MATHEMATICAL MODELING AND COMPUTING, Vol. 12, No. 4, pp. 1320–1332 (2025)

Mixed-Weight Committee Selection in Proof-of-Stake: Tunable Stake-Baseline Mixing with Exponential Tail Guarantees and Incentive Compatibility

Solomka I. R., Liubinskyi B. B., Peniak B. O.

Lviv Polytechnic National University,
12 S. Bandera str., 79013, Lviv, Ukraine

(Received 7 July 2025; Revised 29 November 2025; Accepted 5 December 2025)

Proof-of-Stake (PoS) blockchains often select committees in direct proportion to stake, which makes security sensitive to large validators and stake concentration. In such settings, a purely stake-based lottery can sometimes produce committees whose adversarial share crosses the safety threshold, even if the global adversarial stake remains below one third. This paper introduces a simple mixed-weight rule that combines stake with a bounded baseline distribution through a single mixing parameter λ. The rule leaves committee size, rewards, and VRF-based sortition unchanged, but pulls weight away from highly concentrated positions. Proved that, whenever the adversary is more concentrated than the baseline, the expected adversarial seats fall linearly in λ, while standard concentration bounds show an exponential drop in committee-capture probability. While the mechanism relies on entity-level attribution (or high-cost identities) to prevent Sybil attacks, experiments on ten production PoS networks indicate that modest mixing (around $\lambda = 0.3$) reduces expected adversarial seats by about one quarter and tightens worst-case guarantees by orders of magnitude.

Keywords: *Proof-of-Stake; committee sampling; λ-mixture mechanism; reputation systems; incentives; tail risk.*

2010 MSC: 68M14, 60E15, 68W20, 94A60, 91A80 **DOI:** 10.23939/mmc2025.04.1320

1. Introduction

Proof-of-Stake (PoS) [1, 2] systems derive safety from the assumption that committees rarely contain a destructive fraction of adversarial stake. In practice, several factors conspire against this idealized picture: stake concentration into a small set of operators; correlated failures due to common infrastructure; liveness incidents where a large fraction of voting power goes offline simultaneously; and deliberate attempts to bias leader or committee selection by exploiting predictable weight maps. Even when protocols employ Verifiable Random Functions (VRFs) [3] to choose committees, the input distribution to the lottery is typically fixed and strictly proportional to stakes, leaving limited room to attenuate the most dangerous tail events without increasing the committee size.

This paper develops and analyzes a minimal, auditable modification to the input distribution, namely an affine mixture of stake and a bounded baseline:

$$q_i(\lambda) = (1-\lambda)w_i + \lambda u_i, \quad \sum_i w_i = \sum_i u_i = 1, \quad \lambda \in [0,1].$$

The baseline u acts as an uniform regularizing component that dilutes concentrated positions. It may be uniform or entity-level reputation–aware (incorporating signals such as downtime and slashing). The single knob λ continuously trades off pure stake proportionality against dispersion toward the baseline. It is shown that the adversary's expected seats decrease linearly whenever the adversary holds more stake share than baseline share, and that standard Chernoff/Hoeffding [4] tails shrink exponentially in the resulting safety margin.

Scope and Assumptions. A critical challenge in mixed-weight selection is incentive compatibility. In a fully permissionless setting with low node-creation costs, a uniform baseline incentivizes rational actors to split their stake across multiple identities (Sybil attack [5]) to harvest additional baseline weight. To address this, our mechanism is designed for settings where *Entity Resolution* is present. This includes permissioned ledgers, consortium blockchains, or networks with high economic barriers to entry where the cost of generating a new validated identity exceeds the marginal gain from the baseline parameter λ. Under this assumption of entity-level attribution, it is formally shown that the mechanism preserves incentive compatibility and renders stake-splitting a dominated strategy.

To assess practical impact, daily snapshot tooling was built and simulated risk on ten production PoS networks. The data indicate that even modest λ values produce substantial headroom against committee capture, without modifying committee sizes or finality thresholds.

2. Related work

Committee formation and voting in Proof-of-Stake (PoS) [1] have been explored from multiple angles: stake-proportional or VRF-based selection, profile- or reputation-weighted voting, and hybrid PoW/PoS defenses. Each line of work targets a different layer of the pipeline-who gets into the committee versus how the committee aggregates votes-leaving a gap for a tunable, analytically-tractable selection rule that reduces adversarial concentration without changing committee size or imposing heavy learning machinery. Our proposal fills precisely this gap by introducing an affine stake–reputation mixture ($w_i(\lambda) = (1-\lambda)w_i + \lambda u_i$) with explicit tail-probability bounds for adversarial capture, a knob (λ) that continuously trades off fairness and safety, and straightforward compatibility with standard VRF sortition.

Weighted voting on blockchain. Leonardos [6] studied weighted voting within a formed committee and update the "profiles" of the validator by multiplicative weights to improve the robustness against abstention and faults, keeping the selection mechanism unchanged. Their model shows that weighting votes by learned profiles can raise the probability that a formed committee reaches the correct decision. This work differs in locus and tooling: the selection weights is modified before committees exist, without maintaining per-validator learning dynamics, and provided closed-form Chernoff-type bounds.

VRF-based committee selection. Algorand [7] introduced cryptographic sortition and ephemeral VRF proofs [8] to elect proposers and committees with stake-proportional probability. This design achieves scalable BA finality with low latency, but the stake-to-selection mapping is essentially fixed. Our mixture keeps the VRF machinery intact and merely replaces the input weights to the VRF lottery. As a result, it retains all verifiability and anti-grinding properties of sortition, while adding a policy lever that dilutes concentrated stake by blending it with an entity-aware baseline u_i (uniform or reputation-derived), thereby shrinking the adversarial mean seats $\mu(\lambda)$.

Finality overlays and slashing-based accountability. Casper FFG [9] adds PoS finality atop a proposal mechanism, bringing explicit slashing conditions and safety under $< 1/3$ Byzantine [10] weight. Like weighted voting, it operates after selection and is agnostic to how committees are formed. Our mixture is orthogonal: it reduces the probability that any given committee even approaches the dangerous region where Casper has to rely on slashing to preserve safety.

Hybrid PoW/PoS defenses. Paper [11] propose a hybrid mechanism aimed at double-spending 51% resilience by combining PoW and PoS subsystems. Our approach stays entirely within PoS and does not alter reward semantics or fork choice; instead, it reweights selection inputs to reduce committee capture probability under an assumed adversarial stake fraction (α).

