Proof-of-Biometric-Uniqueness (PoBU): A Scarcity Primitive for One-Human-One-Node Blockchain Consensus

Dato Kavazi, Viktor Vernissage and Humanode Core

2025-12-31

Abstract

Permissionless consensus must impose a participation cost sufficient to resist Sybil attacks [1]. Proof-of-Work and Proof-of-Stake price identities via computation or capital [2, 3], but their scarce resources can concentrate, enabling durable capture by industrial or wealthy coalitions.

This paper formalizes Proof-of-Biometric-Uniqueness (PoBU) as an alternative scarcity primitive in which baseline eligibility for PoBU-weighted roles is bounded by verified unique humans rather than by energy or stake. PoBU is inherently probabilistic: the one-human-one-eligible-account goal holds except with system-defined failure probabilities determined by biometric error rates and operational policy.

We provide: (i) a protocol-level definition of PoBU and its eligibility interface (prove-uniqueness, on-chain notification, wipe/refresh, and optional renewal/revocation); (ii) a mapping from human-bounded baseline weight to representative consensus safety thresholds; (iii) a threat taxonomy emphasizing issuer concentration, eligibility compromise/coercion, and availability; and (iv) an evaluation approach grounded in reproducible chain-derived measurements. Using Humanode as a reference deployment boundary, we report preliminary validator-set, churn, and block-author concentration statistics extracted via direct Substrate RPC (with disclosed sampling); in a recent 90-day window, the top-10 block authors produced 4.5% of sampled blocks.

Keywords. Sybil resistance; scarcity primitives; Proof-of-Work; Proof-of-Stake; Proof-of-Biometric-Uniqueness; confidential computing; remote attestation; biometric information protection; validator decentralization metrics; Substrate.

1 Introduction

Permissionless systems must decide what an "independent participant" means. If identities are cheap to create, redundancy and voting can be subverted by adversaries who instantiate many pseudonyms. The Sybil attack formalizes this baseline obstacle: strong Sybil resistance in open membership systems requires either a scarce-resource test or a certification path that binds eligibility to an external property [1, 4].

Blockchains operationalize that choice through scarcity primitives. Proof-of-Work (PoW) prices eligibility with externally scarce computation/energy, while Proof-of-Stake (PoS) prices it with internally scarce stake [2, 3]. Both can deter Sybil identities yet remain vulnerable to capture because the scarce resource can concentrate (through mining pools, stake compounding, etc.). Section 2 details those capture modes and motivates a third option.

Proof-of-Biometric-Uniqueness (PoBU) adopts that third option: baseline eligibility for PoBU-weighted roles is bounded by verified unique humans, with operational failure probabilities determined by biometric error rates and policy choices. PoBU is about enforcing "one eligible account per human" (with quantified bounds) rather than storing civil identity or biometrics on-chain. The rest of the paper clarifies what PoBU means operationally, how Humanode realizes a reference boundary, and which measurements can test the concept (Sections 4–7).

1.1 Contributions

This paper aims to:

- Formalize PoBU as a protocol property: an eligibility interface (prove-uniqueness, on-chain notification, wipe/refresh, and optional renewal/revocation) and a human-bounded baseline-weight interpretation (Section 4).
- Specify a threat taxonomy for PoBU systems emphasizing issuer concentration, eligibility compromise/coercion, availability, and linkability risks (Sections 5–6).
- Provide a reproducible empirical evaluation approach using public chain data (direct Substrate RPC), and report initial results from a live PoBU chain (Section 7).
- Use Humanode as a reference deployment boundary to make the off-chain/on-chain split explicit without defining PoBU as Humanode-specific (Section 5).

1.2 Paper organization

Section 2 frames the problem and the evaluation questions. Section 3 surveys related work. Section 4 defines PoBU primitives and measurable parameters. Sections 5–7 summarize Humanode's implementation boundary, analyze security, and present a measurement approach with preliminary statistics. Section 8 states limitations and future work, and Section 9 concludes.

2 Motivation and Problem Statement

2.1 Sybil Resistance and Scarcity Primitives

Blockchains select a scarcity primitive that defines eligibility. In PoW/PoS this primitive is a resource whose acquisition and control can be quantified in economic terms (hashpower, stake). That quantifiability creates two practical effects:

- The resource can be pooled, bartered, or compounded, making it easier for a coalition to grow influence even without gaining new distinct humans.
- The same resource remains on-chain or in plain sight, so governance or consensus outcomes become predictable functions of capital ownership, which can lessen the marginal cost of capture.

We therefore ask how PoW and PoS shift the capture surface toward these dynamics and why a different primitive (PoBU) aims to alter the scaling variable.

2.2 PoW: External Scarcity and Industrial Coordination

PoW makes identity replication costly through external computation/energy expenditure. However, decentralization is not guaranteed by protocol mechanics alone. Cong, He, and Li analyze how mining pools and industrial organization reshape the control surface even when the protocol is permissionless [5]. In practice, PoW tends to select for specialized hardware and coordinated operations, moving the dominant capture surface from protocol rules to industrial economics.

2.3 PoS: Internal Scarcity and Wealth Concentration

PoS makes identity replication costly through stake. Two structural concerns then appear:

- **Wealth dynamics.** Fanti et al. model stake compounding and show that common reward allocations can be inequitable even under honest behavior [6].

- **Participation incentives.** In token-voting systems, participation can concentrate among large holders. Buterin describes this dynamic directly:

> "one of them is plutocracy which is that only the rich people really have an incentive to participate." [7]

PoS therefore tends to reduce Sybil risk by using capital as eligibility, while leaving open capture pathways tied to wealth concentration, delegation, and collusion.

2.4 PoBU: Human-Bounded Eligibility

PoBU is motivated by separating two questions that are often conflated:

- **Anti-Sybil:** how expensive is it to create many eligible participants?
- **Anti-capture:** can a small coalition obtain durable control, and how does the cost scale?

Proof-of-Biometric-Uniqueness (PoBU) proposes a different scarcity primitive: treat *verified unique humans* as the scarce resource that bounds baseline eligibility for PoBU-weighted roles. In contrast to PoW/PoS, the intended scaling variable for baseline control is the number of distinct humans an adversary can recruit and maintain, not the amount of capital it can compound.

PoBU necessarily relies on a certification path (the uniqueness/liveness verification process). The core design question becomes how to minimize and constrain that trust surface (confidential verification, auditability, revocation/recovery, issuer decentralization), and how to parameterize its probabilistic failure modes (Section 4).

PoBU can be realized by different mechanisms. In this paper we distinguish two broad classes: *credentialed* PoBU (long-lived attestations with revocation/renewal) and *freshness-gated* PoBU (short-lived authentications requiring periodic re-verification). Both satisfy the same PoBU eligibility interface, but they shift risks differently (Section 4).

2.5 Comparison of PoW, PoS, and PoBU Scarcity Primitives

Table 1 provides a compact comparison of PoW, PoS, and PoBU as scarcity primitives, highlighting how each shifts the dominant capture surface.

Property / risk surface	PoW (one-hashpower-one-vote)	PoS (one-coin-one-vote)	PoBU (one-human-one-eligible account)
Anti-Sybil primitive	External resource cost	Internal capital cost	Real-world uniqueness check
Centralization pressure	Industrial mining, pools [5]	Wealth compounding / cartelization [6]	Issuer concentration / coercion
Bribery / rental	Indirect (hashpower markets)	Direct (delegation / vote markets)	Direct (eligibility rental)
Recovery after compromise	Hard (economic inertia)	Hard (economic inertia)	Possible via expiration/deauthentication and re-verification
Privacy risks	Low (pseudonymous)	Low (pseudonymous)	High unless privacy tech is used

Table 1. Comparison of PoW, PoS, and PoBU as scarcity primitives.

2.6 Evaluation Questions and Evidence

This paper evaluates PoBU using only publicly reproducible, chain-derived measurements (Section 7). Concretely, we ask:

- Does the eligible validator set remain broad over time (set size and churn), or does it collapse to a small stable core?
- How concentrated is block production across validator keys (top-k, HHI, Gini), under disclosed sampling?
- Do publicly disclosed identity-layer operational events correlate with on-chain observables (availability shock indicators), without making causal claims from incomplete incident data?

Measurements that would further strengthen the evaluation but are not reported here include identity-layer aggregates (issuance/revocation counts, uniqueness and PAD error summaries) and governance participation distributions; these remain future work unless publishable as safe aggregates (Section 8).

3 Related Work

PoBU sits at the intersection of Sybil-resistance foundations, resource-based consensus (PoW/PoS), proof-of-personhood approaches, biometric recognition and biometric information protection, privacy-preserving credentials, and confidential-computing-based private verification. We reference these areas only to position PoBU, not to provide a full survey.

3.1 Sybil Resistance: Foundations and Taxonomies

Douceur formalizes the core issue: if identities are cheap to create, redundancy and voting-based mechanisms can be undermined by a single adversary controlling many pseudonyms [1]. Subsequent surveys group defenses into certification, resource testing, and social-network-based approaches, reinforcing that strong Sybil resistance generally requires external assumptions that a protocol must make explicit and constrain [4].

3.2 PoW and PoS as Scarcity Primitives

Resource-based consensus uses scarce-resource tests as an eligibility and weight primitive. The foundational points of comparison are Nakamoto's PoW construction and early PoS proposals [2, 3]. Subsequent work analyzes how industrial coordination (e.g., mining pools) and economic dynamics can concentrate control in PoW and PoS settings [5, 6, 8], and proposes quantitative measures for how "egalitarian" different currencies are under their chosen scarcity primitives [9].

PoBU differs on the axis these works typically hold fixed: the unit of baseline eligibility is bounded by verified unique humans rather than by capital or compute. This shifts the primary system constraints toward certification integrity, eligibility lifecycle, and availability (Sections 4–6), while leaving collusion and markets as residual risks rather than as eliminated phenomena.

