response to improved pricing accuracy; a negative elasticity would imply that traders are averse to better prices, which contradicts revealed preference in any competitive market.*

(R2) **Rationality bound:** $\eta_A \leq \bar{\eta}_A$ and $\eta_E \leq \bar{\eta}_E$, *where $\bar{\eta}_A, \bar{\eta}_E$ are upper bounds derived from the law of demand applied to the price-accuracy-demand curve.*

Then the multiplier M of Proposition 10.3 satisfies

$$8 \leq M(\eta_A, \eta_E) \leq 8(1 + \bar{\eta}_A + \bar{\eta}_E). \quad (48)$$

Under calibrations reported in the experimental-economics literature on prediction-market liquidity response to signal quality [4, 14], plausible values are $\bar{\eta}_A \approx 0.3-0.8$ and $\bar{\eta}_E \approx 1.0-1.5$ (prediction-market volume responds more elastically to signal quality than spot volume does to pricing accuracy because information-seeking participants are volume-elastic by construction), yielding

$$M \in [8, 8 \cdot (1 + 0.8 + 1.5)] = [8, 26.4],$$

which brackets the originally-claimed 5-20 \times band. The lower bound $M \geq 8$ is unconditional (it is the mechanical ratio).

Proof. The lower bound is the direct-pricing-accuracy term alone, attained when $\eta_A = \eta_E = 0$ (no volume response). The upper bound substitutes the regularity upper bounds into (47). \square

Remark 10.5 (What is proved and what remains empirical). Theorem 10.4 proves the following:

1. The functional form of the multiplier is $M = 8(1 + \eta_A + \eta_E)$, a closed-form identity, not a calibration artefact.
2. The lower bound $M \geq 8$ holds unconditionally from non-negativity of elasticities.
3. The upper bound depends on empirically calibrated values of $\bar{\eta}_A, \bar{\eta}_E$.

What remains empirical is the *numerical value* of η_A, η_E , not the existence of the multiplier or the shape of its dependence on volume response. The original 5-20 \times claim is a corollary of Theorem 10.4 under literature-standard elasticity calibrations, with explicit sensitivity to the elasticity inputs.

10.2 Price discovery under coupling

The closed-form multiplier of Proposition 10.3 situates the coupling in the Glosten-Milgrom [44] and Kyle [11] price-discovery frameworks. Hasbrouck [45] defines price discovery as the share of permanent innovation in the fundamental value series attributable to each trading venue. We state the corresponding claim for the coupled AMM-prediction-market system.

Proposition 10.6 (Hasbrouck information share of the prediction market). *Let P_t^A be the AMM spot price and $G_\ell(E_t) \cdot P_t^A$ the prediction-market-adjusted spot price. Treated as independent unbiased estimators of the common fundamental with variances $\sigma_A^2 + \sigma_P^2$ and σ_E^2 respectively (Proposition 8.2), the Hasbrouck information share of the prediction market is*

$$\text{IS}_E = \alpha^* = \frac{\sigma_A^2}{\sigma_A^2 + \sigma_P^2},$$

exactly the minimum-variance coupling weight.

Proof sketch. By Hasbrouck’s variance-decomposition of the common stochastic trend between two venues, the information share equals the variance-weighted contribution each venue makes to the permanent innovation. Under the independence assumption of Proposition 8.2 and the minimum-variance weighting α^* , the prediction-market share is exactly α^* . \square

The economic multiplier $M = 8(1 + \eta_A + \eta_E)$ of Theorem 10.4 is therefore the elasticity-weighted price-discovery contribution of the prediction market to the spot venue, scaled by the LVR baseline. In the Kyle framework, the coupling reduces the spot venue’s price-impact parameter λ by the fraction IS_E when the informative component of the prediction-market flow is observable on the spot book.

Open Problem 4 (Price discovery under adversarial prediction-market manipulation). Characterize the coupled system’s price discovery when the prediction market is subject to wash-trading or attention-gaming of the sort bounded in Proposition 12.1. Specifically: under what parameter regime does the Hasbrouck information share IS_E remain a faithful estimate of the fundamental’s permanent component, and when does manipulation degrade it below the uncoupled baseline?

11 The Smart Asset Virtual Machine: Specification

This section specifies the “Smart Asset Virtual Machine” (SAVM) rigorously: instruction set, calling convention, memory model, and halting guarantee. The result is a small, deliberately-restricted virtual machine whose design is driven by the compliance-gate evaluation requirement of programmable obligations, not by general-purpose computation.