Probabilistic PoS protocols. LaKSA [12] develops a probabilistic PoS design with rigorous safety/liveness analysis and parameter trade-offs. Our mixture rule can serve as a drop-in input to such probabilistic designs, preserving their proofs while lowering the baseline risk of adversarially heavy committees.

Reputation-based node selection (PoIR and variants). Recent work on Proof of Intelligent Reputation (PoIR) [13] brings ML models and centralized reputation databases into consensus. Our design deliberately avoids heavy, opaque estimators and the governance risks of a reputation oracle. The baseline (u_i) can incorporate simple, auditable reputation components but remains bounded and mixed affinely with stake.

Committee size versus security trade-offs. In [14] is analyzed how to improve the size–security trade-off of stake-based committee selection via low-variance assignment. Focus in this paper is different and complementary: the committee size is hold fixed and instead reshape selection weights with a single scalar (λ), proving that the expected adversarial seats decrease and that Chernoff-style tails contract correspondingly.

Why the λ-mixture is preferable as a tunable knob. Prior art either keeps stake-proportional selection fixed [7], or changes the consensus family altogether, or replaces stake with learned reputation. The λ-mixture isolates a minimal, auditable change that provides a continuous safety–fairness trade-off without changing committee size, VRF wiring, or reward rules. Because the transform is linear and bounded, exact means and standard Chernoff/Hoeffding [4,15] bounds for adversarial capture can be written down. The result is a practical control surface for protocol engineers: setting λ tightens worst-case tails in a predictable way.

Table 1 summarizes all the findings.

Table 1. Comparison of Committee Selection and Security Approaches in PoS.

Approach	**Selection Weight**	**Tunable Parameter**	**Tail Bounds / Variance Control**	**Incentive Compatibility**
λ-Mixture (Proposed)	Affine mixture of stake and a baseline: $(1-\lambda)w_i + \lambda u_i$	$\lambda \in [0,1]$ mixes stake with baseline	Provides explicit, closed-form exponential (Chernoff-type) tail bounds	Preserves long-run stake-proportional rewards; robust in identified / permissioned sets
Weighted Voting [6]	Unchanged (operates post-selection)	None for selection	Improves robustness of *formed* committee's decision	Introduces risks of centralization via profiling
Standard VRF Sortition [7]	Directly proportional to stake (s_i)	None	Variance is dictated entirely by the stake distribution	High compatibility, rewards tied to stake
Finality Overlays (Casper) [9]	Agnostic to selection; operates on formed committees	None for selection	Does not prevent malicious committee formation but provides economic finality	Strong economic disincentives (slashing) for misbehavior
Reputation-Based (PoIR) [13]	Reputation scores from ML models	Implicit (ML parameters)	No explicit, closed-form bounds mentioned	Moves away from stake-proportional rewards, creating governance risks

3. Model and Definitions

Validator set and adversary. There are n validators with stake shares $w_i > 0$ and $\sum_i w_i = 1$. The adversarial set of validators $A \subseteq \{1, \dots, n\}$ has aggregate adversarial stake share

$$\alpha = \sum_{i \in A} w_i.$$

Committee and threshold. Each round forms a size-m committee. The dangerous threshold is the Byzantine [10,16] safety cut $\tau = \lceil m/3 \rceil$.

Baseline distribution. A normalized baseline $u = (u_1, \dots, u_n)$ satisfies $\sum_{i=1}^{n} u_i = 1$. Uniform baseline uses $u_i = 1/n$. For a reputation baseline is set

$$u_i = \frac{r_i}{\sum_{j=1}^{n} r_j},$$

where $r_i \in [0,1]$ is a reputation score for validator i.

The aggregate adversarial baseline share is

$$\rho = \sum_{i \in A} u_i.$$

How reputation r_i is computed. Let $d_i \in [0,1]$ be recent downtime, $s_i \geqslant 0$ scaled to $[0,1]$ – a slashing penalty summary, and $c_i \in [0,1]$ a churn/identity-instability score. A transparent instantaneous score is $\widehat{r}_i = \exp\big(-\beta_1 d_i - \beta_2 s_i - \beta_3 c_i\big)$, with policy weights $\beta_k \geqslant 0$. To resist gaming, lets apply exponential smoothing $r_i^{(t)} = (1-\kappa) r_i^{(t-1)} + \kappa \widehat{r}_i^{(t)}$, $\kappa \in (0,1]$, then normalize $u_i = r_i^{(t)} \big/ \sum_j r_j^{(t)}$. Reputation is computed at the entity level (keys aggregated), not per-address.

Assumption: Identified Validators. It is assumed the protocol operates in a regime where Entity Resolution is possible (e.g., KYC, Proof-of-Authority, or high economic barriers). This ensures ρ reflects the true adversarial entity count rather than a Sybil-inflated count.

Algorithm 1 below specifies how entity-level penalties and smoothed reputation weights are incorporated into the mixed selection probabilities.

Algorithm 1 Entity-Level Penalty Application and Renormalization.

```
Require: Entity $e$ with total stake $S_e = \sum_{i=1}^{k_e} w_i$, $k_e$ identities each with stake $w_i$, reputation $r_i$, and
    penalty score $\pi_i$ for $i = 1, \dots, k_e$; global penalty intensity $\gamma \in [0,1]$.
Ensure: Normalized post-penalty stake weights $\overline{w}_i$ and baseline weights $\overline{u}_i$ for all identities.
Step 1 – Entity-level aggregation.
Compute the entity-wide baseline mass $U_e \leftarrow \sum_{i=1}^{k_e} r_i$.
for each identity $i = 1$ to $k_e$
    Allocate baseline mass proportionally to intra-entity stake: $u_i \leftarrow \frac{w_i}{S_e} \cdot U_e$.
Compute the entity-wide penalty signal $\Pi_e \leftarrow f(\pi_1, \dots, \pi_{k_e})$, where $f$ is a policy function; typical
    choices are
        $\Pi_e = \max_i \pi_i (strict), \quad \Pi_e = \mathrm{mean}_i \, \pi_i (smooth)$
    or a convex combination of these.
Step 2 – Post-penalty scaling at entity level.
for each identity $i = 1$ to $k_e$
    Apply the entity-level penalty to stake: $\widetilde{w}_i \leftarrow (1 - \gamma \Pi_e)\, w_i$.
    Apply the same entity-level penalty to the baseline: $\widetilde{u}_i \leftarrow (1 - \gamma \Pi_e)\, u_i$.
Step 3 – Global renormalization across the network.
Compute global post-penalty sums over all identities in all entities:
        $W_{\mathrm{tot}} \leftarrow \sum_j \widetilde{w}_j, \quad U_{\mathrm{tot}} \leftarrow \sum_j \widetilde{u}_j.$
for each identity $i$ in entity $e$
    Normalize the stake weight: $\overline{w}_i \leftarrow \frac{\widetilde{w}_i}{W_{\mathrm{tot}}}$.
    Normalize the baseline weight: $\overline{u}_i \leftarrow \frac{\widetilde{u}_i}{U_{\mathrm{tot}}}$.
Step 4 – Mixed selection probabilities.
for each identity $i$
    Compute the mixed selection probability
        $q_i(\lambda, \gamma) \leftarrow (1 - \lambda)\, \overline{w}_i + \lambda\, \overline{u}_i,$
    for a chosen mixture parameter $\lambda \in [0,1]$.
Invariant. After Steps 1–3, the normalized weights satisfy
        $\sum_i \overline{w}_i = 1 \quad \text{and} \quad \sum_i \overline{u}_i = 1,$
    so that the mixed probabilities also satisfy
        $\sum_i q_i(\lambda, \gamma) = 1 \quad \text{for all} \quad \lambda \in [0,1], \quad \gamma \in [0,1].$
```