3.3 Proof-of-Personhood and Identity-Based Participation

Borge et al. propose proof-of-personhood (PoP) and explicitly position it as a response to PoW/PoS drawbacks [10]. PoP uses event-based "pseudonym parties" and accountable anonymous credentials to issue one participation token per physical attendee while attempting to preserve anonymity [10].

PoBU shares the goal of binding physical entities to virtual eligibility, but differs operationally. Rather than periodic in-person ceremonies, PoBU relies on biometric uniqueness and liveness verification with an online eligibility lifecycle (prove-uniqueness, on-chain notification, wipe/refresh, and optional renewal/revocation) designed for continuous operation in a live chain.

3.4 Biometrics, Template Protection, and Noisy-Key Cryptography

Biometric recognition provides a practical uniqueness signal but introduces privacy and compromise risks if raw data or stable templates are exposed. Jain, Ross, and Prabhakar provide standard background on biometric recognition performance (e.g., FAR/FRR tradeoffs) [11]. Presentation-attack detection is a separate axis of evaluation; ISO/IEC 30107-3 defines testing and reporting for biometric presentation attacks [12].

To bind cryptographic material to noisy biometrics, PoBU can draw on fuzzy extractor style primitives. Dodis et al. formalize how to derive stable cryptographic material from noisy sources and analyze the security and leakage trade-offs of helper data [13]. For long-lived deployments, ISO/IEC 24745 provides a standardization target for biometric information protection, including confidentiality, unlinkability across applications, and revocability/renewability properties [14].

3.5 Privacy-Preserving and Non-transferable Credentials

Anonymous credential systems provide relevant building blocks for the tension between privacy and accountability in PoBU eligibility checks. Camenisch and Lysyanskaya introduce efficient non-transferable anonymous credentials with optional anonymity revocation [15]. This line of work informs PoBU designs that aim to support revocation and auditability without turning eligibility checks into stable cross-application identifiers.

3.6 Confidential Computing and Remote Attestation

PoBU systems that perform biometric verification off-chain must contend with infrastructure trust: the verifier and uniqueness registry may run on hardware not controlled by the participant. Confidential computing and remote attestation provide a mechanism to reduce reliance on the cloud operator by allowing relying parties to verify what code is running in a protected execution environment; cloud documentation for SEV-SNP attestation is one concrete reference point [16]. Humanode uses this approach as part of its reference deployment boundary (Section 5).

4 Definitions and Notation

This section formalizes PoBU as an *eligibility and baseline-weight interface*. The purpose is (i) to separate the consensus protocol (which consumes eligibility and weight) from the certification mechanism (which produces them), and (ii) to make explicit the parameters that determine what "one human, one eligible account" can mean in a probabilistic, adversarial setting.

4.1 PoBU Realization Classes

PoBU can be realized by different verification and eligibility mechanisms. Two broad classes are sufficient to cover current designs:

- **Credentialed PoBU:** eligibility is proven by a long-lived, reusable attestation (e.g., a credential) with explicit revocation or renewal.
- **Freshness-gated PoBU:** eligibility is granted by a short-lived authorization that expires after a fixed window and must be refreshed via periodic re-authentication.

The definitions below specify a common eligibility interface that both classes satisfy; implementation details (where verification runs, exact proof format, or revocation logic) are left abstract.

4.2 Eligibility as a Scarcity Interface

Consensus systems must decide *who is eligible* to participate in a role (e.g., block production, committee membership, governance) and *how much weight* each eligible participant receives. We model this as an eligibility predicate $\mathsf{Eligible}_t(\cdot)$ and a baseline weight function $w_t(\cdot)$ defined over blockchain accounts at time/state t. Unless stated otherwise, the role is implicit in $\mathsf{Eligible}_t$; when needed we write $\mathsf{Eligible}_t^{(r)}$ to emphasize a particular role r (e.g., *validator*).

Entities. Let $\mathcal{H}$ denote the set of physically distinct humans and let $\mathcal{A}$ denote the set of blockchain accounts (or validator keys). We use t to denote a time/state index (e.g., block height, epoch, or session). For $h \in \mathcal{H}$ and $a \in \mathcal{A}$, we write $\mathsf{Ctrl}(h, a)$ to mean that human h can cause account a to act in PoBU-weighted roles; this abstracts away delegation and organizational structures and is used only to state an intended bound on baseline eligibility.

Verification attestation. A verification attestation is an authorization artifact produced by a verification system after (i) verifying liveness and capture integrity and (ii) checking uniqueness against an enrollment or matching substrate. The attestation is bound to an account (or validator public key) and to a deployment/domain d; it may be a long-lived credential (credentialed PoBU) or a short-lived authorization (freshness-gated PoBU). Its validity is bounded either by explicit expiry/revocation or by a wipe/refresh mechanism at the domain level.

4.3 Eligibility Interface via Uniqueness Notifications

The eligibility interface abstracts over implementation details (biometric modality, where verification runs, attestation format). It is sufficient for protocol-level security reasoning as long as the interface is verifiable by the chain. We write d for a PoBU verification domain or deployment (e.g., a chain or instance) and att_d for an attestation scoped to d.

Lifecycle operations. We model a minimal interface:

- $\mathsf{ProveUniqueness}_d(h, a) \to \mathsf{att}_d$ issues an attestation for human h bound to account a within domain d (the verifier may infer a from proof-of-key-possession rather than take it as an explicit input);
- $\mathsf{NotifyUniqueness}_d(\mathsf{att}_d, a)$ records on-chain eligibility based on a valid attestation;
- $\mathsf{NotifyWipe}_d()$ signals a domain-level refresh that invalidates prior notifications (e.g., periodic re-verification or bulk revocation).

Eligibility can be invalidated by an explicit domain-level wipe $\mathsf{NotifyWipe}_d$ and/or by per-notification lifetimes L (expiration without a global wipe). In credentialed systems, invalidation is typically driven by expiry and revocation policy, while in freshness-gated systems it is driven by short lifetimes and periodic re-authentication.

Enrollment (one-time). Most deployments include an enrollment step that registers a public key (or account) and enforces uniqueness before ongoing proofs are accepted. We keep enrollment abstract because some systems treat it as a separate operation, while others embed it into the first successful $\mathsf{ProveUniqueness}_d$ call; in either case, enrollment is the point where duplicates are rejected.

Eligibility predicate. Let $\mathsf{Eligible}_t(a)$ be the predicate that account a is eligible at time/state t for PoBU-weighted roles (validator eligibility, committee membership, voting weight, etc.). Let $W_d(t)$ denote the most recent time $\leq t$ at which $\mathsf{NotifyWipe}_d$ occurred (or 0 if none), and let $\mathsf{Notified}_d(\mathsf{att}_d, a, t')$ denote that $\mathsf{NotifyUniqueness}_d(\mathsf{att}_d, a)$ was accepted at time t'. Let L denote the maximum lifetime of a uniqueness notification before it must be refreshed (a re-verification period); in credentialed systems L can be interpreted as a long expiry horizon or as a bound induced by revocation policy. In the simplest interface,

$$\mathsf{Eligible}_t(a) := \exists \mathsf{att}_d, t' \text{ such that } \mathsf{Notified}_d(\mathsf{att}_d, a, t') \text{ and } W_d(t) \leq t' \leq t \text{ and } t \leq t' + L.$$

For convenience, define the eligible set

$$E(t) := \{a \in \mathcal{A} : \mathsf{Eligible}_t(a)\}.$$

The transaction layer may permit multiple accounts per human; PoBU constrains only eligibility for PoBU-weighted roles.

4.4 PoBU as a Protocol Property

Definition (PoBU system). A protocol is a PoBU system if there exists an eligibility interface such that, for each human $h \in \mathcal{H}$ and time/state t, the number of concurrently eligible accounts controlled by h is bounded by one except with a small system-defined failure probability. Let $\varepsilon_{\mathrm{U}} \in [0, 1]$ denote that uniqueness-failure parameter. Writing

$$C_{h,t} := \{a \in \mathcal{A} : \mathsf{Eligible}_t(a) \wedge \mathsf{Ctrl}(h, a)\},$$

the intended property is

$$\Pr[|C_{h,t}| > 1] \leq \varepsilon_{\mathrm{U}}. \tag{1}$$

Here $\Pr[\cdot]$ is over the randomness and noise in the verification process (and any protocol randomness), and ε_{U} depends on biometric error rates and operational policy (e.g., operating points, re-authentication cadence, and rate limits) [11, 12].

Why this differs from PoW/PoS. Under PoW/PoS, the weight-bearing resource (hashpower or stake) can be split across keys and pooled across organizations. In this model, the definition above asserts that baseline eligibility is bounded at the *human* level rather than at the *key* level. The bound is probabilistic because uniqueness and liveness verification are probabilistic.

4.5 Baseline Weight and Capture Thresholds

Baseline weight. Let $w_t(a)$ denote the baseline weight assigned to account a by the protocol at time/state t for a PoBU-weighted role. A canonical PoBU design sets $w_t(a) = 1$ for eligible accounts and $w_t(a) = 0$ otherwise. More generally, a PoBU system aims to enforce $0 \leq w_t(a) \leq 1$ for eligible accounts, with the intent that capital cannot multiply baseline weight beyond the number of humans who undergo verification. Any additional, non-baseline weighting (e.g., stake-weighted voting layered on top of eligibility) is outside this definition and must be analyzed separately.

Capture thresholds (baseline weight). Define the total baseline weight

$$W_t := \sum_{a \in E(t)} w_t(a).$$

Under $w_t(a) \leq 1$ we have $W_t \leq |E(t)|$. Let $\theta \in (0, 1]$ be the baseline-weight fraction required to dominate a PoBU-weighted role. Then the adversary must obtain at least $\lceil \theta W_t \rceil$ eligible accounts; under the PoBU property this requires recruiting and maintaining that many distinct humans except with failure probability ε_{U}.

