ISSN 2306-5478

УДК 004.75:519.8:519.2 https://doi.org/10.31713/vt2202610

Solomka I. R.[1] ([1:ORCID: 0009-0004-3705-8801])
PhD Student, Corresponding author
Liubinskyi B. B.[1] ([1:ORCID: 0000-0002-0715-8068])
PhD, Associate Professor

[1]Department of Applied Mathematics, Institute of Applied Mathematics and Fundamental Sciences, Lviv Polytechnic National University

A METHOD FOR ASSESSING COMMITTEE-CAPTURE RISK UNDER STAKE CONCENTRATION IN PROOF-OF-STAKE

This paper studies committee selection in committee-based Proof-of-Stake networks under increasing stake concentration. The problem is how to assess committee-capture risk, that is, the probability that a monitored coalition obtains a capturing share of committee seats, directly from protocol state. In existing practice this risk is often discussed qualitatively, so committee parameters such as size, threshold, and admissible corrective mixing are chosen heuristically, and the point at which a chosen risk tolerance is no longer achievable may be crossed without an explicit warning. To address this problem, the paper proposes a state-auditable method for assessing committee-capture risk under stake concentration. The method combines an exact risk certificate for a monitored coalition with a feasibility/saturation boundary that shows when bounded corrective mixing can still keep risk within a chosen budget and when it no longer can. The method is implemented in a reproducible Proof-of-Stake proof of concept. The results show that the proposed certificate provides an informative estimate of committee-capture risk in the operating regime of interest. The practical value of the method lies in monitoring, governance, and parameter audit for committee-based Proof-of-Stake systems. It is intended as a risk-assessment and decision-support instrument rather than as a universal network-security oracle. The proof of concept demonstrates computability from protocol state, the **certificate's** behaviour under controlled concentration drift, and an interpretable warning at budget violation or infeasibility. Its conditions of applicability are a non-strategic concentration process, a monitored top-k coalition, and the treatment of the binomial estimate as a conservative upper bound on the actual

without-replacement committee draw.
Keywords: Proof-of-Stake, committee selection, committee capture, stake concentration, risk assessment, capture certificate, risk budget, feasibility boundary, on-chain auditability, math modeling.

1. Introduction. Proof-of-Stake secures a large and growing share of public blockchain value. The committee-based variant, which delegates each round to a sampled subset of validators, is the dominant design for high-throughput chains.

The safety of such a chain rests on the rule that maps a validator's stake to its chance of selection. Because the committee is a sample, a coalition can cross a committee threshold in some round without a global majority of stake, with a risk that depends jointly on the selection rule, the committee size, and the threshold. This committee-threshold crossing is distinct from three neighbouring notions: stake concentration, the control of many validators by one operator, and mere committee overrepresentation. Only the crossing is a direct safety event for a sampled committee, and it is the event this paper concerns.

The central question is not how to prevent such crossing, but how to assess and signal its risk from observable protocol state. A committee size and a capture threshold are not self-explanatory: the same parameters are safe under one stake distribution and unsafe under another, and concentration changes over time, so a margin judged safe at design time can erode silently.

This paper does not claim that operating networks are near collapse from concentration; it addresses one auditable component of committee security. An operator or governance process needs an indicator that recomputes from current state, raises a red flag when a declared tolerance is no longer met, and is computable without external identity infrastructure, so that parameter decisions rest on a quantity rather than on intuition.

Therefore, research on state-auditable assessment of committee-capture risk under stake concentration in Proof-of-Stake is timely and important.

2. Analysis of literature data and the problem statement. The review proceeds through three strands toward the operational gap the present method fills: static stake-weighting and decentralisation transforms, committee election and capture probability, and adaptive or prior-

 ISSN 2306-5478

author approaches. Two economic analyses fix the setting. [1] establishes Proof-of-Stake as agreement without the computational waste of Proof-of-Work, so stake is the scarce decision resource. Authors [4] show that under reinvested rewards the share distribution behaves as a martingale and need not concentrate, so concentration is a contingent regime that motivates a response which engages conditionally rather than on a fixed schedule.

Static weighting and decentralisation transforms form the first strand. Motepalli and Jacobsen show that concave maps such as the Square Root and Logarithmic Stake Weight raise the Nakamoto coefficient and lower the Gini coefficient relative to linear weighting [2], and that operating chains vary widely in realised decentralization [3]; e-PoS reshapes weighting so that concentrated stake yields less committee influence [6]. These fix or compare a weighting map for distributional fairness, and leave open how much correction to apply at a given network state against a stated target.