Pre-normalization penalties and renormalization. Optional penalties with intensity $\gamma \in [0,1]$ and entity-level signals $\pi_i \in [0,1]$ scale raw weights:

$$\widetilde{w_i} = (1-\gamma\pi_i)\,w_i, \quad \widetilde{u_i} = (1-\gamma\pi_i)\,u_i.$$

After renormalization,

$$\overline{w}_i = \frac{\widetilde{w_i}}{\sum_j \widetilde{w_j}}, \quad \overline{u}_i = \frac{\widetilde{u_i}}{\sum_j \widetilde{u_j}},$$

and adversarial masses become

$$\overline{\alpha} = \sum_{i\in A} \overline{w}_i, \quad \overline{\rho} = \sum_{i\in A} \overline{u}_i.$$

Mixed selection rule. Each seat is sampled independently (with replacement) with probabilities

$$q_i(\lambda,\gamma) = (1-\lambda)\,\overline{w}_i + \lambda\,\overline{u}_i, \quad \lambda \in [0,1].$$

Let X be the number of adversarial members in the committee. If the protocol samples seats without replacement, negative association tightens concentration; the binomial-based tail bounds used below remain conservative. The parameter λ is applied to selection; long-run payouts remain stake-proportional to avoid changing macro-economics and to support incentive compatibility.

4. Main results

Expectation and safety margin. The adversary's expected seats are

$$\mu(\lambda,\gamma) = E[X] = m\big[(1-\lambda)\,\overline{\alpha} + \lambda\,\overline{\rho}\big].$$

Hence

$$\frac{\mathrm{d}\mu}{\mathrm{d}\lambda} = m\,(\overline{\rho} - \overline{\alpha}),$$

so whenever $\overline{\alpha} > \overline{\rho}$ the expectation decreases linearly with λ. The safety gap to the threshold,

$$g(\lambda,\gamma) = \tau - \mu(\lambda,\gamma) = \tau - m\big[(1-\lambda)\,\overline{\alpha} + \lambda\,\overline{\rho}\big],$$

grows affinely in λ at slope $m(\overline{\alpha} - \overline{\rho}) > 0$.

Chernoff upper tail (with penalties). Writing $p(\lambda,\gamma) = (1-\lambda)\,\overline{\alpha} + \lambda\,\overline{\rho}$ and $\mu(\lambda,\gamma) = mp(\lambda,\gamma)$, the standard multiplicative Chernoff bound gives, for any $\tau > \mu(\lambda,\gamma)$,

$$P[X \geqslant \tau] \leqslant \exp\left(-\mu(\lambda,\gamma)\,\phi\left(\tfrac{\tau}{\mu(\lambda,\gamma)}\right)\right), \quad \phi(x) = x\ln x - x + 1.$$

The MGF derivation applies verbatim after renormalization; see Appendix A.1.

Hoeffding's distribution-free bound.

$$P[X \geqslant \tau] \leqslant \exp\left(-\tfrac{2(\tau-\mu(\lambda,\gamma))^2}{m}\right).$$

Both bounds improve exponentially as $g(\lambda,\gamma)$ grows linearly in λ.

Lower bounds and tightness. For binomial [17] (and broadly Poisson-binomial) tails, Chernoff's exponent is sharp in the large-deviation regime:

$$P[X \geqslant (1+\delta)\mu] = \Theta\big(m^{-1/2}\big)\exp\big(-\mu\phi(1+\delta)\big),$$

so the upper bound captures the correct exponential rate. In practice, our Monte Carlo [18] often reports zero unsafe events at moderate λ; this sits below the conservative bound, exactly as expected.

Sensitivity to α near $1/3$. Since

$$p(\lambda,\gamma) = (1-\lambda)\overline{\alpha} + \lambda\overline{\rho}, \quad \tau = \left\lceil \frac{m}{3} \right\rceil,$$

the condition $p(\lambda,\gamma) \leqslant \frac{1}{3} - \varepsilon$ for a target margin $\varepsilon > 0$ holds whenever

$$\lambda \geqslant \frac{\overline{\alpha} - \left(\frac{1}{3} - \varepsilon\right)}{\overline{\alpha} - \overline{\rho}}, \quad \text{assuming} \quad \overline{\alpha} > \overline{\rho}.$$

This gives a clear governance dial: the smaller $\overline{\rho}$ (entity-aware reputation baseline), the smaller λ needed to certify a given gap. Without replacement. The same formulas remain valid as conservative certificates; negative dependence reduces variance, shrinking tails further.

5. Experimental evaluation

This section quantifies the security effect of mixed-weight committee selection on ten production Proof-of-Stake (PoS) networks over a two-month window (September–October 2025). The analysis reports how the mixing parameter $\lambda \in \{0, 0.1, 0.2, 0.3\}$ alters the expected number of adversarial seats and a worst-case tail guarantee for Byzantine capture, under several BFT thresholds.