4.6 Threshold Mapping via Representative Examples

PoBU does not, by itself, define a consensus protocol; it defines an eligibility and baseline-weight interface. To interpret security thresholds, we can map human-bounded baseline weight to a representative consensus safety condition.

Example: BFT-style safety threshold. In Byzantine fault tolerant replication, a common safety condition is $n \geq 3f+1$, where n is the number of voting participants and f is the number of Byzantine participants tolerated [17, 18]. In such a model, safety fails if the adversary controls at least $f+1$ participants.

PoBU interpretation. If a PoBU system runs a BFT-style protocol over n eligible accounts (or over a committee selected from $E(t)$), then breaking safety requires controlling at least $f+1$ eligible participants. Under participation-boundedness, this scales with the number of distinct humans the adversary can recruit and maintain, except with uniqueness failures bounded by ε_{U}.

Example: random committee selection. Many systems select smaller committees from a larger eligible set. Let $N_t := |E(t)|$ and let k with $1 \leq k \leq N_t$ be the committee size. If an adversary controls c eligible accounts at time t (with $0 \leq c \leq N_t$), let X denote the number of adversary-controlled seats in the committee. Let $\tau \in (0, 1]$ be a threshold fraction (e.g., $\tau = 1/3$ for many BFT-style safety bounds). The probability that the adversary controls at least τk committee seats is

$$\Pr[X \geq \lceil \tau k \rceil] = \sum_{i=\lceil \tau k \rceil}^{k} \frac{\binom{c}{i}\binom{N_t - c}{k-i}}{\binom{N_t}{k}},$$

where i counts adversarial seats in the committee and X is hypergeometric. Under PoBU, increasing c requires sustaining more distinct humans (up to uniqueness failures bounded by ε_{U}), while under PoS/PoW increasing c can be achieved by acquiring more capital or compute.

4.7 Biometric Verification and Protection

Measurements and templates. Biometric measurements are noisy [11]. We distinguish a sample (raw capture at time t), a reference (enrollment-time representation used for later matching), and a protected reference (a transformed reference intended to reduce privacy and security risks).

Recognition and error rates (minimal model). Let $\mathsf{Match}_\tau(\cdot)$ denote a biometric decision rule at threshold τ. Standard performance measures include the false accept rate and false reject rate [11]:

$$\mathrm{FAR}(\tau) := \Pr[\mathsf{Match}_\tau(\text{impostor}) = \mathsf{accept}], \tag{2}$$

$$\mathrm{FRR}(\tau) := \Pr[\mathsf{Match}_\tau(\text{genuine}) = \mathsf{reject}]. \tag{3}$$

In PoBU, these rates influence both (i) exclusion of honest users (availability and accessibility) and (ii) uniqueness failures (duplicate issuance/eligibility). The uniqueness check is typically a de-duplication decision (a one-to-many search) rather than a single one-to-one match, so FAR/FRR should be treated as shorthand for the deployed system's false match and false non-match behavior under its operating point. Deployment therefore requires disclosing an operating point and monitoring drift.

Information protection goals. ISO/IEC 24745 articulates privacy requirements for biometric information protection, including confidentiality and unlinkability goals for stored biometric references across applications and databases [14]. Presentation-attack detection (PAD) is a separate axis; ISO/IEC 30107-3 defines testing and reporting for biometric presentation attacks [12].

Noisy-key cryptography. To bind cryptographic material to noisy biometrics, PoBU systems can draw on fuzzy extractor style primitives. Dodis et al. formalize how to derive stable cryptographic material from noisy sources and analyze the security and leakage trade-offs of helper data [13].

4.8 Failure Parameters and Measurable Quantities

PoBU is inherently probabilistic: eligibility is derived from biometric verification and operational policy, and the system must attach explicit failure parameters to its "at most one" claim.

Eligibility lifetime. Recall L, the maximum lifetime of a uniqueness notification before it must be refreshed. Let q denote an upper bound on the number of verification attempts a single physical human can make within one lifetime window (enforced by rate limits, costs, or operational controls).

Uniqueness failure (ε_{U}). Let p_{dup} denote the per-attempt probability that a human with an existing valid eligibility obtains an additional eligible account due to a uniqueness-check failure. Under a simple worst-case union bound (a conservative bound),

$$\Pr[|C_{h,t}| > 1] \leq \min\{1,\ q \cdot p_{\mathrm{dup}}\} =: \varepsilon_{\mathrm{U}}. \tag{4}$$

A reviewer-grade evaluation should either measure or upper-bound p_{dup} and disclose the operational controls that justify q.

Transferability (p_{rent}). Let p_{rent} represent the effective probability that an eligible account can be rented or transferred within a lifetime window at a cost low enough to scale adversarially (Section 6). This parameter highlights a key contrast with PoS: stake is directly transferable and delegable, so vote/validator markets are intrinsic. PoBU does not require $p_{\mathrm{rent}} = 0$; it requires transferability to be treated as a first-class security parameter with explicit mitigations (e.g., short lifetimes, re-verification) and clear disclosure.

Publicly measurable (chain-derived) quantities. This paper focuses on chain-derived measurements over on-chain consensus keys: validator-set size and churn, and concentration over block author keys (top-k, HHI, Gini), which are reproducible from public RPC data (Section 7). These metrics do not directly measure ε_{U} or biometric/PAD error rates, which are identity-layer parameters, and they need not equal the full eligible-human count if the protocol distinguishes *eligible humans* from *active validators*.

5 Humanode PoBU System Overview

This section specifies the publicly documented boundary of Humanode as a reference PoBU deployment at the level needed for threat modeling and empirical interpretation. The aim is not to define PoBU as "Humanode-specific", but to make explicit (i) what is enforced off-chain (capture, liveness, uniqueness) and (ii) what reaches the chain (authentication state and eligibility) so later sections can separate protocol claims from implementation assumptions.

5.1 Public Design Constraints

Humanode states two constraints that delimit the on-chain and infrastructure exposure of its PoBU path:

> "no personally identifiable information (PII) is ever recorded on-chain." [19]

> "protect data from external access - even by system administrators or hypervisors." [19]

These statements imply an architectural split: biometric processing and uniqueness state are kept off-chain, while the chain enforces eligibility for PoBU-weighted roles via an authentication interface.

5.2 Architecture and Interface Mapping

Humanode can be described as implementing the interface in Section 4 with an off-chain verifier that issues short-lived authorizations and an on-chain eligibility predicate driven by active authentications.

Pipeline (high level).

1. **Capture (user device):** the user provides a biometric sample and liveness data.
2. **Sign (validator key):** the full liveness payload (including the biometric sample) is signed locally on the validator node by the validator key to prove key possession.
3. **Verify (robonode service):** liveness is checked, the signature is verified against the matched registry record, and a one-time auth ticket (including a nonce and robonode signature) is issued on success [19, 20].
4. **Authenticate (on-chain):** an authenticate transaction submits the auth ticket; the chain records an active authentication with an expiration time, and the associated validator key becomes eligible for PoBU-gated roles [19, 21].

The auth ticket is one-time use and only exists to create or refresh on-chain eligibility; it is not a reusable credential.

Enrollment and uniqueness (one-time path). Enrollment is a separate operation that registers a validator public key against the biometric registry and rejects duplicates via a matching check; this is where uniqueness is enforced. Subsequent authentications match liveness data against the existing registry and recover the associated public key to validate key possession before issuing an auth ticket.

Mapping to the eligibility interface. At the abstraction level of Section 4, the verifier implements $\mathsf{ProveUniqueness}_d(\cdot)$ by issuing a short-lived attestation (auth ticket) after liveness and uniqueness checks, while the chain implements $\mathsf{NotifyUniqueness}_d(\cdot)$ by accepting the ticket and recording an active authentication with an expiration. Humanode's periodic re-authentication and expiry-driven cleanup serve the role of $\mathsf{NotifyWipe}_d$ at the level of eligibility lifetime [21, 22].

Data placement (publicly claimed). Table 2 summarizes the intended placement of sensitive inputs and the resulting public chain surface.

Location	Primary contents	Publicly observable
User device	Raw capture and liveness interaction	No
Robonode verifier (CVM)	Biometric processing, uniqueness state, auth ticket issuance	No (confidential-computing boundary)
Chain state	Active authentications with an expiry timestamp (`expires_at`) and validator keys	Yes

Table 2. Humanode PoBU boundary at a high level. The chain is intended to expose eligibility state without exposing biometrics or civil identity [19].

5.3 Cryptobiometric Verification Pipeline

Humanode positions PoBU as cryptobiometric verification: a participant proves liveness and uniqueness via an off-chain verification pipeline, and the resulting eligibility is represented on-chain without exposing biometrics or civil identity [19]. In the Humanode reference deployment, the verifier issues one-time auth tickets and the chain records short-lived active authentications; confidentiality and integrity are scoped to the CVM boundary and to a verifiable chain of trust for deployed verifier code [20].

5.4 Confidential Verification Boundary and Attestation

In this architecture, confidentiality and integrity of the verification path depend on (i) the CVM threat model and (ii) the ability for relying parties to validate that the expected verifier code is running. Humanode discusses CVMs and a verifiable trust chain for the deployed verifier [20]. Remote attestation is one concrete mechanism to support such validation in SEV-SNP deployments [16].

Scope and trust assumptions. This paper does not attempt to evaluate confidential-computing hardware security. Instead, it treats CVMs and attestation as explicit trust assumptions and focuses on the protocol-level consequences: what the chain can enforce, and which failure modes shift to the off-chain verification layer.

