SWISSI INSTITUTE FOR AI

Identity-Staked Consensus and Collusion Resistance in Chartered Validator Sets

A Trust Model for Decentralised and Compliant Distributed Settlement Infrastructure

Walter Kurz[1]

[1]Swissi Institute for AI, office@swissi-ai.institute

July 2026

Abstract

Permissioned ledgers are commonly treated as centralised because admission is restricted. This paper separates permissioning from control distribution and proposes identity-staked consensus as a trust model for accountable settlement ledgers operated by chartered validators. The model treats validator identity as externally costly collateral: public legal identity, charter state, institutional reputation, liability, attributable audit exposure, a phase-indexed conditional identity-loss floor, and loss-realisation channels outside the protocol. It distinguishes this construct from proof-of-authority by formalising validator acts as actor constellations, public validator anchoring, affiliation-aware voting caps, threshold class coverage, per-member collusion margins, observer-supported detectability, bootstrap claim discipline, and a consensus/application enforcement boundary. The paper connects the model to a broader identity-infrastructure series: the actor-assurance paper supplies capability-gate evidence, the trust-anchor paper supplies public validator anchoring and assurance-at-time, and the delegated-authority paper can consume the ledger evidence record for mandate and model-attribution records.

Keywords: identity-staked consensus; permissioned ledger; proof-of-authority; validator trust model; collusion resistance; settlement infrastructure; actor assurance; threshold class coverage; ledger evidence record

1 Introduction

Permissioning and control distribution are different design properties. Permissioning answers who may validate: the admission rule, evidence required for admission, and governance process that maintains the validator set. Control distribution answers how validation power, operational failure domains, governance influence, and collusion feasibility are spread after admission. A ledger operated by named validators can concentrate control in one institution, one jurisdiction, one vendor stack, or one founder group; it can also distribute control across mutually independent operators whose public identity makes misbehaviour attributable. The design question is the distribution of accountable control across admitted validators.

The two-axis design space separates the familiar comparators from the proposed model: openness of admission and distribution of post-admission control vary independently (fig. 1). Identity-staked consensus sits in the restricted-admission quadrant, but its decentralisation claim depends on measured control distribution across affiliation classes, attributable validator constellations, and externally realisable loss.

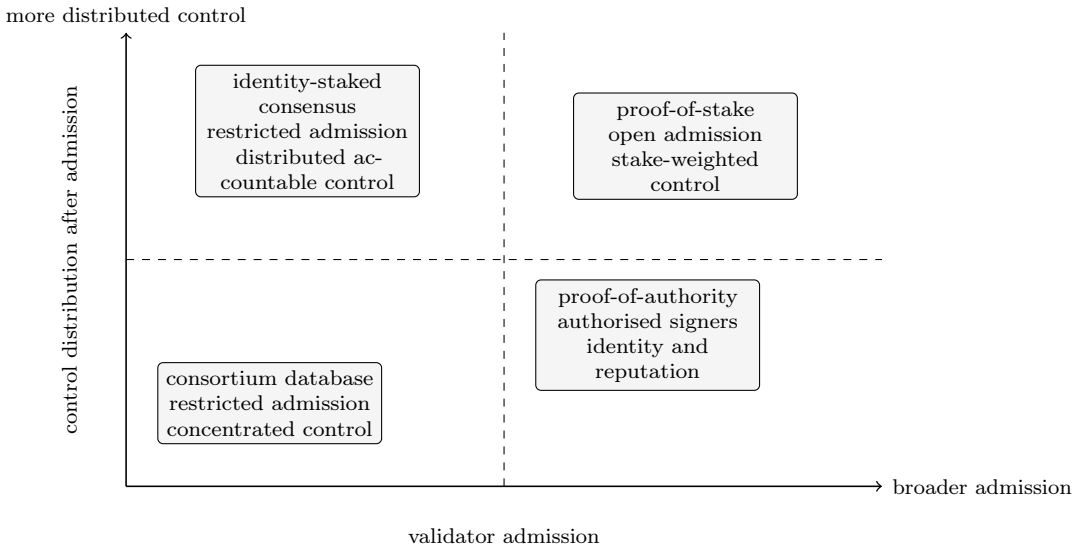

Figure 1: Permissioning and control distribution as separate ledger-design axes.

Settlement infrastructure makes that question sharper. Financial-market-infrastructure principles treat legal basis, governance, risk management, operational reliability, participation requirements, finality, and accountability as system properties for infrastructure design.[1] A digital settlement ledger that aims to serve accountable markets needs known operators, auditability, recourse, and protocol finality, while preserving resilience against unilateral rewriting and hidden capture. The usual opposition between a public anonymous chain and a private institutional database leaves that design space under-specified.

Existing consensus families supply the comparators. Ethereum's proof-of-stake design secures validation through endogenous capital posted inside the protocol and made slashable for dishonest behaviour.[2] Proof-of-authority already stakes something exogenous: authorised signers put identity and reputation behind block production, and Clique maintains the signer set and signer-voting process inside Ethereum-compatible headers.[3, 4] The missing structure is the content of that collateral: how much a given signer stands to lose, through which channel the loss is realised, how correlated the signers are across jurisdiction, ownership, and vendor stack, and how likely misconduct is to be detected, attributed, and acted on.