Daily snapshots of active validator sets and simulated committee outcomes were collected automatically. For each day and network, the simulator evaluated a grid over λ and α, producing per-gridpoint estimates of the mean adversarial seats $\mu(\lambda)$ and the Chernoff upper bound on the capture event $\Pr[X \geqslant \tau]$ with $\tau = \alpha m$ and committee size m. The present section aggregates these daily outputs into day-wise medians per network, and then compares across networks.

The portfolio comprises Avalanche [19], Celestia [20], Cosmos Hub [21], Injective [22], NEAR [23], Osmosis [24], Sei [25], Solana [26], Sui [27], Tezos [28], spanning diverse validator cardinalities and typical committee sizes. In the current empirical evaluation, the snapshots do not yet include downtime, slashing, or churn metrics; therefore, r_i is not instantiated from on-chain data, and the baseline u_i is taken to be uniform.

5.1. Metrics

Let X denote the number of adversarial seats in a committee of size m. Expected adversarial seats is

$$\mu(\lambda) \equiv E[X].$$

This quantity summarizes the typical adversarial influence; percent reductions are reported relative to $\lambda = 0$. Chernoff tail bound. For threshold $\tau = \alpha m$,

$$\Pr[X \geqslant \tau] \leqslant \exp\big(-mD(\alpha \parallel q(\lambda))\big),$$

where $q(\lambda)$ is the per-draw adversarial success probability implied by the mixing law and $D(\cdot \parallel \cdot)$ is the binary KL divergence. Because these probabilities can be extremely small, the results are presented on a base-10 logarithmic scale:

$$\log_{10} \Pr[X \geqslant \tau] \leqslant -\frac{m}{\ln 10} D(\alpha \parallel q(\lambda)).$$

Each unit on this axis equals one order of magnitude; differences are thus directly interpretable as multiplicative improvements in the worst-case guarantee. Remark on empirical tails. For large m and moderate λ, Monte Carlo estimates of $\Pr[X \geqslant \tau]$ frequently underflow to zero at feasible trial counts; the Chernoff bound remains a conservative and comparable statistic across networks and days.

5.2. Aggregate outcomes at $\alpha = 0.25$

Table 2 summarizes, for $\alpha = 0.25$, the median reduction in the expected adversarial seats and the corresponding improvement in worst-case tail guarantees when the mixing parameter increases from $\lambda = 0$ to $\lambda = 0.3$.

Table 2. Cross-network Median Improvements at $\alpha = 0.25$ under λ: 0 → 0.3.

Network	**Threshold** α	**Median** μ ($\lambda = 0$)	**Median** μ ($\lambda = 0.3$)	μ **Reduction** (%)	$\Delta \log_{10}$ **Bound**
Avalanche	0.25	90.8698	65.9403	27.43	4.83
Solana	0.25	96.7064	69.1534	28.49	5.69
NEAR	0.25	33.0855	23.9999	27.46	1.36
Tezos	0.25	30.7990	21.9186	28.83	0.99
Cosmos Hub	0.25	22.7290	16.3903	27.89	0.83
Celestia	0.25	10.7244	8.2734	22.85	0.36
Osmosis	0.25	10.2810	8.0367	21.83	0.37
Sui	0.25	12.2400	10.1117	17.39	0.33
Injective	0.25	5.5833	4.3883	21.40	0.14
Sei	0.25	4.4950	3.7465	16.65	0.07

Notes. μ is the expected number of adversarial seats; entries are medians across daily snapshots. $\Delta\log_{10}$ Bound reports the increase in base-10 orders of magnitude (i.e., tightening) of the Chernoff tail bound for $\Pr[X \geqslant \alpha m]$ moving from $\lambda = 0$ to $\lambda = 0.3$.

Across all ten networks, moving from $\lambda = 0$ to $\lambda = 0.3$ reduces the day-wise median μ by approximately 24% on average (mean across networks), with individual networks ranging from 16.7% (Sei) to 28.8% (Tezos). Thus, even a modest mixed-weight component yields a double-digit decrease in the average number of adversarial seats. The median improvement in the log-tail bound displays a pronounced dependence on committee size:

— Large-committee systems (e.g., Solana and Avalanche) improve by $+5.69$ and $+4.83$ orders of magnitude, respectively. In plain terms, the worst-case per-draw failure upper bound tightens by factors of roughly 4.9×10^5 and 6.7×10^4.

— Moderate-committee systems (e.g., NEAR, Tezos, Cosmos Hub) improve by $\approx 1.0 \pm 0.4$ orders (i.e., about 7 times to 25 times tighter).

— Small-committee systems (e.g., Injective, Sei, Sui, Celestia, Osmosis) register improvements between ≈ 0.07 and ≈ 0.37 orders (i.e., ~ 1.2 times to ~ 2.3 times tighter).

Taken together, the cross-network mean improvement in log-orders is $\approx +1.5$ while the cross-network median is $\approx +0.6$, reflecting a bimodal distribution driven by committee size.

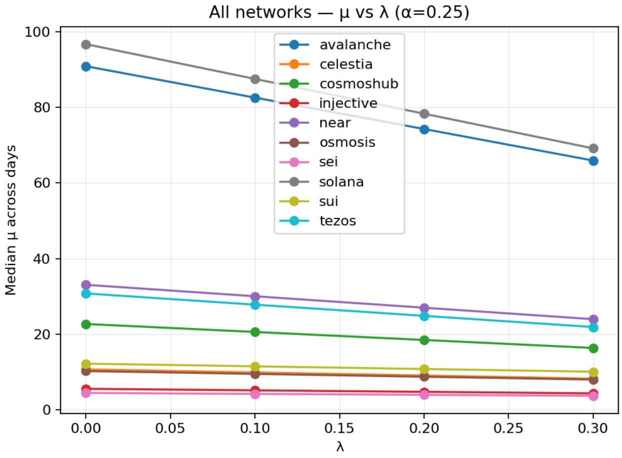

Fig. 1. All networks - median μ vs λ at $\alpha = 0.25$.

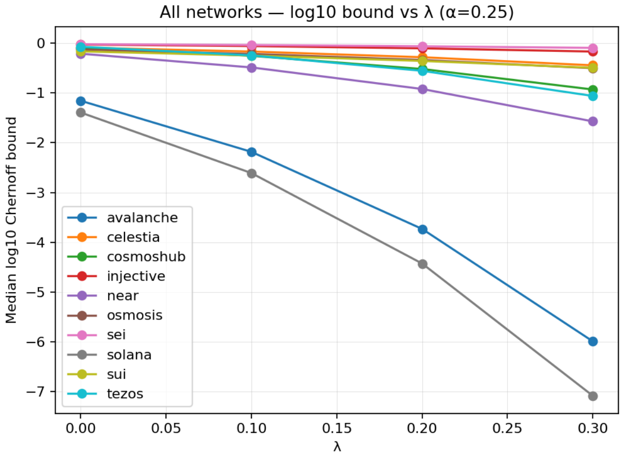

Fig. 2. Median $\log_{10}$ Chernoff bound vs λ at $\alpha = 0.25$.