5.5 On-chain Interface, Lifecycle, and Observables

On-chain interface. This paper treats the exact auth ticket format and the on-chain encoding of eligibility as an implementation detail, and uses only the externally observable interface: (i) eligibility can be checked from chain state via active authentications and expiry, and (ii) PII/biometrics are not recorded on-chain [19]. The black-box view is sufficient for the threat model in Section 6 and the chain-derived measurements in Section 7.

Lifecycle and operations. Humanode documents an authentication lifecycle with periodic re-authentication (used as a security control against long-lived eligibility abuse) and operational procedures such as biometric server maintenance windows and CVM rotations [21, 22]. These details affect availability and exclusion risk, and motivate correlating public identity-layer operational events with on-chain observables (Section 7).

Chain-derived observables. Because biometrics and uniqueness state are off-chain, the most reproducible public evidence for PoBU system behavior is chain-derived: validator-set size and churn, and concentration over block author keys (top-k, HHI, Gini). We report these measurements using direct Substrate RPC extraction in Section 7.

6 Security Analysis

We analyze the PoBU eligibility interface (Section 4) under the Humanode reference deployment boundary (Section 5). The security question is not whether collusion is possible in principle, but whether an adversary can *scale* baseline influence beyond the number of humans it can recruit and maintain, or can impose censorship/monitoring by compromising the verification layer.

6.1 Scope and Evidence Policy

This section is a threat analysis, not a security proof. We treat statements about Humanode's implementation as facts only when supported by cited public documentation. We do not measure biometric model quality (e.g., PAD error rates) or confidential-computing hardware security; instead, we treat them as explicit parameters and assumptions (Section 4, Section 8).

6.2 Security Objectives at the Interface Level

Let eligibility be determined by $\mathsf{Eligible}_t(\cdot)$ (Section 4), derived from verifiable attestations or active authentications. A PoBU system aims to:

- **Bound multi-eligibility:** keep the probability that one human controls multiple concurrently eligible accounts small (captured by ε_{U}).
- **Limit transferability:** treat eligibility rental/coercion as a first-class risk parameter (p_{rent}), and design lifecycle rules that make it costly to scale.
- **Preserve privacy:** prevent biometric samples/templates and civil identity from reaching chain state, and minimize linkability of eligibility proofs.
- **Support recovery:** support refresh and invalidation (lifetimes, wipe/refresh, and/or revocation) so compromise does not permanently corrupt the eligibility set.
- **Maintain availability and avoid exclusion:** ensure enrollment and periodic re-verification do not become a practical censorship bottleneck.

Freshness-gated systems trade long-lived credential risk for periodic liveness requirements. This shifts some risk from transferability to availability: short lifetimes reduce long-term abuse but make re-auth outages and verification-layer denial-of-service more consequential.

Conversely, credentialed systems reduce dependence on frequent re-authentication but increase the importance of credential protection, revocation correctness, and long-term key compromise handling.

6.3 Adversary Model and Assumptions

We consider the following adversaries (not mutually exclusive):

- A1 (remote Sybil adversary): attempts to obtain many eligible accounts via repeated enrollment, automation, or synthetic media.
- A2 (market adversary): rents, buys, or coerces humans to obtain/maintain eligibility (industrial-scale *human farming*).

- A3 (infrastructure adversary): controls cloud/hypervisor/host OS; seeks to read biometric data or tamper with verifier logic.
- A4 (issuer/verifier insider): compromises verifier code, issuance policy, or keys; may attempt mass issuance or selective censorship.
- A5 (observer/adversarial application): links repeated eligibility checks to track participants across time and applications.

Assumptions. We assume standard cryptographic primitives are correctly implemented and that the confidential-computing substrate provides the intended isolation and attestation properties within its documented model [16]. We do not assume perfect biometrics: uniqueness and liveness are probabilistic and must be parameterized (Section 4) [11, 12].

6.4 Attack Surface Decomposition

Under the reference boundary in Section 5, attacks fall into four surfaces: (i) capture and liveness at the client, (ii) verification and issuance inside CVMs, (iii) authentication/eligibility updates on-chain, and (iv) on-chain usage (validation and application-layer eligibility checks).

Table 3 summarizes how key attack classes map to the interface-level parameters introduced in Section 4.

Attack class	Primary parameter(s) affected	Primary surface(s)
Presentation attacks / liveness bypass	multi-eligibility risk (via liveness failures)	client capture; verifier
Duplicate issuance / uniqueness failures	multi-eligibility risk (uniqueness failures)	verifier/registry
Authorization theft / account takeover	eligibility misuse within a lifetime window	client; on-chain submission
Eligibility rental/coercion	transferability and coercion markets	market; lifecycle
Verifier compromise / mass issuance	multi-eligibility risk, censorship, and availability	verifier; issuance keys
Privacy leakage / linkability	cross-context identifiability and surveillance	application; chain usage
DoS / exclusion at verification layer	availability and participation skew	verifier; operations

Table 3. Threat mapping to PoBU parameters (Section 4) and to the reference deployment surfaces (Section 5).

6.5 Presentation Attacks and Liveness Bypass

Threat (A1): an adversary attempts to bypass liveness checks using replay, masks, synthetic media, or weak capture UX. ISO/IEC 30107-3 defines and evaluates presentation-attack detection (PAD) in terms of attacks at the capture device and standardized testing and reporting [12]. A successful liveness bypass increases the effective probability of duplicate issuance and therefore increases ε_{U}.

Mitigation direction (not measured here): disclose PAD evaluation methodology and aggregate performance (e.g., APCER/BPCER), and re-evaluate periodically as synthesis improves. This paper does not claim any PAD performance; it only identifies PAD as a parameter that must be reported for a reviewer-grade PoBU evaluation.

6.6 Duplicate Issuance and Uniqueness Failures

Threat (A1/A4): an adversary obtains multiple valid eligibilities despite uniqueness checks. Causes include threshold misconfiguration, sensor drift, data partitioning, and implementation bugs. Any biometric system has non-zero FAR/FRR [11]; therefore, a PoBU system must treat multi-eligibility as probabilistic and track a bound or estimate for ε_{U} under the deployed operating point (Section 4).

Mitigation direction (not measured here): publish aggregate bounds (or estimates) on multi-eligibility probability together with the operational controls that justify the attempt bound q (rate limits, costs, policy), and disclose incident handling for suspected duplicates.

6.7 Verifier Compromise, Censorship, and Attestation Failures

Threat (A3/A4): the verifier is modified to leak biometrics, relax uniqueness checks, selectively censor verification, or mass-issue authorizations. Humanode positions verification as running inside CVMs and discusses a verifiable trust chain for the deployed verifier [20]; remote attestation is one concrete mechanism to support such verification in SEV-SNP deployments [16].

Mitigation direction (implementation-dependent): require relying parties (users, auditors, or chain governance) to validate attestation evidence and deployed measurements; use reproducible builds and controlled rollouts; and reduce single-operator concentration at the verification layer (multi-issuer/multi-operator), which remains an open problem (Section 8).

6.8 Credential Theft, Account Takeover, and Key Compromise

Threat (A1/A2): attackers steal signing keys, steal eligibility artifacts, or compromise the binding step. Humanode documents short-lived eligibility through periodic re-authentication:

> "you need to reauthenticate every 7 days (168 hours) at the same time." [21]

Short eligibility lifetimes can reduce the window in which stolen eligibility remains useful (reducing effective abuse within L), but they couple security to the availability of re-verification and can increase exclusion risk (Section 8).

6.9 Credential Rental, Coercion, and Human Markets

Threat (A2): an attacker pays humans to maintain eligibility and vote/validate in the attacker's interest. This does not violate the PoBU uniqueness property, but it increases the effective transferability parameter p_{rent} and becomes a scaling path for capture.

Mitigation direction (not measured here): document which lifecycle rules (renewal cadence, proof freshness, device binding) create friction against industrial-scale rental and quantify the trade-off between friction and accessibility/exclusion.

6.10 Availability and Exclusion Attacks

Threat (A3/A4): denial of service at enrollment or renewal creates de facto censorship, reduces participation, and can shift baseline control toward actors with privileged access to the verification pipeline. Humanode publicly announces at least some identity-layer operational events (e.g., biometric server maintenance windows and CVM rotations) [22]. We correlate such publicly disclosed events with on-chain observables (validator-set size, churn, and concentration) in Section 7, and avoid causal claims unless the incident dataset is sufficiently complete.

6.11 Linkability and Privacy Leakage

Threat (A5): repeated eligibility checks become a tracking primitive. Even if the design intent is that biometrics remain off-chain and protected within the confidential verification boundary [19, 20], stable on-chain accounts are linkable by default. A PoBU system must treat privacy as more than "no biometrics on-chain": it must minimize stable eligibility identifiers and provide guidance for application-layer usage that avoids turning eligibility checks into cross-context identifiers.

Mitigation direction (not measured here): specify an eligibility interface that supports verification without revealing a stable proof handle across applications, and align eligibility lifecycle design with biometric information protection goals (e.g., unlinkability and revocability) [14].

6.12 Governance

This paper does not report governance measurements; see Section 8 for scope and future work.

7 Empirical Observations from a Live PoBU Chain

Humanode has operated as a live network long enough to support measurement-driven evaluation of participation breadth and operational stability under a PoBU reference deployment boundary. Because PoBU verification (uniqueness and liveness) is enforced off-chain (Section 5), this section restricts itself to *publicly reproducible, chain-derived observables* and treats biometric error rates and uniqueness failures as explicit parameters rather than measured quantities (Section 4).

7.1 Scope: What These Measurements Show and Do Not Show

We measure two necessary (but not sufficient) signals for broad participation under a PoBU-gated validator set:

- **Validator-set breadth and dynamics:** the size of the active validator set over time and its churn from session to session.
- **Block-production concentration:** how concentrated block authorship is across validator keys, computed from sampled blocks.