This paper proposes *identity-staked consensus*: a trust model that makes those collateral questions explicit for settlement infrastructure. The model consumes the actor-assurance ontology of *Multi-Jurisdictional Actor Identity Assurance for Capability Gating*, including source-graded confirmation and reliance events, capability gates, and jurisdictional anchoring.[5] It also consumes the root/profile anchor and assurance-at-time layer of *Credentials and Triangulated Trust Signals on a Single Accountable Identifier*, which gives the ledger an identity object to host or reference while reusing the identity-tier mechanics.[6] The research question is which conditions allow a permissioned validator set to be decentralised through exogenous identity stake, affiliation-class weighting, observer-supported detectability, and bootstrap governance.

The paper states those conditions as a formal trust model. It defines chartered validators as actor-assured institutions or legal arrangements, separates admission from voting power, models collusion over affiliation-constrained coalitions, and states the loss-realisation channels through which external identity stake can become costly. It then locates the ledger in a base-plus-domain topology: the base layer records stable identity, validator,

ordering, protocol-finality, and ledger-evidence-record invariants, while domain-specific compliance and activity rules remain at higher layers. A downstream delegated-authority paper can then use the ledger as a substrate for mandate records and durable model attribution while the consensus layer stays focused on validator evidence, ordering, and replay.[7]

2 Related Work

Three adjacent literatures motivate the model and leave different parts of the trust question open. The first is consensus security. Practical Byzantine Fault Tolerance gives a baseline for replicated state machines that tolerate arbitrary faulty behaviour among replicas, and later permissioned ledgers inherit this known-participant setting in different forms.[8] Ethereum's proof-of-stake documentation supplies the endogenous-collateral comparator: validators post capital inside the protocol and can lose that capital for specified protocol violations.[2] Proof-of-authority is the closer comparator because authorised signers already place identity and reputation behind block production, while Clique specifies signer authorisation, signer voting, and deauthorisation inside an Ethereum-compatible protocol.[3, 4] The open question for this paper is how exogenous identity stake is measured, diversified, made detectable, and realised as loss across a settlement validator set.

The second literature is permissioned and consortium DLT. Hyperledger Fabric presents a modular permissioned blockchain architecture with configurable consensus, portable membership, and integration with identity-management systems.[9] Corda focuses on synchronised records of agreements among mutually distrusting institutions and people, with nodes operated by a broad range of participants.[10] These systems show that restricted participation, institutional deployment, and distributed records can coexist. This paper contributes a distinct trust model: affiliation-aware voting, minimum-loss coalition reasoning, explicit loss-realisation channels, and observer-supported detectability.

The third literature concerns decentralisation metrics and settlement infrastructure. Recent decentralisation-measurement work treats decentralisation as a measurable property that depends on the resource being measured, the threshold assumed, and the metric applied; minimum-colluding-party measures are useful only after the relevant control resource is specified.[11] Financial-market-infrastructure standards require legal basis, governance, participation criteria, operational reliability, risk management, accountability, and settlement finality as infrastructure properties.[1] CPMI's DLT analytical framework applies similar safety and efficiency questions to restricted ledgers in payment, clearing, and settlement.[12] EU settlement-finality law also keeps legal finality anchored in system rules and applicable law, which is separate from protocol finality in a consensus layer.[13]

Wholesale DLT experiments show why the distinction matters in practice. Project Helvetia explored settlement of tokenised assets in central bank money, including end-to-end experimentation with commercial banks.[14] Project Jura explored cross-border wholesale CBDC transfers between French and Swiss commercial banks on a single third-party-operated DLT platform.[15] These projects support the settlement-infrastructure context; this paper supplies the validator trust model for that context. Finally, cartel-enforcement literature treats coordinated misconduct as an institutional deterrence problem shaped by detection, sanctions, leniency, enforcement capacity, and compensation mechanisms.[16] This paper uses that logic cautiously by analogy: collusion resistance depends on the minimum-loss threshold coalition and on the channels through which detected misconduct becomes costly for each required coalition member.

3 Contribution

The contribution is a formal trust model for accountable settlement ledgers operated by chartered validators. The model treats a validator as an actor-assured juridical entity, legal arrangement, or public body constituted under public law whose admission evidence is recorded through the actor-assurance event grammar. Validator signing is a $J + N + M$ constellation: a juridical actor, an accountable natural officeholder or role, and a machine actor that controls the signing node. Validator governance is a $J + N$ constellation. This is the first distinction from proof-of-authority: the model authorises key use by an accountable actor constellation under a governing-law gate, with the authorised signing key as one field in the security object.