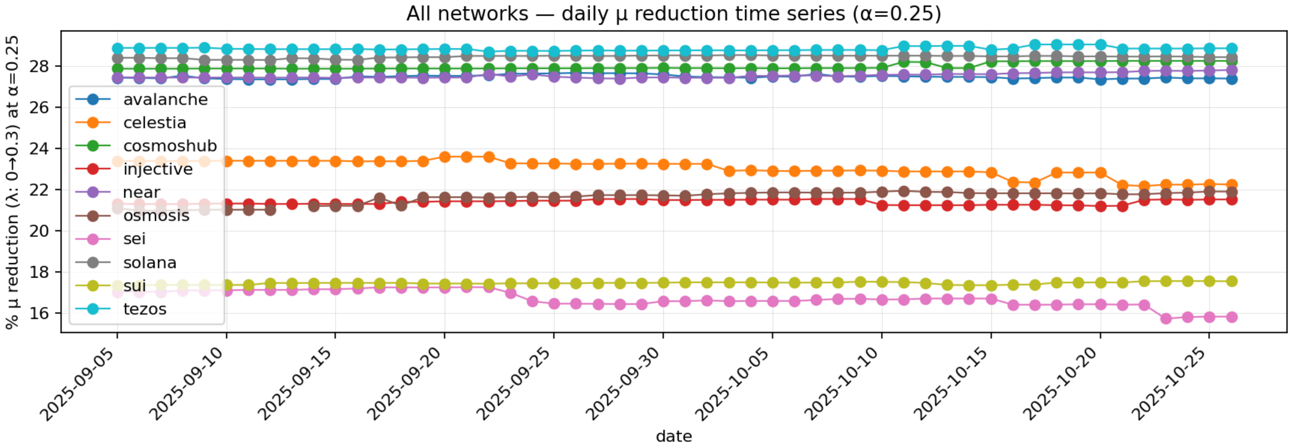

Fig. 3. Daily % μ-reduction time series at $\alpha = 0.25$.

Figure 1 illustrates the median μ as a function of λ at $\alpha = 0.25$ for all networks. Curves decrease almost affinely, indicating stable, near-linear gains in the first moment. Figure 2 displays the median $\log_{10}$ tail bound versus λ at $\alpha = 0.25$. Systems with larger m plunge to around -7 at $\lambda = 0.3$, while small m networks settle near 0 to -1. A value of -7 on the $\log_{10}$ axis corresponds to an upper bound of 10^{-7} for a single committee draw, i.e., one in ten million. Values near -1 correspond to "at most ten percent." The divergence arises because the Chernoff exponent scales linearly in m and increases with the gap $\alpha - q(\lambda)$. In large-committee systems, the same decrease in $q(\lambda)$ (induced by the mixing) is magnified by m, yielding orders-of-magnitude improvements in the worst-case guarantee.

The daily trajectories of the percent reduction in μ for λ: $0 \to 0.3$ at $\alpha = 0.25$ exhibit tight dispersion around their medians; Figure 3 shows all networks on one axis with limited day-to-day variation. This supports the interpretation that the observed gains are structural and not artifacts of particular days' validator snapshots.

5.3. Value and operational meaning

The results indicate that introducing a small mixed-weight component $\lambda \in [0.2, 0.3]$:

- Systematically reduces typical adversarial influence. Median reductions of μ fall in the 16%-29% band across heterogeneous networks.
- Sharpens worst-case guarantees when committees are larger. For networks with m in the hundreds, the same knob increases the Chernoff exponent sufficiently to push the tail guarantee to the sub-micro regime (e.g., $\leqslant 10^{-7}$ per committee draw).
- Offers a tunable security–performance control. The near-affine $\mu(\lambda)$ curves and monotone tail improvements provide transparent predictability as λ is tuned, enabling protocol designers to target specific first-moment reductions or tail-risk orders without redesigning the selection mechanism.

Small-committee systems still benefit in expectation; if stronger tail guarantees are desired, a moderate increase in m combined with $\lambda \approx 0.3$ would move them into the regime where order-of-magnitude improvements emerge, consistent with the exponential dependence on m. Threats to validity and robustness checks:

- Empirical tail probabilities often underflow at feasible trial budgets when m is large. The Chernoff bound remains conservative; in the dataset, empirical estimates never exceeded the bound, confirming expected slackness.
- Day-wise medians reduce sensitivity to single-day outliers; the time-series panel suggests stability over the entire window.
- Without-replacement effects and mild negative dependence tend to improve concentration relative to sampling; the theoretical section discusses why the reported bounds remain valid or conservative in these regimes.

5.4. Evaluation summary

Across two months and ten PoS networks, introducing a modest mixed-weight component consistently reduces the average adversarial presence in committees by about one quarter, and, in large-committee settings, strengthens worst-case guarantees by several orders of magnitude. The effect is monotone in λ, robust across α, and stable over time, providing a simple, tunable mechanism to materially raise committee safety without architectural change.

Reproducibility note. The implementation, including snapshot collection (snapshots.py) and risk evaluation (risk_batch.py), is publicly available at `https://github.com/ihosol/mixweight-committee-consensus`. The repository includes the code for fetching validator weights from network APIs, computing mixed weights $q_i(\lambda)$, $\mu(\lambda)$, Chernoff bounds, and empirical risks via Monte Carlo. Results are generated with parameters $m = 0.4n$, $\alpha \in [0.20, 0.25, 0.30, 0.33]$, $\lambda \in [0, 0.1, 0.2, 0.3]$, and output to risk_result.csv for reproducibility. Commit hash: fdb27a0 (as of September 2025)

6. Broader impact

The mixture mechanism primarily targets the safety of high-value applications and DeFi protocols where the cost of a committee capture event is catastrophic. By reducing the probability that any given committee holds a destructive fraction of adversarial weight, the mechanism lowers the baseline risk of censorship, double-signing, and correlated liveness failures.