11.1 SAVM instruction set architecture (ISA)

Definition 11.1 (SAVM ISA). The Smart Asset Virtual Machine (SAVM) is a stack-based virtual machine with the following instruction classes:

Arithmetic (FixedPoint 1e-18, signed 128-bit mantissa). ADD, SUB, MUL, DIV, NEG, ABS, MIN, MAX, FLOOR, CEIL. Overflow traps are mandatory (no wrap-around); division by zero traps.

Control flow. JMP *addr* (unconditional), JZ *addr* (jump if top of stack is zero), JNZ *addr*, CALL *addr* (push return address, bounded depth), RET, TRAP *code* (abort with explicit error code).

Compliance-gate. GATE {*domain*} (evaluates the compliance predicate for a given domain from the 23-domain lattice; pushes 1 for pass, 0 for fail; traps if the domain is not registered). GATE_MEET {*D1*, *D2*, . . . , *Dk*} evaluates the pointwise meet of multiple domains (returns 1 iff all pass).

State I/O (read-only for obligation evaluation, write-gated elsewhere). LOAD_RESERVE token (reads AMM reserve; read-only in obligation path), LOAD_PRICE token (reads current marginal price), LOAD_CLOCK (reads block height and timestamp; deterministic within a block), LOAD_PARAM *key* (reads a statically declared obligation parameter).

Stack. PUSH *imm*, POP, DUP *n*, SWAP *i j*, LOAD_LOCAL *n*, STORE_LOCAL *n* (reads/writes function-local fixed-point slot; stack per function frame has a statically bounded size declared by the obligation language).

Definition 11.2 (SAVM calling convention). Each SAVM program is a tuple (*code*, *entry_point*, *n_params*, *stack_size*). On entry, parameters are pushed left-to-right; the program reads them via LOAD_PARAM and leaves the return value on the stack. Functions share no state across invocations (no heap;

obligation-pack compiler statically verifies absence of persistent mutable state). Reentrancy is *statically disallowed*: the compiler rejects programs where a CALL could transitively reach its caller (see Lemma 11.5 below).

Definition 11.3 (SAVM memory model). The SAVM memory is partitioned into: (i) a *stack* of bounded depth D , where each stack frame has a statically-bounded local slot count S_i (determined per-function at compile time), giving total stack memory $\sum_{i \in \text{call chain}} S_i \leq D \cdot S_{\max}$; (ii) *read-only obligation parameters*, immutable for the lifetime of the program; (iii) *external state accessors*, which are the LOAD_* instructions and are side-effect-free (they do not mutate external state in the obligation-evaluation path). There is no heap and no general-purpose writable memory beyond the stack.

11.2 Structural guarantees

Lemma 11.4 (Recursion-depth bound). *Every SAVM program admits a compile-time upper bound $D^* \in \mathbb{N}$ on call-stack depth, and every execution satisfies $\text{depth}(t) \leq D^*$ for all t during execution.*

Proof. The obligation-pack compiler rejects any program whose call graph contains a cycle (Lemma 11.5). The call graph of an acyclic program is a DAG, whose longest path length D^* is computed in polynomial time by topological traversal. The SAVM interpreter maintains a call-depth counter and traps when the depth exceeds D^* (which by the static analysis can never occur in a well-formed program). \square

Lemma 11.5 (Static reentrancy exclusion). *The obligation-pack compiler rejects any SAVM program whose call graph contains a cycle.*

Proof. The compiler constructs the call graph $G = (V, E)$ where V is the set of function labels and $(u, v) \in E$ iff the body of u contains a CALL to v . Cycle detection on a DAG of $|V|$ nodes runs in $O(|V| + |E|)$ time via depth-first search. Programs with cycles are rejected with a compile-time error before bytecode emission. \square

Theorem 11.6 (Termination). *Every well-formed SAVM program terminates in bounded time on every input.*

Proof. By Lemmas 11.4 and 11.5, the call graph is a DAG with maximum depth D^* . Each function body is a straight-line sequence of bounded length L^* (by construction: no looping constructs other than CALL are provided in the ISA of Definition 5.2; conditional branches JZ/JNZ are backed by the compiler’s requirement that target addresses be *forward* in the function body, preventing in-function loops). Total instruction count is bounded by the static program bound $L^* \cdot D^*$. Termination follows. \square

Remark 11.7 (Why no in-function loops). The SAVM forbids unbounded loops by design. The obligations encoded in obligation packs (Sharia asset-backing ratios, accreditation checks, compliance-gate meets, revenue-share calculations over a bounded number of cash-flow periods) express as bounded straight-line code with bounded recursion. The restriction yields a strong termination guarantee without loss of expressiveness for the target workload. Programs requiring unbounded iteration lie outside the SAVM and decompose into multi-block obligations coordinated by the kernel’s external state machine.