The first formal move is to separate validator admission from validation power. Let $V = \{v_1, \ldots, v_n\}$ be the validator set. Each validator v_i has an actor identity $I(v_i)$, charter state $C(v_i)$, affiliation vector $d(v_i)$, voting weight ω_i, conditional identity-loss magnitude λ_i, protocol penalty exposure π_i, and one or more loss-realisation channels. The affiliation vector records shared-failure classes such as jurisdiction, sector, ownership or control group, public-grant source, critical vendor, governance affiliation, commercial-counterparty concentration, admission cohort, and hosting or network path. Independence is derived by comparing those vectors across validators. Voting weights are capped by institution and by shared-failure class, so another recognised institution adds validation power only when the relevant caps still have room. If a cap binds, the candidate may join as an observer or service provider with zero validation weight until diversity recovers.

The second formal move is to replace raw validator count with threshold-coalition coverage and per-member deterrence. For finality threshold q, a coalition $S \subseteq V$ is security-relevant when $\sum_{v_i \in S} \omega_i \geq q$. The paper reports the weakest affiliation diversity available to such coalitions, then asks whether each required coalition member expects misconduct to be unattractive. In compact form, the member condition is $p_i(S)\lambda_i + \pi_i > g_i(S)$, where $p_i(S)$ is the joint probability of detection, attribution, and external realisation, λ_i is conditional identity-loss magnitude given realisation, π_i is protocol penalty exposure, and $g_i(S)$ is the member's expected gain. A phase-indexed admission floor $\lambda_{\min}(B_t)$ prevents zero-stake institutions from entering as full validators while leaving calibration to implementation.

The third formal move is to specify loss realisation. External identity loss becomes stake only when governance design makes it reachable. The model requires each admitted validator to have at least one contracted or legally reachable realisation channel: consortium expulsion and forfeiture of access or governance rights, contractual liability to relying parties, litigation exposure for attributable misconduct, reputational discipline in the validator's own market, and, where misconduct falls inside the relevant external body's authority, loss of a public grant, accreditation, licence, or public-law charter. Contracted consortium consequences are design commitments. Licence, accreditation, public-grant, and public-law consequences remain jurisdiction-specific legal questions. A charter needs a realisable channel to contribute stake; pure reputation contributes softer exposure.

The fourth formal move is operator incentive-compatibility. A validator participates when the expected benefit of node operation exceeds operating cost, compliance cost, residual liability, and the option value of remaining outside the validator set. Benefits may include direct settlement access, governance voice, audit access, cost recovery, reputational positioning, and public-interest or charter-aligned duties. The model ties those benefits to service, evidentiary, and governance roles, while validation power remains capped by institution and shared-failure class.

The fifth formal move is to make detectability endogenous to the design. The probability term $p_i(S)$ depends on attributable signatures, data availability, evidence retention, dispute records, and observer capability. The ledger design can provide read access, audit trails, and escalation interfaces to auditors, public-interest observers, and supervisors whose own authority permits participation. Observer activity follows those actors' legal or contractual roles; validation remains separate by default. The scope assumption is fixed: the settlement asset is regulated, validators are chartered, validator operation follows consortium governance, and supervisory authority

over validation arises only where an applicable legal regime creates it.

The sixth formal move is bootstrap honesty. Early validator cohorts have disclosed concentration risk, so the model introduces a bootstrap state B_t with admission sequencing, founder caps, public audit, class-diversity targets, and sunset rules. Stronger decentralisation claims arise only when the measured affiliation classes and minimum-loss threshold coalition support them. Admission state is a consensus-owned state variable for each validator; admission-rule amendments are versioned governance objects that are anchored to the base layer, become effective only through a defined supermajority and delay, and leave historical validator state replayable under the rule version in force at the relevant time.

The seventh formal move is the consensus/application boundary. Let $E(x) \in \{L_1, L_2\}$ classify a candidate enforcement rule or record invariant. The base layer owns invariants whose violation would compromise the ledger's evidentiary substrate: validator admission state, signature validity, ordering, protocol finality, hash integrity, identity-anchor uniqueness, and ledger evidence record integrity. Higher layers own domain rules that change with law, market practice, or supervisory interpretation; they anchor their evidence to the base ledger while preserving their own amendment paths.

4 Research

4.1 Inputs from actor assurance and anchoring

This paper uses the preceding identity papers only at the interface where validator status becomes a consensus-security object. The actor-assurance paper supplies the gate grammar, actor constellations, source-graded confirmation events, reliance events, and jurisdictional anchoring needed to decide whether an institution may enter the validator set.[5] The trust-anchor paper supplies the root/profile anchor and assurance-at-time layer needed to bind validator evidence to a replayable identity object.[6] This paper adds the control-distribution and collusion model that belongs to the ledger layer.