These metrics are computed over *on-chain keys* (validator identities as seen by the chain). They do not directly measure unique humans, off-chain enrollment counts, or the identity-layer uniqueness error rate ε_{U}.

7.2 Why These Measurements Matter

PoBU's central claim is not that biometrics are "on-chain," but that the *unit of eligibility* for consensus is bounded by unique humans (Section 4). That claim has two necessary observable consequences at the chain boundary:

- **Eligibility should not collapse to a small key set.** If the verification/eligibility layer effectively gates participation, we should observe a large active validator set, with churn that reflects many independent operators rather than a stable cartel.
- **Control should not concentrate in block production.** Even with many validators, control can concentrate if a small subset of keys produces most blocks (e.g., due to delegation, key reuse, cartel coordination, or infrastructure bottlenecks). Low concentration is therefore a necessary (but not sufficient) condition for "human-bounded" participation to be credible at the protocol boundary.

This section therefore focuses on (i) validator-set size and dynamics and (ii) block-author concentration. These observables do not prove "one human, one validator," but they do test whether the *visible* consensus control surface is already dominated by a small key oligopoly. In freshness-gated PoBU systems, eligibility is short-lived and must be refreshed; validator-set size and churn therefore reflect both participation and re-authentication cadence.

7.3 What Would Indicate Failure

Empirically, PoW and PoS systems often exhibit concentration pressures driven by industrial coordination or capital compounding. For PoW, mining pools and industrial organization can make decentralization a technological possibility without guaranteeing it as an economic outcome [5]. For PoS, compounding effects can create inequitable stake growth even under honest behavior [6]. A PoBU system does not eliminate collusion, but it aims to change the marginal cost of scaling control from "acquire more capital" to "coordinate more unique humans."

In our metrics, two patterns would be immediate warning signs for PoBU's intended participation bound:

- **Validator-set collapse:** a sustained drop in $|S_t|$ toward a small core, or near-zero churn over long windows, suggesting barriers to entry or centralized issuance.
- **High author concentration:** high top-k shares and high HHI, indicating that a small number of keys dominate block production despite the claimed eligibility model.

7.4 Data Sources

We use public chain data accessed via WebSocket RPC; endpoints are published in Humanode documentation [23].

7.5 Time Windows

We evaluate multiple windows to reduce cherry-picking and to separate recent behavior from longer-horizon regimes:

- W1 (recent): last 90 days;
- W2 (mid-term): last 365 days;
- W3 (long horizon): earliest non-zero on-chain timestamp (post-genesis) $\rightarrow$ present.

For each window we report explicit start/end timestamps and block ranges derived from the chain's on-chain timestamps. Timestamps are reported in ISO 8601; the suffix `Z` denotes UTC.

7.6 Metrics

Validator set size and churn. Let S_t denote the active validator set in session t (i.e., validator keys with active authentication in that session). We report:

- $|S_t|$ (active validator count per session);
- entries $|S_t \setminus S_{t-1}|$ and exits $|S_{t-1} \setminus S_t|$;
- set overlap (Jaccard), and a symmetric churn rate (defined below).

Block production distribution. Over a sampled set of blocks in a window, we report:

- top-k author share (top 1/5/10);
- Herfindahl–Hirschman Index (HHI);
- Gini coefficient over author block counts.

Operational event overlay. We overlay publicly disclosed identity-layer operational events (e.g., biometric server maintenance windows and CVM rotations) onto concentration time series to look for visible discontinuities [22]. This is an exploratory check, not causal attribution.

7.7 Metric Definitions

Let B denote the sampled set of blocks in a window and let $N := |B|$ (we assume $N > 0$). For each author i, let c_i denote the number of sampled blocks authored by i and let $s_i := c_i/N$ denote its share. Let authors be sorted so that $c_1 \geq c_2 \geq \cdots \geq c_m$ where m is the number of distinct observed authors (so $m \geq 1$).

Top-k share. For a chosen cutoff k with $1 \leq k \leq m$ (e.g., $k \in \{1, 5, 10\}$), define:

$$\mathrm{TopK}(k) := \sum_{i=1}^{k} s_i. \tag{5}$$

HHI.

$$\mathrm{HHI} := \sum_{i=1}^{m} s_i^2. \tag{6}$$

Note that $\sum_{i=1}^{m} s_i = 1$. As an interpretive aid, 1/HHI can be read as an "effective number" of equally-sized authors that would yield the same concentration.

Gini coefficient (over author block counts). Let $\mu := \frac{1}{m}\sum_{i=1}^{m} c_i$ denote the mean count. We compute:

$$\mathrm{Gini} := \frac{1}{2m^2\mu} \sum_{i=1}^{m} \sum_{j=1}^{m} |c_i - c_j|. \tag{7}$$

Validator-set overlap and churn. Let S_t denote the validator set in session t. The Jaccard overlap is:

$$J(S_{t-1}, S_t) := \frac{|S_{t-1} \cap S_t|}{|S_{t-1} \cup S_t|}. \tag{8}$$

We report entries $|S_t \setminus S_{t-1}|$ and exits $|S_{t-1} \setminus S_t|$, and define a symmetric churn rate:

$$\mathrm{Churn}(S_{t-1}, S_t) := \frac{|S_t \setminus S_{t-1}| + |S_{t-1} \setminus S_t|}{\frac{1}{2}(|S_t| + |S_{t-1}|)}. \tag{9}$$

We interpret this quantity only when $|S_t| + |S_{t-1}| > 0$ (which holds in all measured windows).

7.8 How to Interpret the Tables and Plots

All metrics in this section are computed over *on-chain keys* (block authors and validator keys), not over verified unique humans.

- **Top-1 / Top-10 share:** the fraction of sampled blocks attributed to the most prolific author (Top-1) or the top 10 authors (Top-10). Higher values mean block production is concentrated into fewer keys; lower values mean production is spread.
- **HHI:** a concentration index over author shares. It equals 1 under monopoly and decreases as production spreads across more keys; 1/HHI is a convenient "effective number" of equally-sized authors.
- **Gini:** inequality over per-author block counts. It is 0 under perfect equality and increases as the distribution becomes more unequal. Because we compute Gini on sampled blocks, its value can change with the sampling step and the window length even if top-k shares remain small.
- **Entries/exits, Jaccard, churn:** session-to-session validator set dynamics. High Jaccard (near 1) indicates that most keys persist across sessions (a stable set), while higher churn indicates more turnover. Stability can be compatible with broad participation if the set is large, but can also indicate barriers to entry if the set is small; interpret dynamics together with set size.

7.9 Extraction Methodology with Direct RPC

All extraction is done via WebSocket RPC against public endpoints, using Substrate API tooling. The repository includes scripts under `metrics/` that:

- convert timestamps to block ranges via binary search over on-chain timestamps;
- extract block authors across a block range;
- extract session validator sets for churn analysis;
- compute concentration metrics (top-k, HHI, Gini) from the extracted datasets.

Reproducibility and artifacts. The exact RPC endpoint, tool versions, commands, datasets, and checksums used to produce the tables and figures in this section are recorded under `metrics/runs/`; scripts are documented in `metrics/README.md`. The primary artifacts are `authors_*.csv` (sampled block authors), `sessions_*.jsonl` (session validator sets), and `summary_*.json` (derived window statistics); these files are sufficient to regenerate the tables and figures in this section without relying on explorer-specific APIs.

Sampling and limitations. Block-author attribution is derived from headers via Substrate API tooling. Author metrics are computed from sampled blocks with an explicitly disclosed step size (to keep RPC extraction tractable). Fixed-step sampling can introduce aliasing; we therefore (i) disclose step sizes and sample counts, (ii) repeat key summaries under multiple step sizes as a sensitivity check, and (iii) avoid interpreting these results as a full census.

7.10 Results: W1 Recent Window

W1 spans `[14831011, 16128021]` (2025-10-02 00:00:00Z to 2025-12-31 00:00:00Z, via on-chain timestamps). For author metrics we sample every 25 blocks (n = 51,881 sampled blocks). Session metrics are computed from session snapshots at session boundaries.

Table 4 summarizes the W1 results.

Metric	Value
Unique block authors (sample)	405
Top-1 share	0.00465
Top-5 share	0.02299
Top-10 share	0.04514
HHI	0.00343
Gini (author block counts)	0.348
Avg. entries per session	1.30
Avg. exits per session	1.60
Avg. Jaccard overlap (sessions)	0.989
Avg. validators per session	258.4 (min 169, max 367)
Avg. churn rate per session	0.0114

Table 4. W1 summary statistics from direct-RPC extraction (author metrics sampled every 25 blocks; session metrics from session snapshots).

W1 interpretation. W1 shows hundreds of distinct block authors in the sample and low block-production concentration (e.g., the top-10 authors account for only a few percent of sampled blocks). The corresponding HHI implies an effective author count on the order of 1/HHI rather than a small oligopoly. On the validator-set side, the per-session counts remain in the hundreds with non-zero churn, indicating that (at least at the key level) participation is not confined to a small static set.

Two plots provide time-resolved context:

- Figure 1 shows the active validator set size per session. Each point is $|S_t|$ at a session boundary, so the plot indicates whether the validator set remains broad or collapses toward a small core.
- Figure 2 shows daily concentration metrics (top-1 share, top-10 share, and Gini) computed from sampled blocks bucketed by UTC day. The y-axis values are fractions (for top-k) or a coefficient (for Gini); lower values correspond to less concentration over block authors.

Bootstrap variability. As a stability check for the sampled author metrics, we bootstrap-resample the sampled block-author labels with replacement (2000 iterations; sample size fixed at $n = 51881$). The resulting percentile intervals (95%) are reported in Appendix A (Table A1). This is a sampling-variability check over the sampled blocks; it does not account for systematic effects from fixed-step sampling and should not be interpreted as a population confidence interval.