Theorem 11.8 (Obligation-pack compilation correctness). *Let \mathcal{O} be a declarative obligation pack (as defined in Section 3), specifying a predicate $\pi_{\mathcal{O}} : \text{State} \rightarrow \{0, 1\}$. Let $C(\mathcal{O})$ be the SAVM bytecode emitted by the obligation-pack compiler. Then for every state s ,*

$$\text{eval}_{\text{SAVM}}(C(\mathcal{O}), s) = \pi_{\mathcal{O}}(s).$$

That is, SAVM evaluation of the compiled bytecode returns the declarative predicate’s value.

Proof sketch. The obligation-pack schema is a first-order language over FixedPoint arithmetic and compliance-domain predicates. Compilation proceeds syntax-directed: each schema construct (arithmetic, boolean combinator, compliance-gate access) maps to a fixed SAVM instruction sequence whose evaluation semantics matches the schema’s denotational semantics. Correctness is proved by structural induction on the schema tree: the base cases (literals, parameter loads, compliance-gate accesses) match by the ISA specification of Definition 5.2, and the inductive cases (arithmetic combinators, boolean combinators) match by compositionality of SAVM evaluation. See [29] for the standard technique. \square

11.3 Consequences of the SAVM specification

With Definitions 5.2, 11.2, 5.3 and Theorems 11.6, 11.8, the specification is complete. The obligation-pack compiler emits SAVM bytecode; the SAVM interpreter evaluates it; the compilation-correctness theorem ensures the SAVM evaluates the declarative obligation.

Remark 11.9 (Relation to the Ethereum EVM and Solana eBPF). The SAVM is restrictive compared with the EVM and Solana’s eBPF target: no unbounded loops (termination is structural, not gas-metered), no heap or persistent mutable state in the obligation-evaluation path (reentrancy vulnerabilities are statically impossible rather than mitigated), and fixed-point-only arithmetic (no floating point, no implicit integer widening). The trade-off is that the SAVM compiles only obligation packs from the declarative schema; general-purpose smart contracts requiring full Turing completeness run on the kernel’s other execution paths. The SAVM’s mandate is the correctness and auditability of programmable-securities compliance predicates, not arbitrary DApp execution.

12 Open Problems

Open Problem 5 (Sharia verification complexity and decidability). Characterise the set of Sharia structural constraints whose verification is decidable in polynomial time under fixed commitment-state size, and separate them from the constraints requiring substantive judgement under AAOIFI *fiqh al-muamalat* standards. The kernel verifies structural SH-01..SH-04 constraints (no riba, no gharar, no maysir, asset backing) computationally; SH-05 (SSB certification) remains a substantive determination requiring human scholars. Whether the kernel’s structural proofs reduce SSB review time from weeks to days, and whether SSBs accept kernel-produced proofs as supplementary evidence rather than substitutes for review, are empirical questions requiring pilot issuances with SSB cooperation.

Open Problem 6 (Cross-jurisdictional corporate-action composability). When two jurisdictions declare overlapping corporate actions on the same multi-harboured instrument (for example, a US merger record date coinciding with an EU tender-offer election window), the corridor mechanism of Definition 4.2 does not provide a priority rule. Characterise the partial order on $CA \times CA$ pairs that minimises investor harm while satisfying each sovereign’s statutory mandate.

Open Problem 7 (Dynamic obligation-pack upgrades for long-dated instruments). For perpetual preferred equity and multi-decade sukuk, the obligation pack in force at issuance may diverge from the pack required at a later coupon date (statutes change; component constraints evolve). Formalise the upgrade protocol that propagates pack identifiers to in-flight instruments, preserves audit-trail immutability, and bounds the re-evaluation cost.

Open Problem 8 (Gas-metered compliance-check MEV surface). The SAVM cost of COMPLIANCE_CHECK is 50 gas per domain (Definition 5.5). Does the per-trade compliance

load concentrate gas consumption in the compliance tier sufficiently to create MEV surfaces (block builders ordering trades by compliance-check cost)? Characterise conditions under which compliance-driven gas variance admits a fair-ordering guarantee.