The validator gate is stricter than ordinary institutional presentation because validation is an infrastructural act. For validator signing, the admissible constellation is $J + N + M$: the juridical institution or public-law entity, an accountable natural officeholder or role, and the machine actor that controls the signing node. For validator governance, the admissible constellation is $J + N$. The validator gate g_V consumes live status, governing-law evidence, a register or public anchor, authority-bearing roles, charter or purpose evidence, an accountability path, record-retention commitments, at least one realisation channel, and a phase-indexed minimum conditional identity-loss floor. Equation (1) turns the general capability-gate form into a validator-admission predicate for consensus security.

$$\mathrm{adm}_{g_V}(v_i, t) = 1 \iff \mathrm{pass}_{g_V}(I(v_i), a_V, t) = 1 \wedge C_t(v_i) = 1 \wedge R_t(v_i) \neq \varnothing \wedge \lambda_i(t) \geq \lambda_{\min}(B_t). \tag{1}$$

Here $\mathrm{adm}_{g_V}(v_i, t)$ is validator admission at time t, $I(v_i)$ is the admitted actor identity, a_V is the validator act being evaluated, $C_t(v_i)$ is live charter state, $R_t(v_i)$ is the set of realisation channels, $\lambda_i(t)$ is the externally realisable loss magnitude conditional on realisation, and $\lambda_{\min}(B_t)$ is the minimum conditional loss required in the current bootstrap phase. The probability of realisation enters through $p_i(S, t)$ in the detectability term.

Validator anchoring also differs from ordinary profile anchoring. The trust-anchor paper keeps ordinary root/profile bindings off-chain to support profile-scoped presentation. Validator admission uses a role-scoped public anchor instead: the validator-specific institutional profile anchor $h_V(v_i)$ is linked to the root institutional anchor for attribution, replay, and retention. That waiver is narrow. It covers the validator role, the validator set, and the evidence needed to replay admission and signing authority. Validator service also sets a retention obligation for

the validator anchor and validator gate during the service period, dispute window, and applicable record-retention window. The ledger stores or anchors the admission result, rule version, evidence pointer, public validator anchor, and retention pointer; the upstream identity papers continue to supply the detailed actor-assurance and profile-binding mechanics.

4.2 Validator state and affiliation

At time t, the validator set is $V_t = \{v_1, \ldots, v_n\}$. Each admitted validator record contains the admission predicate $\mathrm{adm}_{g_V}(v_i, t)$, public validator anchor $h_V(v_i)$, actor identity $I(v_i)$, live charter state $C_t(v_i)$, affiliation vector $d_t(v_i)$, voting weight $\omega_i(t)$, conditional identity-loss magnitude $\lambda_i(t)$, protocol penalty exposure $\pi_i(t)$, and realisation-channel set $R_t(v_i)$. The tuple is kept in prose because the later model consumes the fields directly. Admission, voting, and loss are separate objects: admission states who may validate, voting states how much control the admitted validator carries, and loss states what the validator stands to lose if misconduct is detected, attributed, and realised.

The affiliation vector is a validator-level shared-failure object. It is related to the actor-assurance paper's confirmation-event failure classes, but it works one layer higher: actor assurance discounts correlated evidence, while this paper discounts correlated validators as independent control. The main classes in table 1 identify concentration channels that can make several named validators respond as one control bloc.

Table 1: Affiliation classes used to measure validator-level shared-failure risk.

Class	Reason for inclusion	Example concentration risk
Jurisdiction	Public law, courts, sanctions, and emergency powers can affect validators together	Several validators exposed to one national measure
Sector	Sectoral incentives and business cycles can align conduct	Banks or insurers facing the same market pressure
Ownership or control group	Formal independence can mask common control	Subsidiaries validating as separate institutions
Public mandate source	The same grantor or supervisor can shape institutional incentives	Two banks under one supervisor share more than jurisdiction
Critical vendor	Shared software, custody, or operational provider creates a technical failure mode	Common validator client or key-management provider
Governance affiliation	Common consortium governance can align voting beyond technical validation	Founder bloc retaining amendment control
Commercial-counterparty concentration	Shared reliance on a narrow customer or relying-party set creates commercial pressure	Validators dependent on one settlement sponsor
Admission or founder cohort	Actors admitted under the same origin conditions may share implicit loyalties	Genesis validators preserving early allocation choices
Hosting region or network path	Infrastructure concentration can persist across nominally different vendors	Distinct providers using one availability region or network route

Voting weight is assigned after admission and then constrained by affiliation class. For each capped class $D_k(t)$ extracted from $d_t(v_i)$, Equation (2) prevents admission growth from becoming silent control concentration.

$$\sum_{v_i \in D_k(t)} \omega_i(t) \leq \alpha_k \qquad \forall k \in K. \tag{2}$$

Here K is the set of capped affiliation dimensions, $D_k(t)$ is the validator subset sharing a class value on dimension k, $\omega_i(t)$ is voting weight, and α_k is the maximum aggregate weight allowed for that class. A candidate whose class would breach a cap can be admitted as an observer or service provider with zero validation weight until diversity recovers.

4.3 Coalition coverage and margin

Coalition analysis begins with the threshold that can affect finality, ordering, or rewriting under the selected protocol. Equation (3) defines the set of coalitions that matter for that threshold by voting weight.

$$\mathcal{S}_q(V_t) = \left\{ S \subseteq V_t : \sum_{v_i \in S} \omega_i(t) \geq q \right\}. \tag{3}$$

The term q is the finality or control threshold used for the claim under analysis. A coalition with less than q may still create operational disruption; threshold-control claims in this model use $\mathcal{S}_q(V_t)$.