Applicability and Governance. The immediate utility of this mechanism is highest for *Consortium Blockchains, Enterprise Ledgers, and Proof-of-Authority networks.* In these settings, where identities are known or strictly regulated, the λ-mixture provides a mathematically rigorous way to enforce decentralization and prevent any single member-even one with a majority of tokens-from dominating committee selection.

Risks in Permissionless Settings. In fully anonymous, permissionless environments, introducing a baseline weight (u_i) creates an economic incentive for identity splitting (Sybil attacks). Without the entity-resolution layer assumed in this work, rational actors would fragment their stake to maximize their baseline rewards. Consequently, deployment on public main-nets requires either the high-cost identity barriers discussed in Section 3 or the advanced game-theoretic mitigation strategies reserved for future work.

Performance. From a computational standpoint, the mixture adds negligible overhead. Computing baseline weights and optional penalty scalars is $O(n)$ per epoch, and sampling committee seats remains $O(1)$ per seat. The mechanism is backwards-compatible with standard VRF-based sortition pipelines, allowing existing protocols to adopt it as a soft-fork governance upgrade.

7. Future work

This paper established the probabilistic bounds and safety margins for mixed-weight selection under the assumption of identified entities. Future research will focus on relaxing this assumption to extend the mechanism to fully permissionless environments.

1. Sybil-Resistant Incentive Design. The most critical extension is developing a game-theoretic framework where identity splitting becomes a strictly dominated strategy even without explicit entity resolution. This paper aims to derive a reward-slashing function $S(k)$ dependent on observable behavioral correlations (e.g., correlated uptime/downtime) that makes the cost of maintaining k Sybil nodes exceed the marginal gain from the baseline parameter λ.

2. Decentralized Reputation Oracles. While this work assumes a transparent baseline u_i, future iterations will explore decentralized oracle networks to compute reputation scores. This includes integrating heuristic signals-such as inter-node latency and topology mapping-into the baseline without introducing a centralized point of failure.

3. Dynamic Governance of λ. It is planned to model an adaptive control loop where λ acts as an automated response to network stress. For instance, if the protocol detects high stake concentration or low participation, it could automatically increase λ to force greater committee diversity, provided the safety condition $\overline{\alpha} > \overline{\rho}$ holds.

4. Hybridization. Finally, intention is to combine the λ-mixture with low-variance assignment schemes [14]. By applying low-variance sampling over the mixed distribution $q(\lambda)$, anticipating further tightening of the tail bounds, potentially allowing for smaller committee sizes without compromising security.

8. Summary

This paper introduces a λ-mixture mechanism for PoS committee selection, addressing the risks of stake concentration by affinely blending stake weights w_i with a bounded baseline u_i. The model defines the selection probability as $q_i(\lambda) = (1-\lambda)w_i + \lambda u_i$, yielding an expected adversarial seat count $\mu(\lambda)$ that decreases linearly and safety margins that widen affinely whenever the adversary is more concentrated than the baseline.

Our theoretical analysis (Section 4) derives closed-form exponential Chernoff and Hoeffding bounds, proving that even modest mixing significantly shrinks the probability of committee capture in the tails. Empirically (Section 5), simulations using historical snapshots from ten production PoS networks demonstrate a $\approx 24\%$ reduction in expected adversarial seats and an improvement of over 2 orders of magnitude in worst-case tail guarantees at $\lambda = 0.3$ and $\alpha = 0.25$.

Crucially, this mechanism is framed within the context of *identified validator sets* (e.g., consortium or high-compliance networks). Under the assumption of entity attribution, formally is shown that the mechanism is incentive-compatible and that stake-splitting is a dominated strategy. This establishes the λ-mixture as a powerful, tunable control surface for protocol governance, creating a mathematical foundation that future work can extend to fully permissionless environments.

Key Achievements:

- *Tunable Safety:* A single parameter λ provides linear control over adversarial presence without altering committee size or finality logic.
- *Rigorous Guarantees:* Proven exponential suppression of tail risks (probability of Byzantine majority) using standard concentration inequalities.
- *Conditional Incentives:* Formal demonstration that stake-splitting is economically irrational under the assumption of entity-level penalties.
- *Empirical Validation:* Reproducible quantification of risk reduction across diverse mainnet topologies.

A. Appendix full proofs (Chernoff via MGF, penalties, and sampling schemes)

A.1. Setup and notation

Let n validators have stake weights $w_i > 0$ with $\sum_{i=1}^{n} w_i = 1$. Let $A \subseteq \{1,\ldots,n\}$ be the adversarial set with $\alpha = \sum_{i\in A} w_i$. A normalized baseline $u = (u_1,\ldots,u_n)$ satisfies $\sum_i u_i = 1$ and $\rho = \sum_{i\in A} u_i$. The mixed selection weights for each committee seat are

$$q_i(\lambda) = (1-\lambda)w_i + \lambda u_i, \quad \lambda \in [0,1].$$

Write X for the number of adversarial seats in a committee of size m. Define

$$p(\lambda) = (1-\lambda)\alpha + \lambda\rho, \quad \mu(\lambda) = E[X] = m\,p(\lambda).$$

Pre-normalization penalties with intensity $\gamma \in [0,1]$ and penalty scores $\pi_i \in [0,1]$:

$$\widetilde{w_i} = (1-\gamma\pi_i)w_i, \quad \widetilde{u_i} = (1-\gamma\pi_i)u_i,$$

followed by renormalization to

$$\overline{w}_i = \widetilde{w_i}\Big/\textstyle\sum_j \widetilde{w_j} \quad \text{and} \quad \overline{u}_i = \widetilde{u_i}\Big/\textstyle\sum_j \widetilde{u_j}.$$

Then the operative weights are

$$q_i(\lambda,\gamma) = (1-\lambda)\overline{w}_i + \lambda\overline{u}_i,$$

with adversarial masses

$$\overline{\alpha} = \sum_{i\in A}\overline{w}_i \quad \text{and} \quad \overline{\rho} = \sum_{i\in A}\overline{u}_i,$$

hence

$$\mu(\lambda,\gamma) = m\big[(1-\lambda)\overline{\alpha} + \lambda\overline{\rho}\big].$$