The second strand concerns committee formation and stability. A verifiably secure committee-election rule holds formation to explicit security and proportionality criteria, but for a fixed election instance [5]; a stability analysis treats concentration as a control input that warrants adaptation as it grows [7]. Closer still, [12] improves the size–security tradeoff by seating large stakeholders deterministically from the current stake distribution, and a deterministic-bounds analysis caps a committee's adversarial share through saturation limits tied to the weight distribution [13]. These are design-time selection results: none recomputes a committee-capture probability from runtime state, defines a budget on it, or exposes a feasibility boundary on a corrective control.

The third strand adapts selection using non-stake state. Reputation-based node selection [8] and grouped fairness-and-Sybil scoring [9] add information beyond stake, but rest on off-chain features or on infrastructure heavier than a minimal permissionless protocol assumes; grounding the non-stake component in on-chain age, as the persistence baseline does, avoids that dependence. The authors' predecessors are closest: a static mixed-weight rule lowers capture probability when the coalition is more concentrated in stake than in the baseline [10], and a service-age extension makes the correction adaptive and resistant to fresh-stake splitting [11]. Both react to a decentralisation signal rather than to a declared risk budget.

A common gap remains that the decentralisation indices these

lines rely on, chiefly the Nakamoto and Gini coefficients, measure stake concentration, not the probability that a sampled committee of a given size is captured. No prior work supplies a quantity a protocol can recompute from its own state to certify that probability, a closed-form boundary that says when a bounded correction can still hold it within a declared budget, and a warning when it cannot. The present paper fills this gap with an exact on-chain capture certificate, a closed-form account of the minimal mixing it admits, and a feasibility boundary readable from protocol state.

The gap left open above is the absence of an auditable account of committee-capture risk and of the limit of any bounded correction. The object of study is accordingly narrowed: not stake concentration in general, but the committee-threshold crossing risk of a monitored coalition under a specified sampling rule. The monitored coalition A is a declared set of validators, taken here as the top-k by stake, abstracted as a Bernoulli source whose committee-selection mass the rule fixes. With committee size m and capture threshold ϑ, capture is the event $X_A \geq q$ in which the coalition takes at least $q = \lceil \vartheta m \rceil$ of the m seats; the risk to be governed is $\Pr[X_A \geq q]$, and a budget declares a tolerance $\Pr[X_A \geq q] \leq \epsilon$.

Three things are then required. First, a certificate that expresses $\Pr[X_A \geq q]$ from protocol state, on a scale where a tolerance can meaningfully be set. Second, a closed-form account of both the smallest corrective mixing that keeps the certificate within ϵ and the exact concentration threshold beyond which none does. Third, this feasibility – saturation boundary exposed as a first-class output, so that the transition from a safe to a violated or infeasible regime is detected rather than inferred. The earlier controllers reacted to a decentralisation signal rather than to such a declared event, and the earlier tail analyses left the certificate implicit; neither supplies the auditable quantity and the explicit boundary this requires.

3. The aim and objectives of the study. The aim of the study is to develop a method for assessing committee-capture risk under stake concentration in Proof-of-Stake: an exact risk certificate computable from protocol state, together with a closed-form boundary that delimits when a bounded corrective mixing can hold the risk within a declared budget. This will make it possible for an operating chain and its

 ISSN 2306-5478

governance to audit committee-capture risk from protocol state, and to recognise the concentration level beyond which committee size or cap must be changed structurally, without external identification or manual reconfiguration.

To achieve this aim, the following objectives were set:

- to express committee-capture risk as a certificate computable from on-chain state, with the minimal mixing that holds it within a declared budget;
- to establish, in closed form, the concentration threshold up to which the budget is attainable and the saturation point beyond it;
- to build a reproducible drift instrument and evaluate the certificate's informativeness, its on-chain reproduction, and the feasibility saturation detection.

4. Materials and methods. The object of research is the process of committee selection in permissionless Proof-of-Stake under a varying stake distribution. The main hypothesis is that committee-capture risk admits an exact, state-computable certificate whose monotonicity in the mixing level yields a closed-form minimal intervention and an explicit feasibility threshold, so that the concentration regime in which a bounded mixing meets a declared budget is separated from the regime in which none does. The study adopts three assumptions: the binomial tail conservatively bounds the actual without-replacement committee draw; the monitored coalition is a fixed top-k stake prefix; and the concentration process is non-strategic. The accepted simplifications are a deterministic per-epoch controller whose only state is the previous mixing level, and evaluation on a controlled localnet of twenty-four validators rather than a production network.