Class caps are governance rails, while coverage is the publishable decentralisation metric. Equation (4) reports the weakest diversity present in threshold coalitions along each affiliation dimension.

$$c_{q,k}(V_t) = \min_{S \in \mathcal{S}_q(V_t)} \left| \{ d_{i,k} : v_i \in S \} \right|. \tag{4}$$

Here k names an affiliation dimension, and $d_{i,k}$ is validator v_i's class value on that dimension. A low $c_{q,k}$ means some threshold coalition can be formed with too few distinct class values on dimension k, even if the validator count appears large.

A single headline value can report the weakest class dimension when the governance process needs one public number. Equation (5) compresses the coverage vector to the lowest threshold diversity across capped classes.

$$n_{\text{eff}}(q,t) = \min_{k \in K} c_{q,k}(V_t). \tag{5}$$

The headline value is a minimum threshold-diversity report over the affiliation classes that governance has chosen to measure. A consortium can inflate the value by choosing weak classes, so class definitions must be published and audited before the value supports a decentralisation claim.

Coverage says whether threshold coalitions are diverse enough to support a decentralisation claim. Margin asks whether each required member expects misconduct to be unattractive. Equation (6) avoids the average-stake error by evaluating the expected deterrence condition for each member of a candidate coalition.

$$m_i(S,t) = p_i(S,t)\lambda_i(t) + \pi_i(t) - g_i(S,t). \tag{6}$$

Here $m_i(S,t)$ is validator v_i's margin inside coalition S, $p_i(S,t)$ is the probability of detection, attribution, and external realisation, $\lambda_i(t)$ is conditional identity-loss magnitude, $\pi_i(t)$ is protocol penalty exposure, and $g_i(S,t)$ is the validator's expected gain from the coalition. The coalition is deterred only when the relevant members have positive margins.

The model's system margin is the weakest member in the weakest threshold coalition. Equation (7) identifies the governance repair target by focusing on the lowest exposed member across threshold coalitions.

$$\Phi_q(V_t) = \min_{S \in \mathcal{S}_q(V_t)} \min_{v_i \in S} m_i(S,t). \tag{7}$$

A positive $\Phi_q(V_t)$ supports the claim that every threshold coalition is unattractive for each required member under the model's estimates. A non-positive value identifies a weak coalition and points governance toward higher detectability, higher conditional loss, a stronger protocol penalty, lower correlated weight, or different admission.

4.4 Detectability, observers, and participation

The probability term in the member margin is itself a design object. Equation (8) fixes the convention: $\lambda_i(t)$ is a conditional loss magnitude, and $p_i(S,t)$ carries the probability that misconduct is detected, attributed, and externally realised.

$$p_i(S,t) = \Pr(D_i \cap A_i \cap X_i \mid S,t). \tag{8}$$

The event D_i is detection of conduct relevant to validator v_i, A_i is attribution to the validator or its authorised $J+N+M$ signing constellation, and X_i is realisation through a channel in $R_t(v_i)$. Data availability, attributable signatures, key-control records, rule-version history, evidence retention, and dispute records raise D_i and A_i. Contract drafting, procedural enforceability, market visibility, and external authority determine X_i.

Observers raise detectability while carrying zero validation weight. Let O_t be the observer set. An observer admitted under gate g_O has attributable identity, audit or reporting rights, and zero voting weight. Observer independence is measured through the same affiliation vector used for validators, because an observer affiliated with the coalition it monitors adds little to detection or attribution. Observer reports are L1-anchored and evidentiary: they create durable evidence and escalation paths while leaving finality unchanged.

External loss realisation remains partly contractual and partly jurisdiction-specific. Consortium expulsion, forfeiture of access or governance rights, and contractual liability are channels that the ledger governance can design directly. Litigation exposure, reputational discipline, and loss of a public grant, accreditation, licence, or public-law charter depend on governing law, institutional form, and the authority of the external body. The paper treats those channels as model requirements and future evidence-template work; governing-law analysis determines whether a given validator faces a specific legal consequence.

The validator also needs an honest reason to participate when voting weight is capped. Equation (9) states the participation condition as a separate trade-off from collusion deterrence.

$$U_i(t) \geq K_i(t) + H_i(t). \tag{9}$$

In this condition, $U_i(t)$ is the validator's expected participation benefit, including settlement access, governance voice, audit access, cost recovery, reputational positioning, and public-interest or charter-aligned duties. $K_i(t)$ is operating and compliance cost, and $H_i(t)$ is residual liability from honest participation. Benefits come from service, evidentiary, and governance roles under capped validation power.

4.5 Bootstrap claims and enforcement boundary

Bootstrap governance is a phase-indexed claim-control mechanism. The bootstrap state B_t carries the phase, founder caps, class-diversity targets, phase-specific $\lambda_{\min}(B_t)$, observer rights, sunset rules, and audit cadence. Early phases can support replayable evidence, attributable validation, and published concentration metrics. Mature decentralisation claims require class caps, threshold coverage, positive coalition-margin review, observer capability, and phase exits governed by published rules. The claim ladder in table 2 assigns each phase the claim level supported by its evidence.