A.2. Chernoff upper tails for Poisson-binomial committees

Model each seat $j = 1,\ldots,m$ by an indicator Y that the seat is adversarial. Under independent draws (with replacement),

$$\Pr\left(Y_j = 1\right) = p(\lambda) \quad \text{(aggregated view)},$$

or more generally Poisson-binomial with seat-specific success probabilities p_j satisfying $\sum_j p_j = \mu$. In either case, $X = \sum_{j=1}^{m} Y_j$, $\mu = E[X]$. For any $t > 0$, independence gives

$$E\left[e^{tX}\right] = \prod_{j=1}^{m} E\left[e^{tY_j}\right] = \prod_{j=1}^{m}\left(1 - p_j + p_j e^t\right).$$

Hence

$$P(X \geqslant \tau) \leqslant \exp\Bigg(-t\tau + \sum_{j=1}^{m}\ln\left(1 - p_j + p_j e^t\right)\Bigg).$$

Optimizing over $t > 0$ yields the Cramér–Chernoff bound

$$P(X \geqslant \tau) \leqslant \exp\left(-\sup_{t>0} \left\{ t\tau - \sum_{j=1}^{m} \ln\left(1 - p_j + p_j e^t\right) \right\} \right).$$

When $p_j \equiv p$ (binomial case), $\sum_j \ln\left(1 - p + pe^t\right) = m \ln\left(1 - p + pe^t\right)$ and the optimizer satisfies

$$\tau = m \frac{p\, e^t}{1 - p + p\, e^t} = \mu \frac{e^t}{1 - p + p\, e^t}.$$

Writing $\tau = (1+\delta)\mu$ gives $t^* = \ln(1+\delta)$, and the standard multiplicative Chernoff form

$$P\left[X \geqslant (1+\delta)\mu\right] \leqslant \exp\left(-\mu\phi(1+\delta)\right), \quad \phi(x) = x \ln x - x + 1.$$

For our mixture with replacement,

$$\mu = \mu(\lambda) = m\, p(\lambda), \quad \text{so for any} \quad \tau > \mu(\lambda),$$

$$P[X \geqslant \tau] \leqslant \exp\left(-\mu(\lambda)\, \phi\left(\tfrac{\tau}{\mu(\lambda)}\right)\right).$$

Penalties $\gamma > 0$ after renormalization. All steps remain valid upon replacing $p(\lambda)$ by

$$p(\lambda, \gamma) = (1-\lambda)\, \overline{\alpha} + \lambda\, \overline{\rho}$$

and $\mu(\lambda)$ by $\mu(\lambda, \gamma)$. The MGF derivation depends only on the success probabilities per seat, not on their provenance. Therefore the same Chernoff bound holds verbatim for $\gamma > 0$.

A.3. Hoeffding's (distribution-free) tail

Since $0 \leqslant Y_j \leqslant 1$, Hoeffding's inequality gives for any $\tau > \mu$

$$P(X \geqslant \tau) \leqslant \exp\left(-\tfrac{2(\tau-\mu)^2}{m}\right).$$

This requires neither identical p_j nor any particular generation mechanism, hence applies to Poisson-binomial and to our mixed model (with or without penalties) immediately.

A.4. Sampling without replacement (hypergeometric case)

If the committee is sampled without replacement from a fixed set with adversarial fraction $p(\lambda)$, the selection indicators are negatively associated. Classical results imply that hypergeometric tails are no larger than binomial tails with the same mean and m (variance is smaller). Consequently, both the Chernoff and Hoeffding [29,30] bounds derived above remain conservative for without-replacement sampling in our setting.

A.5. Incentive compatibility under entity attribution

Proposition 1. Under the assumption of Entity Resolution (Permissioned/Identified settings), stake-splitting is a weakly dominated strategy.

Proof. Let Entity E control total stake S_E and total baseline weight U_E. The entity's total selection probability mass is $Q_E(\lambda) = (1-\lambda)S_E + \lambda U_E$. If the entity splits into k identities, the protocol (knowing the entity) ensures the sum of the new baseline weights equals U_E. Thus, the total selection mass remains invariant:

$$\sum_{i=1}^{k} q_i = (1-\lambda) \sum_{i=1}^{k} s_i + \lambda_{i=1}^{k} \sum_{i=1}^{k} u_i = (1-\lambda)S_E + \lambda U_E.$$

Since splitting incurs operational costs without increasing selection probability, rational actors will not split. ■

A.6. Asymptotic tightness (binomial/Poisson-binomial)

For $X \sim \mathrm{Bin}(m, p)$ and any fixed $\delta > 0$,

$$P\left[X \geqslant (1+\delta)\mu\right] = \Theta\left(\tfrac{1}{\sqrt{m}}\right) \exp\left(-\mu\phi(1+\delta)\right), \quad m \to \infty,$$

where $\mu = m\, p$ and ϕ is the Chernoff rate function above. Thus Chernoff's exponent is sharp; the exact probability differs by at most a polynomial prefactor (typically $m^{-1/2}$) in the large-deviation regime.

A more explicit local-limit-style lower bound holds for integer $\tau \geqslant \mu$:

$$P(X \geqslant \tau) \geqslant \frac{1}{(1+\delta)\sqrt{2\pi\mu(1+\delta)(1-p)}} \exp\left(-\mu\phi(1+\delta)\right),$$

valid for sufficiently large μ) (moderate deviations). The same exponential rate ($\phi(1+\delta)$) appears on both sides, demonstrating tightness. For Poisson-binomial $X = \sum_j Y_j$ with means p_j, Cramér's theorem yields the large-deviation rate as the Legendre transform of $\Lambda(t) = \sum_j \ln\left(1 - p_j + p_j e^t\right)$.

When the p_j are not too dispersed, the binomial-style rate with $\mu = \sum_j p_j$ is a good proxy; in any case, the Chernoff construction already optimizes over T and remains the canonical exponential upper bound.

A.7. Practical reading of tightness

In our experiments, for large networks and moderate λ, empirical Monte Carlo with $T \in \left[2 \times 10^4, 10^6\right]$ often observes zero unsafe samples while the Chernoff bound reports a small but positive probability (e.g., $10^{-6} - 10^{-8}$). This is exactly expected: the true probability sits below the conservative bound, and once it is smaller than $1/T$ Monte Carlo loses resolution while the bound continues to certify safety. Conversely, in smaller networks or at higher α close to 1/3, empirical frequencies are measurable and still lie below the bound, confirming its conservatism

[1] Bentov I., Gabizon A., Mizrahi A. Cryptocurrencies Without Proof of Work. Financial Cryptography and Data Security. 142–157 (2016).

[2] Nakamoto S. Bitcoin: A Peer-to-Peer Electronic Cash System (2008).