The method is a per-epoch assessment over on-chain state

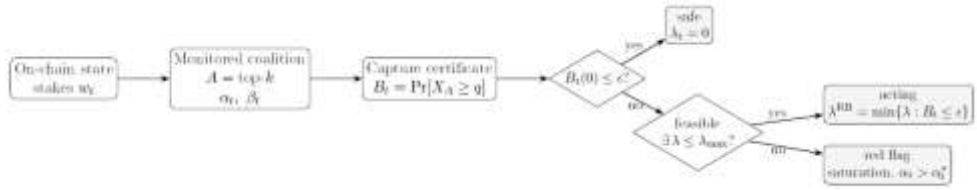

Figure 1 Per-epoch assessment over on-chain state

Algorithm of the method is shown on Fig 1. Each epoch it reads the validator stakes, identifies the monitored coalition A as the top-k by stake, and computes the committee-capture certificate B_t from the

coalition's masses. It then tests the certificate against the declared budget ϵ and reports one of three states. The state is safe when the budget holds at zero mixing. It is acting when a bounded corrective mixing λ^{RB} restores the budget. It is a red flag when no admissible mixing does, that is, when the coalition's stake has crossed the saturation threshold. The corrective mixing draws on a baseline. The persistence baseline of [11] is retained for the inherited fresh-entry path, while the concentration baseline analysed here addresses drift. Learning the regime online is left to future work.

For an epoch-indexed validator set $V_t = \{1, \dots, N_t\}$ with normalised stake $w_i(t) \geq 0, \sum_i w_i = 1$, and protocol age $\tau_i(t)$, the model carries two normalised baselines. The persistence baseline passes age through a concave capped transform $\tilde{\tau}_i = \max(1, \min(\tau_i, \tau_{\max})^{\zeta}), \zeta \in (0,1]$, and weights stake by it,

$$u_i^{\mathrm{P}}(t) = \frac{\tilde{\tau}_i(t)\, w_i(t)}{\sum_{j \in V_t} \tilde{\tau}_j(t)\, w_j(t)}, \qquad (1)$$

so a freshly minted identity receives almost no mass however much stake it posts. Under drift, however, $\tilde{\tau}_i$ is nearly constant among mature validators and (1) collapses to a stake-proportional rule that moves with concentration; the model therefore adds a concentration baseline whose mass is a capped function of stake,

$$u_i^{\mathrm{C}}(t) = \frac{\min(w_i(t), c)}{\sum_{j \in V_t} \min\left(w_j(t), c\right)}, \qquad c \in (0,1), \qquad (2)$$

where a validator above the cap c contributes a constant, so its mass stops growing with its share. Any concave map shares this anti-concentration sign; the Square Root Stake Weight transform [9] is one such choice, and the uniform baseline of [10] is the degenerate case $c \to 0$. The active baseline $u_i(t)$ is u_i^{P} under fresh-entry stress and u_i^{C} under drift, and the committee-selection weights are the affine mixture

$$q_i(t; \lambda_t) = (1 - \lambda_t) w_i(t) + \lambda_t u_i(t), \qquad \lambda_t \in [0, \lambda_{\max}]. \qquad (3)$$

 ISSN 2306-5478

The risk path monitors a fixed top-k coalition $A_t = \{(1), \dots, (k)\}$ ordered by non-increasing stake, with stake mass $\alpha_t = \sum_{i \in A_t} w_i(t)$ and baseline mass $\beta_t = \sum_{i \in A_t} u_i(t)$, so the coalition's committee-selection mass is

$$p_t(\lambda) = (1 - \lambda)\alpha_t + \lambda\beta_t. \quad (4)$$

The path is meaningful precisely when $\beta_t < \alpha_t$, that is, the active baseline is less concentrated on the coalition than stake is; the concentration baseline delivers this inequality under drift while the persistence baseline does not, which is the routing requirement stated above. With committee size m_t and capture threshold ϑ (BFT default $1/3$, majority proxy $1/2$), the coalition is abstracted as a Bernoulli source and the exact binomial upper tail serves as a policy-facing certificate. This choice is deliberate: Chernoff-style exponential tails are analytically convenient but are driven by a Kullback–Leibler exponent and may be too loose around the moderate probabilities relevant to governance thresholds. Accordingly, the certificate is defined as

$$B_t(\lambda) = \Pr[X \geq q], \qquad X \sim \text{Binomial}\big(m_t, p_t(\lambda)\big), \qquad q = \lceil \vartheta\, m_t \rceil, \quad (5)$$

the exact upper-tail probability that the coalition takes a capturing share of the committee when each of the m_t seats is an independent Bernoulli draw with mass $p_t(\lambda)$. By Hoeffding's reduction the without-replacement committee draw is no less concentrated than this binomial source, so B_t is a conservative upper bound on the realised capture.