The enforcement boundary classifies candidate rules by replayability and amendment locus. Equation (10) keeps mutable domain law out of consensus by assigning a rule to the base layer only when breaching it would make prior records unreplayable and when its amendment path belongs to the ledger's versioned protocol process.

$$E(x) = L_1 \iff \mathrm{sub}(x) = 1 \wedge \mathrm{stable}(x) = 1. \tag{10}$$

The predicate $\mathrm{sub}(x) = 1$ means that violating x breaks replay of prior records. The predicate $\mathrm{stable}(x) = 1$ means that historical records replay under the rule version active at record time and that the rule is amended through the ledger's versioned protocol process. Domain-authority amendment places x at L_2 or a higher layer.

Table 2: Bootstrap phases and the claims each phase can support.

Phase	Supported claim	Required evidence	Claim boundary
Pilot	Replayable records and attributable validation	Named validators, rule versions, signatures, retention, and disclosed concentration	Concentration metrics define the claim ceiling
Guarded growth	Measured control distribution under published caps	Affiliation classes, cap compliance, observer access, phase-specific $\lambda_{\min}$	Independence claims require class evidence
Mature operation	Distributed accountable control under current metrics	Threshold coverage, coalition-margin review, observer reports, sunset completion, and public audit cadence	Legal-finality and immunity claims require external proof

The boundary is easiest to audit by applying both tests to candidate invariants and domain rules. Table 3 classifies each example by whether a violation breaks replay and where amendment authority sits.

Table 3: Consensus and application boundary under replayability and amendment-locus tests.

Invariant or rule	Layer	Replayability test	Amendment locus
Validator admission state	L_1	Past signatures require historical admission state	Versioned ledger governance
Signature validity	L_1	Invalid signature rules break attribution of prior records	Versioned ledger governance
Ordering and protocol finality	L_1	Reordering breaks record sequence and finality replay	Versioned ledger governance
Hash integrity	L_1	Payload and predecessor hashes are the replay substrate	Versioned ledger governance
Identity-anchor uniqueness	L_1	Duplicate anchors break attribution and state reconstruction	Versioned ledger governance
Evidence-record integrity	L_1	Missing evidence pointers break audit and dispute replay	Versioned ledger governance
Product eligibility	L_2	Past records remain replayable as records	Domain law, market rule, or supervisor
Reporting format	L_2	Payload meaning can be transformed while record order persists	Domain authority or application governance
Market conduct rule	L_2	The ledger records the event while rule authority sits elsewhere	Domain law or market governance
Supervisory interpretation	L_2	Interpretation changes leave historical ledger facts replayable	External authority

4.6 Ledger evidence record

The base layer stores or anchors the fields needed to replay who acted, under which validator set, under which rule version, and with which evidence pointers. Equation (11) defines the ledger evidence record as a substrate tuple for domain payloads.

$$e = (e_{\mathrm{id}}, t_r, t_v, \mathrm{type}, h_a, V_t, r_t, E(x), h_p, h_{-1}, \sigma, \rho, \delta, o). \tag{11}$$

Here e_{id} is the record identifier, t_r is record time, t_v is valid time where applicable, type is the record class, h_a is the actor or validator anchor, V_t identifies the validator set, r_t is the rule version, $E(x)$ is the enforcement layer, h_p is the payload hash, h_{-1} is the predecessor hash, σ contains attributable signatures, ρ points to retention metadata, δ points to dispute metadata, and o points to observer reports where present.

The record is generic enough for settlement events, validator-state transitions, observer reports, reliance records, and later delegated-authority records. The downstream delegated-authority paper can bind mandate and model-

attribution semantics to the payload or downstream layer, while this paper supplies only the attributable, ordered, replayable substrate.[7]

5 Discussion

The model changes the meaning of decentralisation for settlement infrastructure. In an anonymous token-staked system, the control resource is endogenous stake. In a chartered validator system, the control resource is the combination of voting weight, affiliation class, operational independence, detectability, and conditional external loss. A restricted validator set can be decentralised when threshold coalitions remain diverse across measured classes and when the weakest coalition remains unattractive for each required member. Permissioning is then an admission property; decentralisation is a measured distribution property.

The core distinction from proof-of-authority is the actor constellation. A PoA design can maintain authorised signing keys and signer votes. Identity-staked consensus adds a legal and evidentiary wrapper around key use: the validator signature is attributable to a $J + N + M$ constellation, governance acts are attributable to $J + N$, and both are admitted through a gate whose evidence remains replayable. The security claim rests on verifiable actor conditions: public identity, charter state, retention obligation, realisation channel, affiliation class profile, and conditional loss floor.