[3] Micali S., Rabin M., Vadhan S. Verifiable random functions. 40th Annual Symposium on Foundations of Computer Science (Cat. No.99CB37039). 120–130 (1999).

[4] Hoeffding W. Probability Inequalities for Large Sums of Bounded Random Variables. Journal of the American Statistical Association. **58** (301), 13–30 (1963).

[5] Douceur J. R. The Sybil Attack. International Workshop on Peer-to-Peer Systems (IPTPS). 251–260 (2002).

[6] Leonardos S., Reijsbergen D., Piliouras G. Weighted Voting on the Blockchain: Improving Consensus in Proof of Stake Protocols. International Journal of Network Management. **31** (2), e2128 (2021).

[7] Gilad Y., Hemo R., Micali S., Vlachos G., Zeldovich N. Algorand: Scaling Byzantine Agreements for Cryptocurrencies. SOSP '17: Proceedings of the 26th Symposium on Operating Systems Principles. 51–68 (2017).

[8] Goldberg S., Reyzin L., Papadopoulos D., Včelák J. Verifiable Random Functions (VRFs). Internet Research Task Force (IRTF), RFC 9381 (2023).

[9] Buterin V., Griffith V. Casper the Friendly Finality Gadget. Preprint arXiv:1710.09437 (2017).

[10] Castro M., Liskov B. Practical Byzantine Fault Tolerance. 3rd Symposium on Operating Systems Design and Implementation (OSDI 99). 173–186 (1999).

[11] Akbar N. A., Muneer A., ElHakim N., Fati S. M. Distributed Hybrid Double-Spending Attack Prevention Mechanism for Proof-of-Work and Proof-of-Stake Blockchain Consensuses. Future Internet. **13** (11), 285 (2021).

[12] Reijsbergen D., Szalachowski P., Ke J., Li Z., Zhou J. LaKSA: A Probabilistic Proof-of-Stake Protocol. Proceedings of the Network and Distributed System Security Symposium (NDSS). 1–18 (2021).

[13] Windiatmaja J. H., Hanggoro D., Salman M., Sari R. F. PoIR: A Node Selection Mechanism in Reputation-Based Blockchain Consensus Using Bidirectional LSTM Regression Model. Computers, Materials & Continua. **77** (2), 2309–2332 (2023).

[14] Gaži S., Kiayias A., Russell P. Fait Accompli Committee Selection: Improving the Size-Security Tradeoff of Stake-Based Committees. CCS '23: Proceedings of the 2023 ACM SIGSAC Conference on Computer and Communications Security. 845–858 (2023).

[15] Klenke A., Mattner L. Stochastic ordering of classical discrete distributions. Advances in Applied Probability. **42** (2), 392–410 (2010).

[16] Wood G. Ethereum: A Secure Decentralised Generalised Transaction Ledger. Ethereum Project Yellow Paper (2014).

[17] Teerapabolarn K. Binomial approximation for a sum of independent hypergeometric random variables. Global Journal of Pure and Applied Mathematics. **11** (4), 1967–1970 (2015).

[18] Metropolis N., Ulam S. The Monte Carlo Method. Journal of the American Statistical Association. **44** (247), 335–341 (1949).

[19] Ava Labs. Avalanche Platform Documentation (2020).

[20] Celestia Foundation. Celestia: A Modular Data Availability Network. Technical Specification v1.0 (2023).

[21] Interchain Foundation. Cosmos Hub Documentation: Tendermint and Staking (2019).

[22] Injective Labs. Injective Protocol: A Layer-1 Optimized for DeFi. Whitepaper (2021).

[23] NEAR Foundation. NEAR Protocol: Nightshade Sharding and PoS Mechanism. Technical Paper (2020).

[24] Osmosis Labs. Osmosis: AMM Chain on the Cosmos SDK. Protocol Documentation (2021).

[25] Sei Labs. Sei Network: Parallelized EVM Layer-1. Technical Whitepaper (2023).

[26] Solana Foundation. Solana: A New Architecture for High Performance Blockchain. Whitepaper v0.8 (2018).

[27] Mysten Labs. Sui: A High-Throughput Object-Centric Blockchain. Whitepaper (2022).

[28] Tezos Foundation. Tezos: Self-Amending Cryptographic Ledger. Whitepaper (2018).

[29] Chernoff H. Measure of Asymptotic Efficiency for Tests of a Hypothesis Based on the sum of Observations. Annals of Mathematical Statistics. **23** (4), 493–507 (1952).

[30] Cramér H. On a new limit theorem in probability theory (Sur un nouveau théorème-limite de la théorie des probabilités). Actualités scientifiques et industrielles. **736**, 2–23 (1938).

Добір комітету зі змішаними вагами в системах Proof-of-Stake на основі λ-змішаної моделі ваг з експоненціальними оцінками хвостових ймовірностей та узгодженістю стимулів

Соломка І. Р., Любінський Б. Б., Пеняк Б. О.

Національний університет "Львівська політехніка",
вул. С. Бандери, 12, 79013, Львів, Україна

У блокчейнах на основі Proof-of-Stake (PoS) комітети валідаторів зазвичай формуються пропорційно до розміру їхнього стейку, що робить безпеку мережі залежною від великих гравців та концентрації капіталу. За таких умов, жеребкування, що базується суто на стейку, іноді може призвести до створення комітету, де частка супротивника перевищує поріг безпеки, навіть якщо його загальна частка в мережі менша за одну третину. Робота пропонує простий механізм добору зі змішаними вагами, який поєднує стейк із обмеженим базовим розподілом за допомогою єдиного параметра змішування λ. Цей підхід не змінює розмір комітету, схему винагород чи алгоритми VRF-сортування, однак "розбавляє" вагу надмірно сконцентрованих позицій. В роботі доводиться, що коли частка стейку супротивника більша за його частку в базовому розподілі, очікувана кількість його місць у комітеті спадає лінійно зі зростанням λ, а стандартні оцінки концентрації показують експоненційне зменшення ймовірності захоплення комітету. Експерименти на десяти діючих PoS-мережах свідчать, що помірне змішування (при $\lambda = 0.3$) зменшує очікувану кількість місць супротивника приблизно на чверть та посилює гарантії безпеки для найгірших сценаріїв на кілька порядків.

Ключові слова: *Proof-of-Stake: добір комітету; змішані ваги; системи репутації; узгодженість стимулів; хвостові ймовірності.*