Fixing a budget $\epsilon \in (0,1)$, the risk target is the smallest feasible mixing when one exists and the best-effort minimiser otherwise,

$$\lambda_t^{\text{RB}} = \begin{cases} \inf\{\lambda \in [0, \lambda_{\max}] : B_t(\lambda) \leq \epsilon\}, & B_t(\lambda_{\max}) \leq \epsilon, \\ \arg\min_{\lambda \in [0, \lambda_{\max}]} B_t(\lambda), & \text{otherwise}, \end{cases} \quad (6)$$

which keeps λ_t^{RB} defined even when severe concentration makes the budget unreachable. The tail $p \mapsto \Pr[\text{Binomial}(m_t, p) \geq q]$ is continuous and strictly increasing on $(0,1)$, so the budget threshold, the feasibility frontier and the saturation point are exact analytical objects defined through a single monotone root.

Proposition 1 (Monotone certificate and closed-form target). *If* $\beta_t < \alpha_t$,

then $p_t(\cdot)$ *is affine strictly decreasing and* $B_t(\cdot)$ *is continuous non-increasing on* $[0, \lambda_{\max}]$, *and* $B_t(\lambda) \le \epsilon \Leftrightarrow p_t(\lambda) \le p_t^*$, *where* $p_t^* \in (0,1)$ *is the unique root of* $\Pr[\text{Binomial}(m_t, p) \ge q] = \epsilon$. *Hence the feasible set is* $[\underline{\lambda}_t, \lambda_{\max}]$ *with*

$$\underline{\lambda}_t = \frac{\alpha_t - p_t^*}{\alpha_t - \beta_t}, \tag{7}$$

and the implemented target is $\lambda_t^{\mathrm{RB}} = \min\{\max\{\underline{\lambda}_t, 0\}, \lambda_{\max}\}$.

Theorem 1 (Feasibility frontier and saturation). *Let* $g_t = \alpha_t - \beta_t > 0$. *A budget-meeting admissible mixing exists iff* $p_t(\lambda_{\max}) \le p_t^*$.

(i) For fixed $g_t, m_t, \vartheta, \epsilon, \lambda_{\max}$ *this holds iff* $\alpha_t \le \alpha_t^* := p_t^* + \lambda_{\max} g_t$, *so the saturation flag* $sat_t = 1\{B_t(\lambda_t^{RB}) \le \epsilon\}$ *is* 1 *for* $\alpha_t \le \alpha_t^*$ *and* 0 *above it.*

(ii) When infeasible, λ_t^{RB} *minimises* B_t *over* $[0, \lambda_{\max}]$ *(equal to* $\lambda_{\max}$ *when* $\beta_t < \alpha_t$ *and to* 0 *when* $\beta_t \ge \alpha_t$*), so the controller never moves* λ *in the capture-increasing direction.*

Proof sketch. $p \mapsto \Pr[\text{Binomial}(m_t, p) \ge q]$ is strictly increasing on $(0,1)$ with limits 0 and 1, so $B_t \le \epsilon \Leftrightarrow p_t \le p_t^*$ for a unique p_t^* (intermediate value theorem). Solving $p_t(\lambda) = p_t^*$ gives (7); substituting $\beta_t = \alpha_t - g_t$ into $p_t(\lambda_{\max}) \le p_t^*$ gives (i), and monotonicity gives the minimisers in (ii).

The saturation flag thus marks the exact concentration threshold α_t^* beyond which a structural change (a tighter cap c or a higher $\lambda_{\max}$) is required. The on-chain keeper computes λ_t^{RB} by a fixed-point grid search rather than evaluating (7) directly, to keep the update in integer state-machine arithmetic; by Proposition 1 the result lies within one grid step of $\underline{\lambda}_t$, as result section confirms on-chain.

The model retains the signal-reactive logic of [11]: a decentralisation indicator $s_t \in [0,1]$ gives $\lambda_t^{\mathrm{sig}} = s_t \lambda_{\max}$, combined conservatively with the risk target as

$$\lambda_t^{\star} = \max\left\{\lambda_t^{\mathrm{sig}}, \lambda_t^{\mathrm{RB}}\right\}, \tag{8}$$

 ISSN 2306-5478

and then passed through a bounded filter with a hysteresis floor to give the realised $\lambda_t \in [0, \lambda_{\max}]$. The only state carried across epochs is λ_{t-1}, so the controller is a memoryless map inside a deterministic transition.