Open Problem 9 (Cross-instrument correlation under multi-harbor). When a single entity holds multiple IP assets across multiple jurisdictions, the instruments linked to those assets are correlated through both the entity structure and the corridor network. Characterizing the correlation structure induced by multi-harbor topology, and incorporating it into the clearinghouse’s portfolio margin formula, requires extending the Ising interaction model from token pairs to instrument-jurisdiction pairs. The state space grows from $O(M^2)$ to $O((M \cdot |\Phi|)^2)$, and the computational tractability of the mean-field approximation in this expanded space is unclear.

Open Problem 10 (Optimal corridor sequencing for instrument coverage). Given a target set of instrument types and a set of candidate jurisdictions, the optimal corridor deployment sequence maximizes the tradeable instrument-participant surface at each step. This is a combinatorial optimization over directed graphs with heterogeneous edge costs (each corridor has a different negotiation cost and timeline). The problem is at least as hard as the minimum spanning arborescence problem, and the objective function (tradeable surface) depends on the interaction between corridor topology and compliance lattice in a non-trivial way.

Open Problem 11 (Gaming of attention intensity). If attention intensity $I(e, t)$ influences coupling strength, an adversary can inflate I by wash trading the prediction market (high volume, many addresses, artificial price volatility). The cost of wash trading is bounded below by the prediction market’s bid-ask spread and the LMSR’s loss function. Formalizing the minimum cost of inflating I by a factor λ as a function of the market parameters, and proving that this cost exceeds the profit from the manipulated coupling, is open. The defense requires that the LMSR backstop parameter b and the coupling attenuation α be jointly calibrated so that manipulation is unprofitable.

Proposition 12.1 (Attention-intensity gaming, two-sided bound). *Let $I(e, t)$ denote the attention intensity, assumed to be an increasing concave function of observed prediction-market trading volume V_e : specifically $I = V_e^\gamma$ for $\gamma \in (0, 1)$ (calibrated empirically; $\gamma \approx 0.5$). Let s be the bid-ask spread and b the LMSR liquidity. To inflate I by a factor $\lambda > 1$ the adversary must inflate V_e by $\lambda^{1/\gamma}$. The adversary’s cost to inflate I by factor λ is bounded below by*

$$C_{\text{wash}}(\lambda) \geq s \cdot V_e^{\text{base}} \cdot (\lambda^{1/\gamma} - 1) + b \ln 2, \quad (49)$$

and the adversary’s extracted profit from the inflated coupling is bounded above by

$$\Pi_{\text{wash}}(\lambda) \leq \alpha \cdot (\lambda - 1) \cdot I^{\text{base}} \cdot R_M, \quad (50)$$

where α is the coupling attenuation. Gaming is UNPROFITABLE whenever

$$b \cdot \ln 2 + s \cdot V_e^{\text{base}} \cdot (\lambda^{1/\gamma} - 1) > \alpha \cdot (\lambda - 1) \cdot I^{\text{base}} \cdot R_M. \quad (51)$$

For $\gamma = 0.5$ and $\lambda \geq 2$, the left-hand side grows as λ^2 while the right-hand side grows linearly, so there exists a critical $\lambda^*(b, s, \alpha, R_M)$ above which gaming is NEVER profitable.

Proof sketch. Lower bound on cost. Each wash-trade round-trip pays the bid-ask spread s on notional; inflating volume to $\lambda^{1/\gamma} V_e^{\text{base}}$ requires $V_e^{\text{base}} (\lambda^{1/\gamma} - 1)$ additional notional. The LMSR’s bounded-loss property [66] adds a separate $b \ln 2$ cap on any sub-sequence of prediction-market trades.

Upper bound on profit. The spot extraction is linear in I at rate $\alpha \cdot R_M$; inflating I from I^{base} to λI^{base} yields at most $\alpha (\lambda - 1) I^{\text{base}} R_M$ spot profit before the LMSR self-calibrator

updates $\kappa \rightarrow \kappa'$ (Proposition 12.1 can be tightened when combined with EMA-calibrator response; in this form it is a worst-case no-calibrator bound).