Sensitivity to sampling. We include a step-size sensitivity check for W1 (step=25/50/100) in Appendix A.

Event record. For event overlays in Figure 2, we maintain a small, auditable event log in `metrics/data/events.csv` (date, label, source key) and generate plot tick marks from that file. In the current dataset, the only W1 ops announcement recorded is dated 2025-10-09 [22]. The absence of a visible discontinuity around this marker should not be over-interpreted: the event log is not guaranteed to be complete, and some effects may not be visible at daily aggregation.

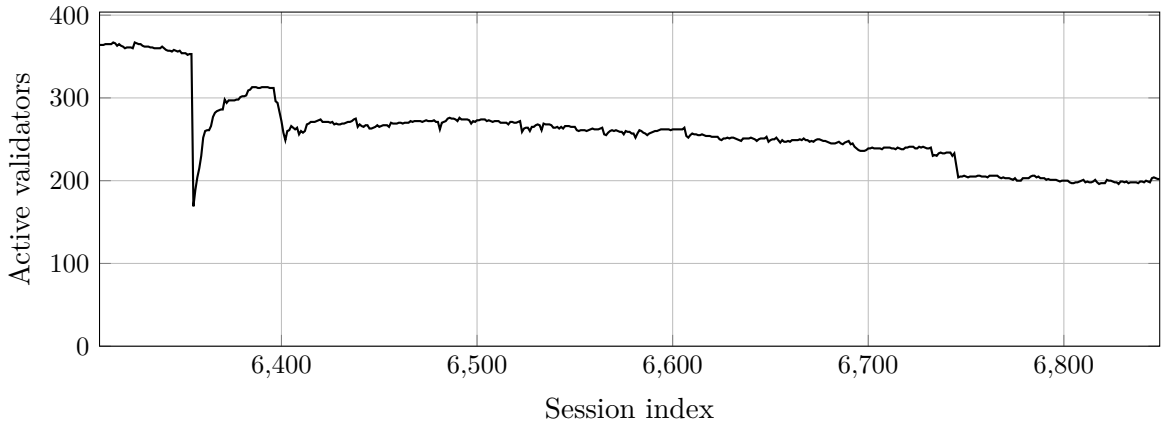

Figure 1. W1 active validator set size per session (direct RPC extraction). Each point is a session snapshot; the x-axis is the session index and the y-axis is the number of active validators.

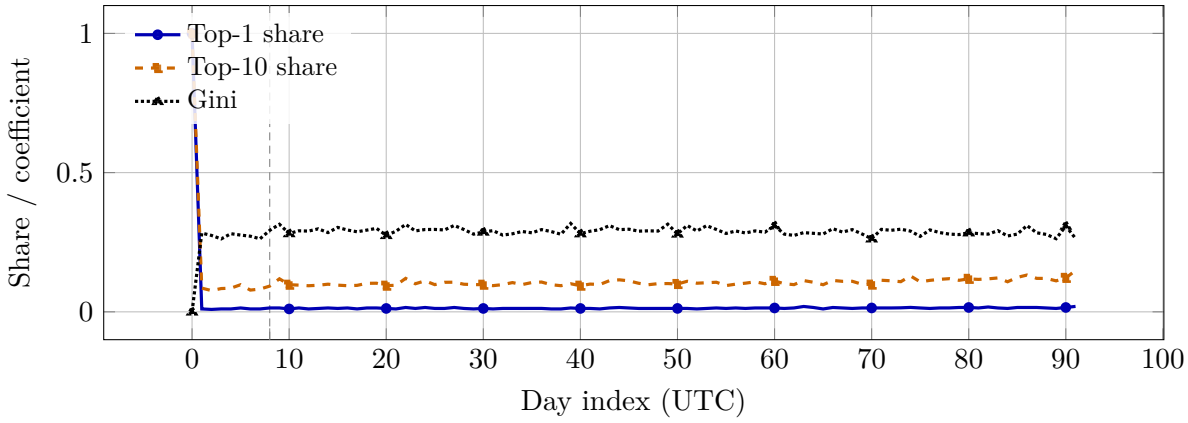

Figure 2. W1 daily block-author concentration metrics (sampled every 25 blocks; UTC day buckets). Dashed vertical lines denote operational events recorded in `metrics/data/events.csv`.

7.11 Results: W2 Mid-term Window

W2 spans `[10900503, 16128883]` (2024-12-31 00:00:00Z to 2025-12-31 00:00:00Z). For author metrics we sample every 1000 blocks ($n = 5{,}229$ sampled blocks). Session metrics are computed from all session snapshots in the window.

Table 5 summarizes the W2 results.

Metric	Value
Unique block authors (sample)	1717
Top-1 share	0.00306
Top-5 share	0.0142
Top-10 share	0.0273
HHI	0.00105
Gini (author block counts)	0.431
Avg. entries per session	5.27
Avg. exits per session	5.87
Avg. Jaccard overlap (sessions)	0.989
Avg. validators per session	1044.6 (min 169, max 1767)
Avg. churn rate per session	0.0112

Table 5. W2 summary statistics (author metrics sampled every 1000 blocks; session metrics from session snapshots).

W2 interpretation. W2 reflects a much larger validator set on average, and the author sample covers many more distinct authors. Top-k shares and HHI remain low, indicating that block production is not dominated by a small subset of keys at the scale visible in this sample. The higher Gini coefficient in W2 relative to W1 is consistent with a long window plus coarse author sampling: with many authors appearing only a few times in the sample, the distribution of sampled counts becomes more unequal even when top-k shares remain small.

To help interpret these aggregates, Figure 3 plots the validator set size per session over W2 (downsampled for readability in the plot). A weekly concentration time series for W2 is provided as supplementary material in Appendix A (Figure A1); the weekly bucketing reduces noise introduced by the coarse author sampling step.

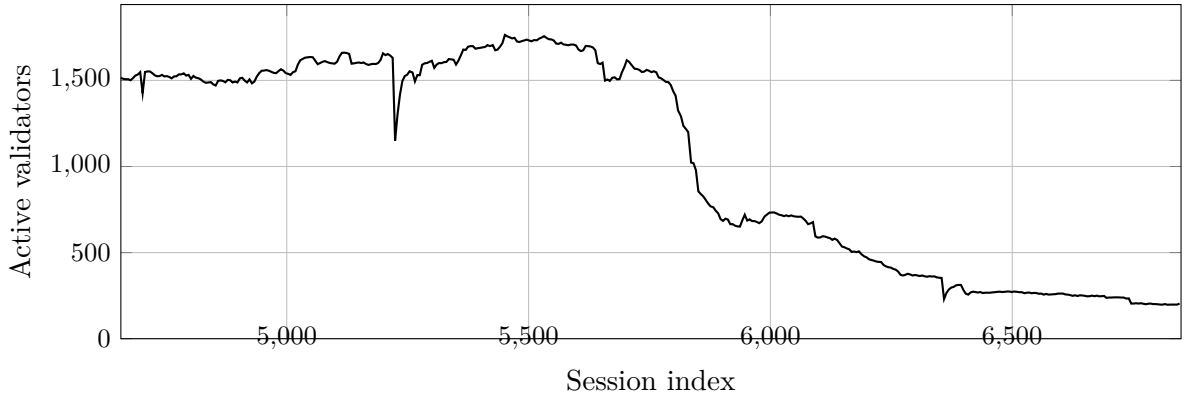

Figure 3. W2 active validator set size per session (direct RPC extraction). For readability, the plot is downsampled (every 5th point) while statistics are computed from the full session series.

7.12 Results: W3 Long-Horizon Snapshot

W3 spans `[1, 16128883]` (from the first non-zero on-chain timestamp, 2022-11-15, to 2025-12-31 00:00:00Z). For author metrics we sample every 5000 blocks ($n = 3{,}226$ sampled blocks). We report W3 as a snapshot because the long horizon mixes distinct protocol and validator-set regimes; interpreting time series without explicit regime segmentation is misleading (Appendix A).

Table 6 summarizes the W3 snapshot results.

Metric	Value
Unique block authors (sample)	1487
Top-1 share	0.00403
Top-5 share	0.0195
Top-10 share	0.0363
HHI	0.00120
Gini (author block counts)	0.392
Avg. entries per session	5.39
Avg. exits per session	5.36
Avg. Jaccard overlap (sessions)	0.982
Avg. validators per session	714.4 (min 3, max 1767)
Avg. churn rate per session	0.0183

Table 6. W3 snapshot statistics (author metrics sampled every 5000 blocks; session metrics from session snapshots).

W3 interpretation. W3 aggregates across multiple operational regimes (early network, growth phases, and later behavior). We therefore avoid claiming a single time-horizon trend. Nonetheless, the snapshot remains informative as a coarse consistency check: it does not exhibit extreme author concentration that would suggest long-term collapse into a small key oligopoly.

To provide visual context for the long horizon, Figure 4 plots the validator set size per session over W3. This figure is not used to argue a single-regime trend; it is included to show the scale of validator-set expansion and regime changes over the full observed history.