The evaluation tests whether the method detects the loss of committee-security margin under a controlled concentration drift. It runs on a Cosmos-SDK localnet of $N = 24$ validators with a single drift-pool delegator: at each epoch boundary the pool redelegates a fixed fraction $\rho = 0.03$ of bonded stake from a donor outside the top ranks to a receiver inside the top-k set, so the coalition mass α_t rises along one reproducible trajectory. Each epoch logs the coalition masses α_t, β_t, the certificate $B_t(0)$ and $B_t(\lambda_t)$, the closed-form target $\underline{\lambda}_t$, the realised mixing λ_t, and the saturation flag.

The drift carries the system through three operational phases. It is safe while the certificate stays within the budget at zero mixing, degrading while the budget is met only by a bounded corrective mixing, and violated once no admissible mixing meets it. A red flag is the entry into the violated phase, recorded by the saturation flag together with the certificate trajectory. The method is judged by whether it detects these transitions.

The matrix reports means over three seeds for four controllers on one trajectory: a $\lambda = 0$ reference, the signal-only controller of [11], and the risk-budget controller with the capped and with the concave baseline, under a competitive committee ($m = 9$ of $N = 24$, top-five coalition, $\vartheta = 1/3$, $\epsilon = 0.5$, $\lambda_{\max} = 0.55$). The seeds share the trajectory and differ only in the committee-draw realisation, so they check robustness rather than add scenarios. The capped and concave baselines are deployment variants of one certification framework, since the certificate, the closed-form law and the feasibility machinery are identical for either. Three targets are measured: whether the binomial certificate is informative where the Chernoff-Kullback-Leibler estimate is not; whether the deployed keeper reproduces the closed-form target of Proposition 1 within the grid resolution; and whether the feasibility frontier and the saturation flag of Theorem 1 detect the concentration threshold beyond which no admissible mixing meets the budget.

5. Results. On a competitive committee ($\boldsymbol{m = 9}$ of $\boldsymbol{N = 24}$, top-five coalition, $\boldsymbol{\vartheta = 1/3, \epsilon = 0.5, \lambda_{\max} = 0.55}$) the drift concentrates stake into the coalition over thirty epochs; the quantities below are means over three seeds, reported as measurements. The method carries the

scenario through the three operational phases and raises the red flag at the predicted point (Fig. 2, Table 1). The run is safe while the certificate stays within the budget at zero mixing, acting while a bounded mixing still holds it, and violated at epoch seven, where the coalition mass crosses the saturation threshold $\boldsymbol{\alpha_t^*}$ and the flag is raised. The uncontrolled certificate $\boldsymbol{B_t(0)}$ climbs from $\mathbf{0.37}$ to $\mathbf{0.84}$ across the drift and the controlled $\boldsymbol{B_t(\lambda_t)}$ from $\mathbf{0.37}$ to $\mathbf{0.72}$, both staying above the budget after saturation. The realised outcome confirms the warning: the **coalition's** Byzantine-threshold capture rises from about $\mathbf{0.23}$ to $\mathbf{0.75}$ along the same trajectory.