Critical λ^ .* For $\gamma = 0.5$, the cost bound reduces to $C_{\text{wash}}(\lambda) \geq sV_e^{\text{base}}(\lambda^2 - 1) + b \ln 2$, while the profit bound is $\Pi_{\text{wash}}(\lambda) \leq \alpha(\lambda - 1)I^{\text{base}}R_M$. Equating the leading terms (dropping the $b \ln 2$ constant): $sV_e^{\text{base}}(\lambda^2 - 1) = \alpha(\lambda - 1)I^{\text{base}}R_M$, which factors as $sV_e^{\text{base}}(\lambda + 1)(\lambda - 1) = \alpha(\lambda - 1)I^{\text{base}}R_M$, giving $\lambda^* = -1 + \alpha I^{\text{base}}R_M / (sV_e^{\text{base}})$ for $\lambda > 1$. For the canonical calibration ($\alpha = 0.3$, $I^{\text{base}} = 10^4$, $R_M = 2.5 \cdot 10^5$, $s = 0.003$, $V_e^{\text{base}} = 10^4$) this yields $\lambda^* = -1 + 0.3 \cdot 10^4 \cdot 2.5 \cdot 10^5 / (0.003 \cdot 10^4) = -1 + 7.5 \cdot 10^8 / 30 = -1 + 2.5 \cdot 10^7$, so $\lambda^* \approx 2.5 \cdot 10^7$. The bound at the stated calibration is *vacuously strong*: any inflation factor $\lambda < 2.5 \cdot 10^7$ is unprofitable, which covers every realistic adversary by an astronomical margin. In practice the more informative statistic is the minimum calibration under which gaming becomes even marginally profitable: reducing $\alpha I^{\text{base}}R_M$ by four orders of magnitude (e.g., to a thin-attention-event regime with $I^{\text{base}} = 10^0$ and $R_M = 2.5 \cdot 10^1$) yields $\lambda^* \approx 1 + 7.5/30 = 1.25$, a regime where the adversary can profitably inflate attention by 25% before the bound bites. Calibrating $(b, s, \alpha, R_M, V_e^{\text{base}})$ so that λ^* sits just above practically-observed wash-trade inflation factors (typically 1.1-2) is the relevant design problem; the bound's *vacuity* at the default calibration is itself evidence that the default calibration over-parameterizes the protection. \square

Remark 12.2 (Tightness). The bounds (49) and (50) match up to the pre-calibration constants. The EMA coupling calibrator of the companion coupling paper [5] (persistent-manipulation upper bound) can only strengthen the upper bound (50) by making α adaptive, so the “gaming is unprofitable” corollary is robust.

A Instrument-Class Prerequisite Ordering

This appendix states the *logical prerequisite ordering* on the instrument classes of Section 3, induced by their compliance-surface requirements (Section 4) and their temperature-tier prerequisites (Definition 4.10).

A.1 Prerequisite ordering on instrument classes

Each instrument class of Section 3 requires a specific subset of the 23 compliance domains to be active on the issuing entity's composed compliance surface, and a specific minimum temperature tier. The prerequisite relation \preceq is defined by: $I_1 \preceq I_2$ iff every domain and tier requirement of I_1 is required by I_2 . This partial order induces three equivalence classes of instrument-domain activation, which we label Class I, Class II, Class III:

- **Class I (minimal compliance surface).** Instrument classes whose compliance requirement is AML/KYC plus at most one venue-internal obligation. Binary event contracts (Definition 3.1) and experiential tokens (Definition 3.9) fall in this class; both trade at the Cold tier and do not require the entity to be securities-registered in any jurisdiction.
- **Class II (multi-harbor compliance composition).** Instrument classes whose compliance requirement activates securities, derivatives, Sharia, or collective-investment-scheme domains and therefore requires the entity to be harbored in a jurisdiction that statutorily recognizes that domain. Revenue-linked notes, contingent licensing options, and sukuk (Definitions 3.3, 3.6, 3.7) fall in this class; they trade at the Warm tier or above and require meet-semilattice composition across at least two harbors with an active corridor.
- **Class III (full-surface graduation).** Instrument classes whose compliance surface must resolve across the full 23-domain tensor before open secondary clearing is admissible. IP index products (Definition 3.5) fall here; they trade at the Hot tier and

require the entity's corridor graph to span enough jurisdictions that the index fund vehicle's composed surface admits each underlying asset's harbor.

Proposition A.1 (Class ordering is transitive). *The relation “ I_1 is in an earlier class than I_2 ” is transitive: if I_1 requires a strict subset of I_2 's compliance domains and a \leq -smaller temperature tier, and likewise I_2 versus I_3 , then I_1 's requirements are a strict subset of I_3 's. Transitivity is immediate from the lattice structure of the compliance surface and the total order on temperature tiers.*

Remark A.2 (Logical structure, not commercial sequence). Proposition A.1 states the *logical prerequisite structure* of the instrument taxonomy. Which class any specific deployment reaches, through which partners, at what capital cost, and under which jurisdictional counterparties is outside the scope of this paper.