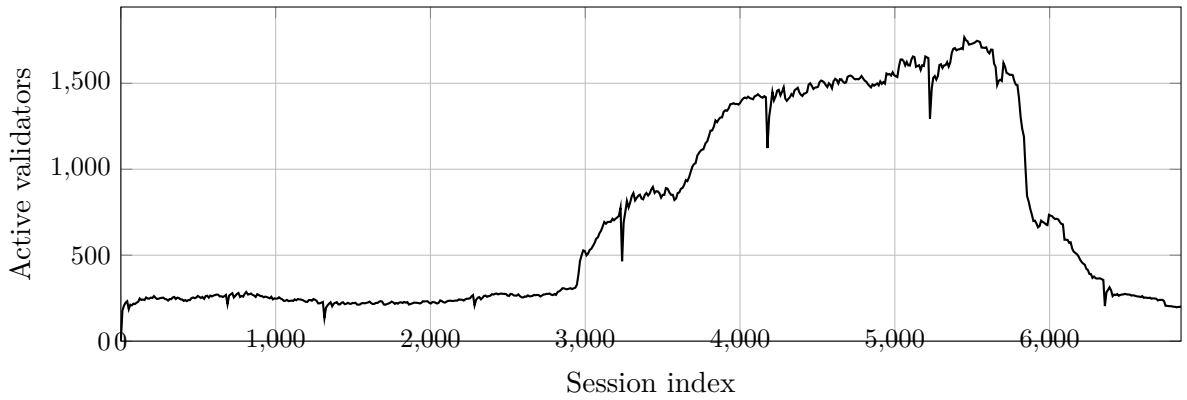

Figure 4. W3 active validator set size per session (direct RPC extraction). For readability, the plot is downsampled (every 10th point) while statistics are computed from the full session series.

8 Limitations and Future Work

PoBU is an eligibility primitive, not a guarantee of benign behavior. This paper also evaluates PoBU using chain-derived measurements, which are necessary for reproducibility but do not directly observe off-chain identity-layer error rates. We therefore state limitations explicitly and distinguish (i) limitations of PoBU as a consensus class from (ii) limitations of this paper's evidence.

8.1 Collusion, Coercion, and Markets Remain Possible

PoBU does not prevent cartels of real humans from coordinating, nor does it prevent bribery and coercion markets from emerging. The contribution is narrower: it shifts the scaling variable for *baseline* control from pure capital compounding toward the operational costs of recruiting and maintaining human participation (Section 4).

Participation asymmetries. Even with one-human-bounded baseline eligibility, participation can still concentrate among well-resourced actors due to attention, coordination, and incentive asymmetries; token-weighted overlays can also reintroduce plutocratic dynamics [7].

8.2 Certification Remains a Structural Dependency

Strong Sybil resistance still requires a certification path or a scarce-resource test. PoBU chooses a certification path and therefore inherits certification risks: issuer concentration, policy capture, selective denial, and insider compromise. Minimizing and decentralizing this trust surface remains a core open problem [1, 4].

8.3 Biometric Systems are Probabilistic and May Be Non-uniform

Uniqueness and liveness checks necessarily operate at non-zero error rates (false accepts and false rejects) and must choose operating points [11]. Presentation-attack detection (PAD) adds an additional evaluation axis with standardized testing and reporting requirements [12]. These factors influence the effective multi-eligibility probability (Section 4) and the exclusion rate for honest users.

Moreover, biometric performance can vary across demographic groups and acquisition conditions; a PoBU deployment should therefore treat performance as context-dependent and publish evaluation summaries appropriate for its target population and sensor set [24].

8.4 Privacy is Not Automatic

Even if biometric data is kept off-chain, stable on-chain accounts are linkable by default and eligibility checks can become a tracking primitive at the application layer. A PoBU system must treat privacy as a first-class requirement and align with biometric information protection goals such as confidentiality, unlinkability across contexts, and lifecycle management (revocability/renewability) [14].

8.5 Confidential-Computing Trust is Not Absolute

Confidential computing reduces reliance on infrastructure operators, but it introduces reliance on hardware roots of trust and their firmware supply chains. Remote attestation can improve auditability of what code is running in a protected environment, but cannot eliminate all side-channel, implementation, or supply-chain risks [16]. These risks must be treated as explicit assumptions and mitigated operationally (diversity, patch cadence, incident response), not assumed away.

8.6 Evidence Limitations of This Paper

This paper reports only chain-derived measurements (Section 7). As a consequence:

- concentration and churn are computed over validator keys and do not directly measure unique humans or identity-layer failures;
- block-author statistics are computed from sampled blocks with disclosed step sizes and may be sensitive to sampling aliasing;
- the operational event overlay uses a small, auditable event log and is not guaranteed to include all relevant identity-layer incidents.

These limitations are not unique to Humanode; they reflect the general problem that PoBU's critical security parameters partly live off-chain.

8.7 Future Work

The most important next steps for PoBU as a consensus class are:

- Publish a precise protocol interface for prove-uniqueness, on-chain notification, eligibility checks, wipe/refresh, and any renewal/revocation mechanisms (beyond the black-box interface specified here).
- Publish safe, aggregate identity-layer measurements (e.g., enrollment/renewal success rates, PAD evaluation summaries, and duplicate-attempt handling) and map them to bounds on ε_{U} under disclosed operational controls (Section 4).

- Explore multi-issuer decentralization and governance mechanisms for adding/removing issuers and verifier operators without creating unilateral control or regional censorship.
- Strengthen privacy-preserving eligibility proofs (e.g., selective disclosure / anonymous credentials) while preserving revocation and audit needs [15].
- Upgrade empirics from snapshots to regime-aware analysis (explicit segmentation across protocol epochs) and consider randomized sampling or full-census author extraction where computationally feasible.

9 Conclusion

This paper defines Proof-of-Biometric-Uniqueness (PoBU) as a scarcity primitive for permissionless consensus: baseline eligibility for PoBU-weighted roles is bounded by verified unique humans rather than by computation or capital. PoBU does not eliminate collusion or coercion; it changes the scaling variable for baseline capture and shifts the dominant trust assumptions to certification integrity, eligibility lifecycle, and privacy.

We formalize PoBU as an eligibility and baseline-weight interface with explicit failure parameters (Section 4), and we analyze the resulting attack surface with a parameter-driven threat model (Section 6). A central requirement is that privacy must be treated as a first-class property—not merely as "no biometrics on-chain"—and should align with biometric information protection goals such as confidentiality, unlinkability, and lifecycle management [14].

Using Humanode as a reference deployment boundary, we make the off-chain/on-chain split explicit (Section 5) and report initial chain-derived measurements with direct-RPC reproducibility (Section 7) [19, 20]. The next step for PoBU as a consensus class is to publish safe, aggregate identity-layer measurements (e.g., enrollment/renewal success rates, PAD summaries, and duplicate-attempt handling) and connect them to explicit bounds on multi-eligibility probability, while continuing to strengthen issuer decentralization and privacy-preserving eligibility proofs (Section 8).

A Appendix: Supplementary Material

A.1 Humanode Documentation Excerpts

This paper treats Humanode as a reference deployment boundary (Section 5) rather than as the definition of PoBU. To keep the main text focused on the PoBU interface and on-chain-derived evidence, we collect a small set of short documentation excerpts that motivate trust-boundary assumptions.

No operator access into CVMs.

> "We do not allow any operator access into the CVMs – as it would allow for potential data leakage or security compromise." [20]

This motivates a *deployment intent* that reduces routine insider access, while Section 6 still models infrastructure compromise as an adversary capability.

Immutability stance (no in-place code changes).

> "the code in the Humanode CVM is never changed after launch." [20]

This motivates treating verifier updates as redeployments rather than in-place patching, and informs the incident-response discussion in Section 6.

Update posture (redeploy over patch).

> "we'd rather redeploy the whole thing, losing all the private data, than upload new code to a CVM." [20]

This motivates a conservative operational posture as a mitigation against privileged maintenance channels (Section 6).

A.2 Supplementary Empirical Material

This appendix contains additional plots and robustness checks that support Section 7 but are not required for the main narrative.

A.2.1 W2 weekly concentration series

W2 block-author concentration is sampled coarsely (every 1000 blocks), making daily series noisy. We therefore bucket sampled blocks into 7-day windows and plot the resulting weekly time series (Figure A1). The three curves correspond to top-1 share, top-10 share, and the Gini coefficient as defined in Section 7. The purpose of this figure is interpretability: it makes regime shifts and sustained trends visible without overclaiming precision from sparse sampling.

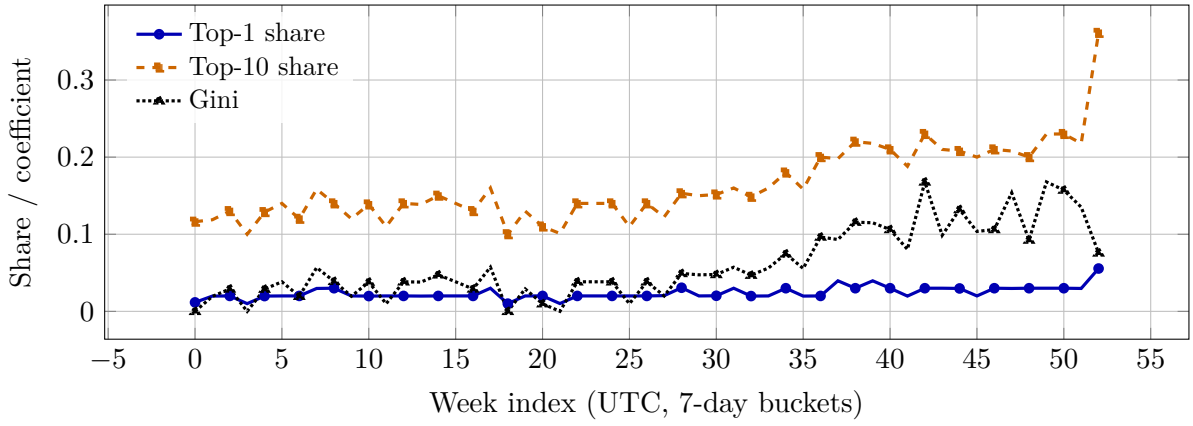

Figure A1. W2 weekly block-author concentration metrics (sampled every 1000 blocks; 7-day buckets).

A.2.2 On-chain `bioauth.activeAuthentications` series

Some identity-lifecycle signals may be exposed on-chain depending on the runtime. We sample the size of the on-chain storage value
`bioauth.activeAuthentications` once per UTC day over W1. Figure A2 plots the resulting series. Each point is the vector length (number of entries) at the sampled block. This is not a measure of unique humans or biometric performance; it is included only as a purely on-chain operational signal that may help interpret changes in validator dynamics when such storage items are present.