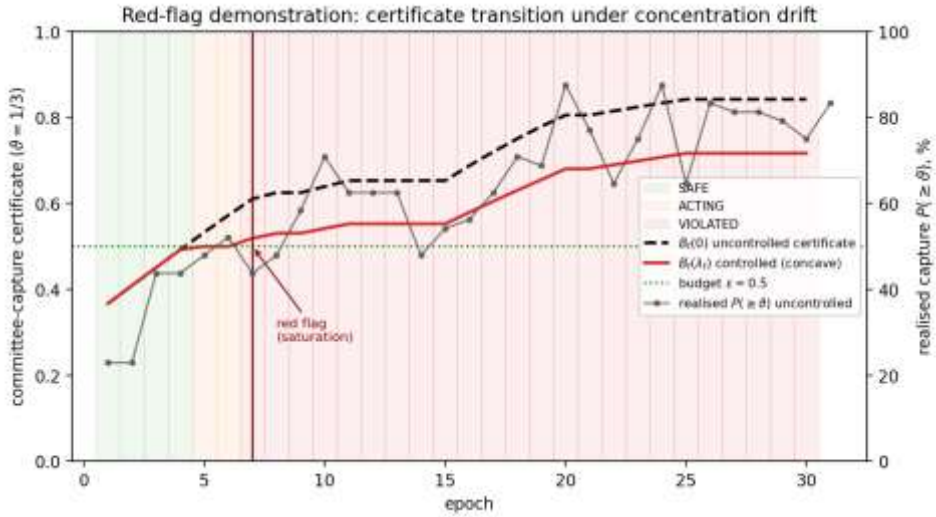

*Figure 2 Red-flag demonstration (**ϑ**=1/3, **ϵ**=0.5*

Fig. 2 demonstrates ($\vartheta = 1/3, \epsilon = 0.5$) the certificate $B_t(0)$ and $B_t(\lambda_t)$ rise with the coalition mass through the safe, acting, and violated phases (background), the vertical line marking saturation; the realised capture $P(\geq \vartheta)$ (right axis) follows.

Two derived quantities make the warning legible (Fig. 2, Table 1). The concentration headroom $\alpha_t^* - \alpha_t$ falls from 0.09 to zero at the red flag, while the stake-Gini is only 0.18 there: a global inequality index misses the top-k concentration that the coalition-keyed certificate detects.

Results in Table 1 shows as the coalition mass α_t grows, the headroom $\alpha_t^* - \alpha_t$ falls to zero and the certificate $B_t(0)$ crosses the

 ISSN 2306-5478

budget, taking the status from safe through acting to violated.

Table 1

Per-epoch status for the concave variant

epoch	α_t	headroom	$B_t(0)$	status
1	0.238	0.090	0.37	SAFE
5	0.298	0.030	0.53	ACTING
7	0.328	0.000	0.61	VIOLATED
15	0.346	0.000	0.65	VIOLATED
30	0.444	0.000	0.84	VIOLATED

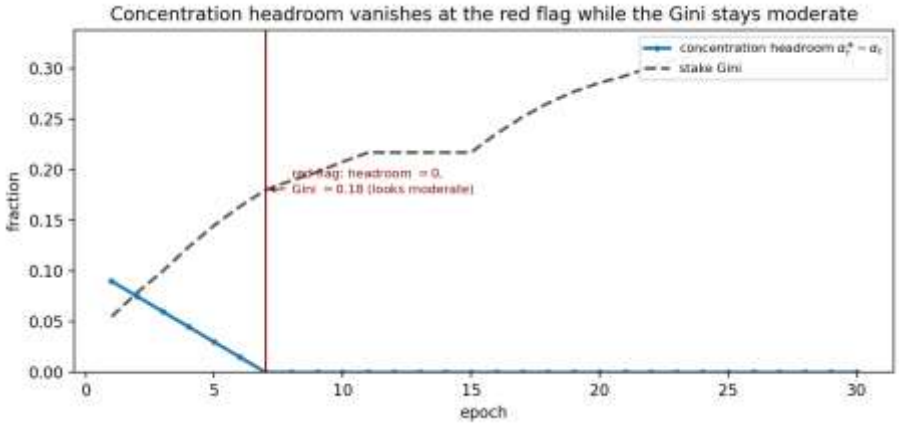

Figure 3 Concentration headroom

Concentration headroom on Fig 3 $\alpha_t^* - \alpha_t$ (solid) counts down to zero at the red flag, while the stake-Gini (dashed) is only 0.18 there.

Saturated-regime outcomes on Table 2 means over three seeds, $\alpha_t \approx 0.44$), $m = 9$, $N = 24$, top-five coalition, $\vartheta = 1/3$, $\epsilon = 0.5$. The mass β_t is defined only for the risk-budget controllers; the budget-satisfied fraction is over the full run.

Table 2

Saturated-regime outcomes

Controller	Mixing		Observed outcome			Budget
	$\bar{\lambda}$	β_t	seat share	$P(\geq 1/3)$	$P(\geq 1/2)$	sat. frac.
Reference	0.000	—	0.369	79.4%	9.9%	0.00
Signal	0.040	—	0.374	83.1%	9.1%	0.20
RB, capped	0.550	0.311	0.331	68.2%	6.0%	0.20
RB, concave	0.550	0.316	0.315	63.0%	4.7%	0.20

Byzantine-threshold capture probability (per cent) in the saturated regime ($\alpha \approx 0.44, m = 9$): the Chernoff–Kullback–Leibler bound, the binomial certificate, and the realised capture.

Table 3.