A.2 Why the ordering matters theoretically

The Class I \rightarrow II \rightarrow III progression is the order in which an entity's composed compliance surface must *grow* (pointwise raise its grades and possibly acquire additional harbors) to admit successive instrument classes. Because the lattice-homomorphism $\Phi : \mathcal{L}_{23} \rightarrow \mathcal{T}$ (Definition 5.24) is meet-preserving and fails to preserve joins, growth of the compliance surface is the only mechanism by which an entity can move to a higher temperature tier on the composed surface: adding a harbor with weaker compliance in a given domain pulls the meet down, never up. This asymmetry is what makes the class ordering a theoretical invariant of the architecture, rather than an operational preference. The corridor mechanism (Definition 4.2; corridor algebra formalized in [8]) is the construction that allows the meet over a harbor set to stay high in domains a sovereign does not recognize, at the cost of restricting the tradeable participant set to those who satisfy the carried-forward surface.

References

- [1] Lima Consulting Group. Global licensing industry survey. *Licensing International*, 2024.
- [2] International Islamic Financial Market. IIFM sukuk report 2024: A comprehensive study of the global sukuk market. *IIFM Annual Report*, 9th edition, 2024.
- [3] Newzoo. Global games market report 2025. *Newzoo Market Intelligence*, 2025.
- [4] Polymarket. 2024 annual volume report. <https://polymarket.com>, 2024.
- [5] R. Lorgat. Event-atomic settlement and probability-weighted market making. *Moxie Technical Paper*, 2026.
- [6] R. Lorgat. Liquidity aggregation and temperature-graduated exchange design. *Moxie Technical Paper*, 2026.
- [7] R. Lorgat. Discovering causal structure in event-linked markets. *Moxie Technical Paper*, 2026.
- [8] R. Lorgat. The algebra of institutional compliance. *Momentum Research*, 2026.
- [9] R. Lorgat. Mass sovereign-kernel design. *Momentum Research*, 2026.
- [10] C. W. J. Granger. Investigating causal relations by econometric models and cross-spectral methods. *Econometrica*, 37(3):424-438, 1969.

- [11] A. S. Kyle. Continuous auctions and insider trading. *Econometrica*, 53(6):1315-1335, 1985.
- [12] R. A. Fisher. On the “probable error” of a coefficient of correlation deduced from a small sample. *Metron*, 1:3-32, 1921.
- [13] J. Milionis, C. Moallemi, T. Roughgarden, and A. L. Zhang. Automated market making and loss-versus-rebalancing. Working paper, 2022.
- [14] R. Hanson. Logarithmic market scoring rules for modular combinatorial information aggregation. *Journal of Prediction Markets*, 1(1):3-15, 2007.
- [15] R. Rana, V. Nadkarni, N. Moshrefi, and P. Viswanath. ParlayMarket: Automated market making for parlay-style joint contracts. arXiv:2603.22596, 2026.
- [16] SEC v. W. J. Howey Co. 328 U.S. 293, 1946.
- [17] G. Angeris, H.-T. Kao, R. Chiang, C. Noyes, and T. Chitra. An analysis of Uniswap markets. *Cryptoeconomic Systems*, 1(1), 2021.
- [18] A. Park. The conceptual flaws of constant product automated market making. Working paper, University of Toronto, 2021.
- [19] E. Budish, P. Cramton, and J. Shim. The high-frequency trading arms race: Frequent batch auctions as a market design response. *Quarterly Journal of Economics*, 130(4):1547-1621, 2015.
- [20] P. Daian, S. Goldfeder, T. Kell, Y. Li, X. Zhao, I. Bentov, L. Breidenbach, and A. Juels. Flash boys 2.0: Frontrunning in decentralized exchanges, miner extractable value, and consensus instability. In *IEEE Symposium on Security and Privacy*, pages 910-927, 2020.
- [21] H. Adams, N. Zinsmeister, M. Salem, R. Keefer, and D. Robinson. Uniswap v3 core. Technical report, Uniswap Labs, 2021.
- [22] Compound Labs. Compound Chain white paper. Technical report, Compound Labs, 2020.
- [23] L. Frohwein, M. Knecht, and J. Zhong. Centrifuge: Tokenized real-world asset financing. Centrifuge white paper, 2019 (revised subsequently).
- [24] D. Vorick and L. Champine. Sia: Simple decentralized storage. Nebulous Inc., 2014; Siafund specification, 2015.
- [25] Terraform Labs. Mirror Protocol: Synthetic assets on Terra. Technical documentation, 2020; v2 updates 2021-2022.
- [26] Securities and Exchange Commission. *SEC v. Terraform Labs Pte. Ltd. and Do Hyeong Kwon*. Complaint filed 16 February 2023; Case No. 1:23-cv-01346 (S.D.N.Y.); summary judgment on unregistered-securities claims 28 December 2023; jury verdict on fraud claims April 2024.
- [27] H. Adams, M. Salem, N. Zinsmeister, S. Reynolds, A. Oskooei, M. Toda, L. Roberts, E. Ulkhanov, A. Karys, S. Son, J. Robinson, A. Sabnis, R. Keefer, and D. Robinson. Uniswap v4 core. Draft technical report, Uniswap Labs, 2024.
- [28] Balancer Labs. Balancer v3: A generalized AMM with the Vault abstraction and smart pools. Technical documentation, 2024.