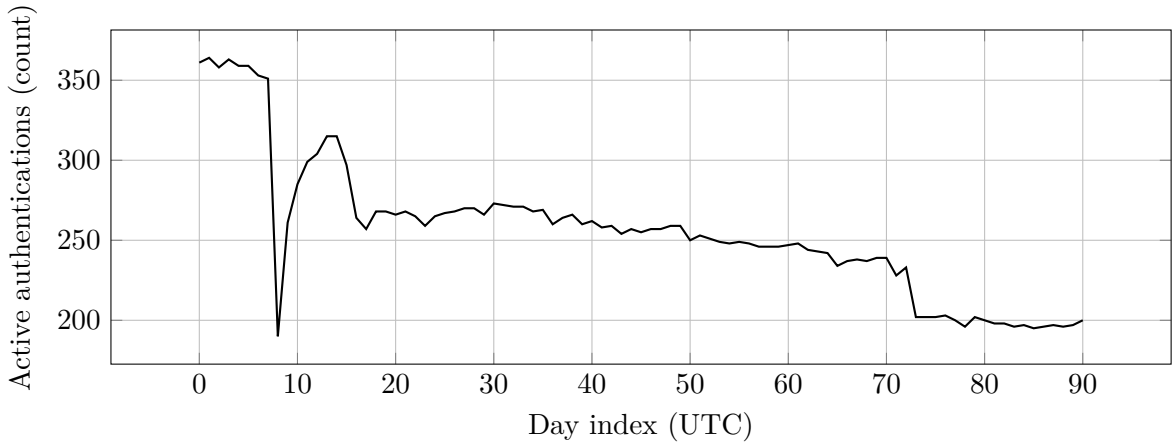

Figure A2. W1 on-chain `bioauth.activeAuthentications` series (daily samples). Each point is the vector length (number of entries) at the sampled block for the corresponding UTC day.

A.2.3 W1 bootstrap variability

To assess variability due to finite sampling of block authors, we bootstrap-resample the sampled per-block author labels with replacement (2000 iterations; sample size fixed). Table A1 reports 95% percentile intervals. Because top-k shares are maxima (order statistics), bootstrap distributions can be biased upward relative to the point estimate; treat these intervals as a robustness check rather than as a population confidence interval. This does not model systematic effects from fixed-step sampling.

Metric (W1, step=25)	Observed	95% CI (percentile)
Top-1 share	0.00465	[0.00478, 0.00540]
Top-5 share	0.02299	[0.02344, 0.02513]
Top-10 share	0.04514	[0.04614, 0.04865]
HHI	0.00343	[0.00343, 0.00346]
Gini (author block counts)	0.348	[0.349, 0.357]

Table A1. Bootstrap variability for W1 author concentration metrics under resampling of the sampled blocks (2000 iterations).

A.2.4 W1 sensitivity to fixed-step sampling

Fixed-step sampling can introduce aliasing. To probe sensitivity, we repeat the W1 author extraction with step sizes 25/50/100 and recompute concentration metrics (Table A2). Coarser steps yield slightly higher concentration estimates, suggesting aliasing; nonetheless, all runs indicate low concentration at the scale relevant to the main claim.

Step	Sampled blocks	Unique authors	Top-1	Top-10	Gini
25	51,881	405	0.00465	0.04514	0.348
50	25,941	403	0.00532	0.04749	0.349
100	12,971	399	0.00571	0.05011	0.351

Table A2. W1 sensitivity to author-sampling step size. Metrics are computed over sampled block authors; session churn statistics are unchanged because they use session snapshots.

A.2.5 W3 concentration time series

We do not include a W3 *concentration* time-series plot because the long horizon mixes distinct protocol and validator-set regimes. Interpreting long-horizon trends without explicit regime segmentation risks overclaiming; one reasonable approach is to segment by protocol/runtime changes and by major shifts in validator-set size. We include the W3 validator-set plot in Section 7 only to provide long-horizon context for regime changes, and we report W3 author concentration as a snapshot.

References

[1] John R. Douceur. The sybil attack. In *Peer-to-Peer Systems: First International Workshop (IPTPS 2002)*, volume 2429 of *Lecture Notes in Computer Science*, pages 251–260. Springer, 2002.

[2] Satoshi Nakamoto. Bitcoin: A peer-to-peer electronic cash system. `https://bitcoin.org/bitcoin.pdf`, 2008. Accessed 2026-01-03.

[3] Sunny King and Scott Nadal. Ppcoin: Peer-to-peer crypto-currency with proof-of-stake. `https://peercoin.net/assets/paper/peercoin-paper.pdf`, 2012. Accessed 2026-01-03.

[4] Aziz Mohaisen and Jung Kim. The sybil attacks and defenses: A survey. *Smart Computing Review*, 3(6), 2013.

[5] Lin William Cong, Zhiguo He, and Jiasun Li. Decentralized mining in centralized pools. `https://www.nber.org/papers/w25592`, 2019. NBER Working Paper 25592. DOI: 10.3386/w25592. PDF: `https://www.nber.org/system/files/working_papers/w25592/w25592.pdf`. Accessed 2026-01-27.

[6] Giulia Fanti, Leonid Kogan, Sewoong Oh, Kathy Ruan, Pramod Viswanath, and Guang Wang. Compounding of wealth in proof-of-stake cryptocurrencies. *Financial Cryptography and Data Security*, 2019. Also available as arXiv:1809.07468. `https://arxiv.org/abs/1809.07468`. Accessed 2025-12-31.

[7] Rob Wiblin. Vitalik buterin on effective altruism, better ways to fund public goods, the blockchain's problems so far, and how it could yet change the world.

https://80000hours.org/podcast/episodes/vitalik-buterin-new-ways-to-fund-public-goods/, 2019. Podcast transcript. Accessed 2025-12-31.

[8] Jonah Brown-Cohen, Arvind Narayanan, Christos-Alexandros Psomas, and S. Matthew Weinberg. Formal barriers to longest-chain proof-of-stake protocols. In *Proceedings of the 2019 ACM Conference on Economics and Computation*, 2019. Also available as arXiv:1809.06528 (2018). https://arxiv.org/abs/1809.06528. Accessed 2025-12-31.

[9] Dimitris Karakostas, Aggelos Kiayias, Christos Nasikas, and Dionysis Zindros. Cryptocurrency egalitarianism: A quantitative approach. *arXiv preprint*, 2019. https://arxiv.org/abs/1907.02434. Accessed 2025-12-31.

[10] Martín Borge, Eleftherios Kokoris-Kogias, Philipp Jovanovic, Linus Gasser, Nicolas Gailly, and Bryan Ford. Proof-of-personhood: Redemocratizing permissionless cryptocurrencies. In *2017 IEEE European Symposium on Security and Privacy Workshops (EuroS&PW)*, 2017.

[11] Anil K. Jain, Arun Ross, and Salil Prabhakar. An introduction to biometric recognition. *IEEE Transactions on Circuits and Systems for Video Technology*, 14(1):4–20, 2004.

[12] ISO/IEC. Iso/iec 30107-3: Biometric presentation attack detection — part 3: Testing and reporting. https://www.iso.org/standard/67381.html, 2017. Standard. Accessed 2025-12-31.

[13] Yevgeniy Dodis, Rafail Ostrovsky, Leonid Reyzin, and Adam Smith. Fuzzy extractors: How to generate strong keys from biometrics and other noisy data. *SIAM Journal on Computing*, 38(1):97–139, 2008.

[14] ISO/IEC. Iso/iec 24745: Biometric information protection. https://www.iso.org/standard/52946.html, 2011. Standard; later revisions may exist. Accessed 2025-12-31.

[15] Jan Camenisch and Anna Lysyanskaya. An efficient system for non-transferable anonymous credentials with optional anonymity revocation. In *Advances in Cryptology — EUROCRYPT 2001*, volume 2045 of *Lecture Notes in Computer Science*, pages 93–118. Springer, 2001.

[16] Amazon Web Services. Attest an amazon ec2 instance with amd sev-snp. https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/snp-attestation.html, 2024. Accessed 2025-12-31.

[17] Miguel Castro and Barbara Liskov. Practical Byzantine fault tolerance. In *Proceedings of the Third Symposium on Operating Systems Design and Implementation (OSDI)*, 1999.

[18] Cynthia Dwork, Nancy Lynch, and Larry Stockmeyer. Consensus in the presence of partial synchrony. *Journal of the ACM*, 35(2):288–323, 1988.

[19] Humanode. Proof-of-biometric-uniqueness (pobu): The consensus of the living. https://whitepaper.humanode.io/whitepaper/proof-of-biometric-uniqueness-pobu-the-consensus-of-the-living, 2025. Accessed 2025-12-31.

[20] Humanode. Confidential vms (cvms). https://cvm.humanode.io/confidential-vms.md, 2024. Accessed 2025-12-31.

[21] Humanode. Humanode mainnet guide: Faqs. https://gitbook.humanode.io/mainnet-guide/faqs, 2024. Accessed 2025-12-31.

[22] Humanode. Regular biometric servers & cvm rotations are scheduled in 24 hours. https://blog.humanode.io/regular-biometric-servers-cvm-rotations-are-scheduled-in-24-hours/, 2025. Operational announcement. Accessed 2025-12-31.

[23] Humanode. Humanode docs (gitbook): Chains. https://gitbook.humanode.io/docs/chains, 2025. Public documentation for chain endpoints. Accessed 2025-12-31.

[24] Patrick Grother, Mei Ngan, and Kayee Hanaoka. Face recognition vendor test (frvt) part 3: Demographic effects. Technical Report NISTIR 8280, National Institute of Standards and Technology, 2019. Accessed 2025-12-31.