Byzantine-threshold capture probability (per cent)

Controller	Chernoff–KL	binomial cert.	realised capture
reference ($\lambda = 0$)	100.0	84.1	79.4
RB, capped	100.0	70.9	68.2
RB, concave	100.0	71.5	63.0

The deployed keeper realises the closed-form law: its mixing matches $\underline{\lambda}_t$ of Proposition 1 within one grid step, and its saturation flag matches the frontier of Theorem 1. The risk path saturates both variants at $\lambda_{\max} = 0.55$, while the signal-only controller barely moves (mean near 0.04) and tracks the reference (Fig. 4).

The corrective mixing's residual effect is modest and is reported as a measured quantity (Fig. 5, Table 2). In the saturated regime the coalition's Byzantine-threshold capture falls from 0.79 to 0.68 for the capped and 0.63 for the concave variant, and the realised seat share from 0.369 to 0.331 and 0.315; the two baselines are not robustly separated at three seeds.

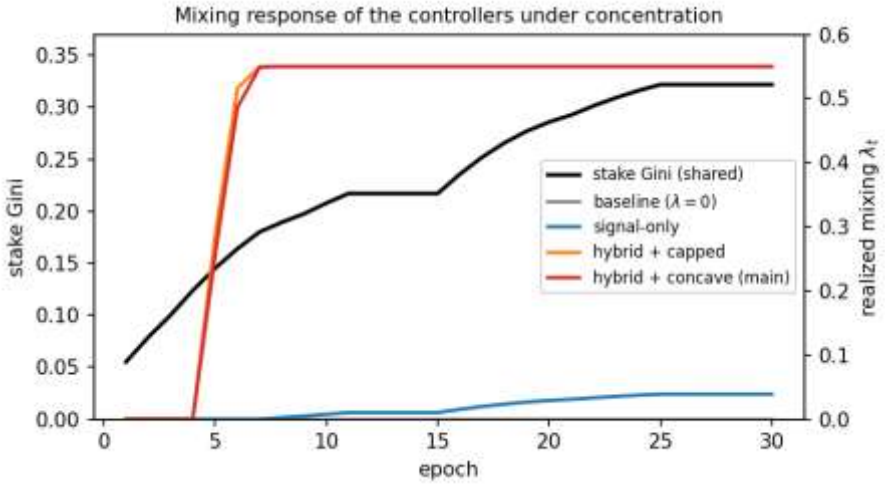

Figure 4 Responsiveness under concentration

 ISSN 2306-5478

As the stake-Gini (black, left axis) rises, the signal-only controller barely lifts λ (right axis) while both risk-budget variants engage the risk path and saturate at $\lambda_{\max} = 0.55$. The reference ($\lambda = 0$) coincides with the horizontal axis.

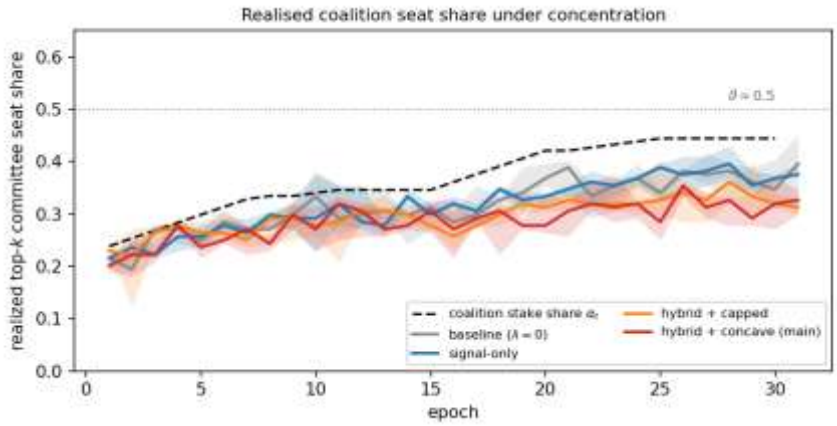

Figure 5 Realised top-five coalition seat share under concentration

Realised top-five coalition seat share under concentration, with the stake share α_t (dashed) for context; the risk-budget variants induce a modest measured reduction. Bands span three seeds.

6. Discussion. The obtained results are explained by the analytical structure of the certificate. Because the capture probability is the binomial tail (5), strictly monotone in the mixing level, the budget threshold and the saturation point are exact roots rather than fitted quantities; this is why the keeper reproduces the closed-form target within one grid step and the red flag falls at the predicted epoch (Fig. 2, Table 1). The same exactness explains why the certificate is informative where the Chernoff-Kullback-Leibler bound is vacuous (Table 3): the binomial tail is exact at the operating point, whereas the Chernoff bound saturates at one.