- [29] B. C. Pierce. *Types and Programming Languages*. MIT Press, 2002.
- [30] Tokeny and the ERC-3643 Association. T-REX: A standard for regulated security tokens (ERC-3643). EIP and reference implementation, 2021-2023.
- [31] Securitize. DS Protocol: A compliance platform for digital securities. Technical documentation, 2024.
- [32] tZERO Group, Inc. ATS-registered secondary market for digital securities. Operator disclosures, 2024.
- [33] Tokeny Solutions. Institutional-grade tokenization platform. Technical documentation, 2024.
- [34] INX Limited. F-1 registration statement for INX Token. U.S. Securities and Exchange Commission filing, 2020; listing August 2021.
- [35] D. A. Zetsche, D. W. Arner, and R. P. Buckley. Decentralized finance. *Journal of Financial Regulation*, 6(2):172-203, 2020.
- [36] D. A. Zetsche, R. P. Buckley, D. W. Arner, and J. N. Barberis. From FinTech to TechFin: The regulatory challenges of data-driven finance. *New York University Journal of Law and Business*, 14:393-446, 2018.
- [37] T. G. Massad. The regulatory architecture for tokenized securities. Working paper, Brookings Institution, 2022.
- [38] L. W. Cong and Z. He. Blockchain disruption and smart contracts. *Review of Financial Studies*, 32(5):1754-1797, 2019.
- [39] A. Volkov. A regulatory taxonomy for tokenized securities. Working paper, 2023.
- [40] C. R. Harvey, A. Ramachandran, and J. Santoro. *DeFi and the Future of Finance*. Wiley, 2021.
- [41] O. Hart and J. Moore. Incomplete contracts and renegotiation. *Econometrica*, 56(4):755-785, 1988.
- [42] O. Hart. *Firms, Contracts, and Financial Structure*. Oxford University Press, 1995.
- [43] P. Bolton and M. Dewatripont. *Contract Theory*. MIT Press, 2005.
- [44] L. R. Glosten and P. R. Milgrom. Bid, ask and transaction prices in a specialist market with heterogeneously informed traders. *Journal of Financial Economics*, 14(1):71-100, 1985.
- [45] J. Hasbrouck. One security, many markets: Determining the contributions to price discovery. *Journal of Finance*, 50(4):1175-1199, 1995.
- [46] J. Hasbrouck. *Empirical Market Microstructure*. Oxford University Press, 2007.
- [47] D. Easley and M. O'Hara. Price, trade size, and information in securities markets. *Journal of Financial Economics*, 19(1):69-90, 1987.
- [48] Y. Amihud. Illiquidity and stock returns: Cross-section and time-series effects. *Journal of Financial Markets*, 5(1):31-56, 2002.
- [49] R. D. Huang and H. R. Stoll. The components of the bid-ask spread: A general approach. *Review of Financial Studies*, 10(4):995-1034, 1997.