The series argument is the same independence discipline applied at two layers. The trust-anchor paper discounts correlated evidence when a relying party evaluates assurance-at-time for a gate. This paper discounts correlated validators when a ledger evaluates threshold control. At both layers, repetition from one failure mode has low probative value: several credentials from one failure mode leave the identity presentation weak, and several validators inside one affiliation class leave the decentralisation claim weak.

The hard case is loss realisation. Institutional identity has security value only when misconduct can be detected, attributed, and made costly. Reputation alone is a weak and uneven form of stake. Contractual liability, expulsion, loss of access, litigation exposure, and public-grant consequences create stronger channels, but each channel depends on drafting, governing law, evidence retention, and procedural enforceability. This is why $p_i(S, t)$ is inside the member margin and why $\lambda_i(t)$ is conditional on realisation. A ledger lacking attributable evidence lowers the expected cost of collusion even when every validator is named.

The observer role is narrower than the technology may suggest. The ledger can expose read access, audit trails, dispute records, and escalation interfaces. Auditors, public-interest observers, and supervisors act through their own legal or contractual authority after admission through g_O. Their reports can strengthen detection and attribution while carrying zero validation weight by default. This keeps the claim hierarchy intact: the settlement asset may sit inside a regulated framework, validators may already be chartered institutions, and validator operation is governed by the consortium and by any legal perimeter that applies to it. Supervisory status remains a separate legal question.

Bootstrap governance is the main adoption risk. Early validator cohorts will be smaller, more correlated, and more dependent on founder choices than a mature validator set. The model handles this by requiring phase-specific claims: a pilot can claim replayable evidence, attributable validation, and published concentration metrics; mature decentralisation claims require affiliation caps, threshold coverage, observer capability, realisation channels, and a positive coalition-margin review. The same logic applies when admitted institutions share a vendor stack, national exposure, ownership group, public-grant source, founder cohort, or hosting region. Each shared-failure class reduces effective independence until caps and governance repair the concentration.

Within the paper series, this paper supplies the record substrate between identity assurance and delegated machine action. The actor-assurance paper supplies evidence for validators and relying parties. The trust-anchor

paper supplies the root/profile anchor and assurance-at-time structure. This paper adds the validator trust model, the layer boundary, and the ledger evidence record. The delegated-authority paper can then bind agent activity and durable model attribution to records whose consensus layer already has accountable validators, replayable rule versions, and a defined enforcement boundary.

6 Conclusion

This paper has defined identity-staked consensus as a trust model for settlement ledgers operated by chartered validators. Its core move is to separate admission from decentralisation: admission determines who may validate, while decentralisation depends on how voting weight, operational dependencies, governance influence, detectability, and externally realisable identity loss are distributed after admission. A named validator set becomes security-relevant only when those post-admission properties can be measured and replayed.

The model makes that claim inspectable. Actor constellations determine which juridical, natural, and machine actors may sign or govern; public validator anchoring binds that role to replayable institutional evidence; affiliation caps and threshold coverage test whether control is distributed across meaningful classes; and the member-margin equations ask whether each required coalition member faces sufficient expected loss. Bootstrap governance then ties public claims to the evidence supported by the current phase, so pilot operation, guarded growth, and mature operation carry different evidentiary burdens.

For the paper series, this paper supplies the ledger substrate between legal identity assurance and delegated machine action. Its evidence record preserves who acted, under which validator set and rule version, with which signatures, evidence pointers, dispute pointers, and observer reports. The delegated-authority paper can bind mandate semantics and durable model attribution to that substrate while leaving consensus focused on accountable validators, ordering, replay, and enforcement boundaries. The result is a consensus account in which institutional exposure becomes a measurable resource for settlement trust.

7 Limitations and Future Research

The model is conceptual and formal; empirical validator census work remains future work. Future work must test candidate validator populations against the affiliation vector, shared-failure caps, threshold class coverage, and minimum-loss threshold-coalition margin. This includes sectoral, jurisdictional, ownership, commercial-counterparty, founder-cohort, hosting, vendor, and governance correlations.

The minimum conditional identity-loss floor is a calibration problem. This paper places $\lambda_{\min}(B_t)$ inside validator admission and leaves phase values unset for banks, insurers, universities, public bodies, chambers, standards organisations, and other chartered validators. A high floor strengthens deterrence and may exclude useful institutions whose participation improves diversity, observability, or public legitimacy. A low floor improves inclusion and may weaken the smallest threshold coalition. Implementation must calibrate that trade-off before the ledger claims mature collusion resistance.

The threshold coverage metric depends on governance-quality class definitions. A consortium that defines affiliation classes too coarsely can hide concentration; one that defines them too finely can inflate diversity. Future work must specify class-definition governance, external audit of class assignments, and methods for measuring higher-order correlations across jurisdiction, sector, ownership, public-grant source, vendors, counterparties, founder cohort, and hosting path.

Loss-realisation channels require jurisdiction-specific legal analysis. Consortium expulsion, contractual liability, litigation exposure, reputational discipline, and public-grant consequences differ across banks, insurers, universities, public bodies, chambers, and standards organisations. The paper states the model requirement; future work must map concrete evidence templates, trigger evidence, realising actors, and enforceability rules.