- [50] M. C. Jensen and W. H. Meckling. Theory of the firm: Managerial behavior, agency costs and ownership structure. *Journal of Financial Economics*, 3(4):305-360, 1976.
- [51] A. Shleifer and R. W. Vishny. A survey of corporate governance. *Journal of Finance*, 52(2):737-783, 1997.
- [52] B. Holmström. Moral hazard and observability. *Bell Journal of Economics*, 10(1):74-91, 1979.
- [53] R. C. Merton. Theory of rational option pricing. *Bell Journal of Economics and Management Science*, 4(1):141-183, 1973.
- [54] R. C. Merton. Option pricing when underlying stock returns are discontinuous. *Journal of Financial Economics*, 3(1-2):125-144, 1976.
- [55] S. G. Kou. A jump-diffusion model for option pricing. *Management Science*, 48(8):1086-1101, 2002.
- [56] A. M. Turing. On computable numbers, with an application to the Entscheidungsproblem. *Proceedings of the London Mathematical Society*, s2-42(1):230-265, 1936.
- [57] C. A. R. Hoare. Proofs of correctness of data representations. *Acta Informatica*, 1(4):271-281, 1972.
- [58] J. He, C. A. R. Hoare, and J. W. Sanders. Data refinement refined. In *ESOP '86*, LNCS 213, pages 187-196. Springer, 1986.
- [59] N. Atzei, M. Bartoletti, and T. Cimoli. A survey of attacks on Ethereum smart contracts (SoK). In *Principles of Security and Trust*, LNCS 10204, pages 164-186. Springer, 2017.
- [60] A. Haas, A. Rossberg, D. L. Schuff, B. L. Titzer, M. Holman, D. Gohman, L. Wagner, A. Zakai, and J. F. Bastien. Bringing the web up to speed with WebAssembly. In *PLDI 2017*, pages 185-200, 2017.
- [61] G. Wood. Ethereum: A secure decentralised generalised transaction ledger. Ethereum Yellow Paper, 2014 (with subsequent revisions).
- [62] X. Leroy. Formal verification of a realistic compiler. *Communications of the ACM*, 52(7):107-115, 2009.
- [63] K. Kasprzyk. The concept of smart contracts from the legal perspective. *Review of European and Comparative Law*, 34(3):101-118, 2018.
- [64] P. Lamela Seijas and S. Thompson. Marlowe: Financial contracts on blockchain. In *ISoLA 2018*, LNCS 11247, pages 356-375. Springer, 2018.
- [65] A. Cook, A. Gillioz, F. Vogelsteller, and G. Roets. ERC-1400: Security token standard. Ethereum Improvement Proposal, 2018.
- [66] R. Hanson. Combinatorial information market design. *Information Systems Frontiers*, 5(1):107-119, 2003.
- [67] J. Lange. Demand-based adjustable return financial products. U.S. Patent 6,321,212 (filed 7 July 1999; issued 20 November 2001).
- [68] J. Lange and N. Economides. A parimutuel market microstructure for contingent claims. *European Financial Management*, 11(1):25-49, 2005.

- [69] M. Peters, A. M.-C. So, and Y. Ye. A convex parimutuel formulation for contingent claim markets. Working paper, Stanford University, 2006.
- [70] K. Baron and J. Lange. *Parimutuel Applications in Finance: New Markets for New Risks*. Palgrave Macmillan, 2007.
- [71] Swiss Financial Market Supervisory Authority. Guidelines for enquiries regarding the regulatory framework for initial coin offerings (ICOs). FINMA publication, 16 February 2018.
- [72] Monetary Authority of Singapore. Payment Services Act 2019 (No. 2 of 2019) and subsequent amendments. Republic of Singapore, 2019.
- [73] U.S. Securities and Exchange Commission, Division of Corporation Finance. No-action letters: TurnKey Jet, Inc. (2019); Pocketful of Quarters, Inc. (2019); INX Limited (2020). SEC public letters, 2019-2020.
- [74] H. Adams, D. Robinson, and M. Salem. An introduction to automated market makers and decentralized exchanges. In P. Lestyán, editor, *Decentralized Finance: A Technical Primer*. 2023.
- [75] K. Qin, L. Zhou, B. Livshits, and A. Gervais. Attacking the DeFi ecosystem with flash loans for fun and profit. In *Financial Cryptography and Data Security, 2021* (preprint 2020).
- [76] H. P. Young. *Individual Strategy and Social Structure: An Evolutionary Theory of Institutions*. Princeton University Press, 1998.
- [77] H. P. Young. *Strategic Learning and Its Limits*. Oxford University Press, 2001.
- [78] B. M. Friedman. Financial innovations, monetary policy, and the efficiency of capital formation. *Journal of Monetary Economics*, 27(1):163-185, 1991.
- [79] P. Tufano. Financial innovation. In G. M. Constantinides, M. Harris, and R. M. Stulz, editors, *Handbook of the Economics of Finance*, volume 1A, pages 307-335. Elsevier, 2003.
- [80] M. A. Heller. The tragedy of the anticommons: Property in the transition from Marx to markets. *Harvard Law Review*, 111(3):621-688, 1998; anticommons applied to innovation in M. Heller, *The Gridlock Economy*, Basic Books, 2008.