Implementation and proof obligations remain open. A deployed protocol would need precise validator-key governance, evidence-retention rules, observer access control, dispute procedure, amendment mechanics, data-availability guarantees, and a formal security proof under stated network and adversary assumptions. Bootstrap governance also needs live measurement: a ledger should publish the claims supported by its current phase while mature-set claims are reserved for the measured mature state.

References

1. Committee on Payment and Settlement Systems and Technical Committee of the International Organization of Securities Commissions. Principles for Financial Market Infrastructures. Bank for International Settlements and International Organization of Securities Commissions, 2012. Available from: https://www.bis.org/cpmi/publ/d101a.pdf [Accessed on: 2026 Jul 24]
2. Ethereum.org. Proof-of-stake (PoS). 2026. Available from: https://ethereum.org/developers/docs/consensus-mechanisms/pos/ [Accessed on: 2026 Jul 24]
3. Ethereum.org. Proof-of-authority (PoA). 2026. Available from: https://ethereum.org/developers/docs/consensus-mechanisms/poa/ [Accessed on: 2026 Jul 24]
4. Szilágyi P. EIP-225: Clique proof-of-authority consensus protocol. 2017. Available from: https://eips.ethereum.org/EIPS/eip-225 [Accessed on: 2026 Jul 24]
5. Kurz W. Multi-Jurisdictional Actor Identity Assurance for Capability Gating. Manuscript in the Swissi core-concept paper series, Paper 14. 2026
6. Kurz W. Credentials and Triangulated Trust Signals on a Single Accountable Identifier. Manuscript in the Swissi core-concept paper series, Paper 15. 2026
7. Kurz W. Durable Model-Configuration Attribution and Recovery-Conditioned Mandates for Autonomous Economic Agents. Manuscript in the Swissi core-concept paper series, Paper 17. 2026
8. Castro M and Liskov B. Practical Byzantine Fault Tolerance. *Proceedings of the Third Symposium on Operating Systems Design and Implementation (OSDI 99)*. New Orleans, LA: USENIX Association, 1999 Feb. Available from: https://www.usenix.org/conference/osdi-99/practical-byzantine-fault-tolerance [Accessed on: 2026 Jul 24]
9. Androulaki E, Barger A, Bortnikov V, Cachin C, Christidis K, De Caro A, Enyeart D, Ferris C, Laventman G, Manevich Y, Muralidharan S, Murthy C, Nguyen B, Sethi M, Singh G, Smith K, Sorniotti A, Stathakopoulou C, Vukolić M, Weed Cocco S, and Yellick J. Hyperledger Fabric: A Distributed Operating System for Permissioned Blockchains. Proceedings of the Thirteenth EuroSys Conference. 2018. DOI: 10.1145/3190508.3190538. Available from: https://arxiv.org/abs/1801.10228 [Accessed on: 2026 Jul 24]
10. R3. The Corda Platform: An Introductory White Paper. 2018. Available from: https://r3.com/the-corda-platform-an-introduction-whitepaper/ [Accessed on: 2026 Jul 24]
11. Ovezik C, Karakostas D, Kiayias A, and Woods DW. SoK: Measuring Blockchain Decentralization. 2025. Available from: https://arxiv.org/html/2501.18279v1 [Accessed on: 2026 Jul 24]
12. Committee on Payments and Market Infrastructures. Distributed Ledger Technology in Payment, Clearing and Settlement: An Analytical Framework. CPMI Papers No 157. Bank for International Settlements, 2017. Available from: https://www.bis.org/cpmi/publ/d157.htm [Accessed on: 2026 Jul 24]

13. European Parliament and Council of the European Union. Directive 98/26/EC of 19 May 1998 on Settlement Finality in Payment and Securities Settlement Systems. 1998. Available from: https://eur-lex.europa.eu/eli/dir/1998/26/oj/eng [Accessed on: 2026 Jul 24]

14. BIS Innovation Hub, Swiss National Bank, and SIX. Project Helvetia: A Multi-Phase Investigation on the Settlement of Tokenised Assets in Central Bank Money. 2026. Available from: https://www.bis.org/about/bisih/topics/cbdc/helvetia.htm [Accessed on: 2026 Jul 24]

15. BIS Innovation Hub, Banque de France, and Swiss National Bank. Project Jura: Cross-Border Settlement Using Wholesale CBDC. 2026. Available from: https://www.bis.org/about/bisih/topics/cbdc/jura.htm [Accessed on: 2026 Jul 24]

16. Organisation for Economic Co-operation and Development. Recommendation of the Council Concerning Effective Action against Hard Core Cartels. OECD/LEGAL/0452. OECD, 2025. Available from: https://legalinstruments.oecd.org/api/print?ids=652&lang=en [Accessed on: 2026 Jul 24]

Contact

Dr. Walter Kurz, MBA, MSc
Präsident des Verwaltungsrates

Swissi Institute for AI
Swissi Holding AG
Baarerstrasse 78
6300 Zug
Switzerland

UID CHE-189.603.726