The distinctive feature against existing work is that the method reports a sampled committee-capture probability keyed to a monitored coalition, not a global concentration index. The decentralisation metrics of [2], [3] and the Nakamoto and Gini coefficients generally, can stay moderate while the committee budget is already breached (Fig. 3); the committee-election guarantees of [5] hold for a fixed instance rather than a drifting state; and the prior mixed-weight controllers [10] and [11] react to a decentralisation signal rather than to a declared budget.

The results hold within stated bounds. The certificate is a with-

replacement bound, conservative several-fold at the majority threshold and tight at the Byzantine threshold, so the budget ϵ must be calibrated to the realised scale and treated as a governance parameter. The reported outcomes are reproducible across the three seeds of the controlled localnet ($N = 24, m = 9$), which share one drift trajectory; their magnitude depends on the committee size and the per-member concentration of the coalition, and is established only for the non-strategic drift studied here.

A shortcoming distinct from these bounds is the limited reach of the corrective mixing: it acts only through the gap $\alpha_t - \beta_t$, which collapses when an adversary splits its stake across many validators below the cap, so the validator-level certificate then under-reports the risk. This can be removed by evaluating the certificate at the level of beneficial-owner entities rather than validators. A second shortcoming is the narrow empirical base of two baseline variants on a single trajectory, which a sweep over drift rates and committee sizes would widen.

The natural development is an entity-level certificate driven by an estimate of correlated control, which would recover the corrective gap that identity-splitting destroys. The difficulty is methodological: such an estimate is an adversarially evadable inference from on-chain provenance and off-chain correlation, with no ground-truth identity in a permissionless setting. A policy-aware adversary that times accumulation against the budget is a further open problem of the same kind.

7. Conclusions

An exact committee-capture certificate $B_t(\lambda) = \Pr[X_A \geq q]$, with the closed-form minimal mixing $\underline{\lambda}_t = (\alpha_t - p_t^*)/(\alpha_t - \beta_t)$ that holds it within a declared budget, was obtained and made computable from protocol state. Unlike the implicit tails of prior mixed-weight rules and the global decentralisation indices, it is a sampled-coalition capture probability: at the Byzantine threshold it reads 0.84, which is explained by the binomial tail being exact at the operating point while the Chernoff bound saturates.

 ISSN 2306-5478

The concentration threshold $\alpha_t^* = p_t^* + \lambda_{\max}(\alpha_t - \beta_t)$ up to which the budget is attainable, and the saturation point beyond it, were established in closed form (Proposition 1, Theorem 2). The distinctive feature is that feasibility is an analytical boundary rather than a tuning heuristic; the deployed keeper reproduced the target within one grid step, which follows from the strict monotonicity of the tail in the mixing level.

A reproducible Cosmos-SDK instrument was built and used to evaluate the method under concentration drift over three seeds. The method raised the red flag at the predicted epoch, where the concentration headroom reached zero while the stake-Gini was only 0.18, and the realised Byzantine-threshold capture rose from 0.23 to 0.75, confirming the warning. This is explained by the certificate keying on the monitored coalition's mass, which a global decentralisation index does not track.

Taken together, the results turn the otherwise silent erosion of committee-security margin into an auditable saturation signal that separates parameter tuning from structural redesign; extending the certificate from validators to beneficial-owner entities is the natural continuation.

1. Saleh F. (2021) Blockchain without Waste: Proof-of-Stake. *The Review of Financial Studies*. Vol. 34, No. 3. P. 1156–1190. https://doi.org/10.1093/rfs/hhaa075. 2. Motepalli S., Jacobsen H.-A. (2024) How Does Stake Distribution Influence Consensus? Analyzing Blockchain Decentralization. *2024 IEEE International Conference on Blockchain and Cryptocurrency (ICBC)*. P. 343–352. https://doi.org/10.1109/ICBC59979.2024.10634400. 3. Motepalli S., Jacobsen H.-A. (2025) Decentralization in PoS Blockchain Consensus: Quantification and Advancement. *IEEE Transactions on Network and Service Management*. Vol. 22, No. 4. P. 2930–2943. https://doi.org/10.1109/TNSM.2025.3561098. 4. Roșu I., Saleh F. (2021) Evolution of Shares in a Proof-of-Stake Cryptocurrency. *Management Science*. Vol. 67, No. 2. P. 661–672. https://doi.org/10.1287/mnsc.2020.3791. 5. Cevallos A., Stewart A. (2021) A Verifiably Secure and Proportional Committee Election Rule. *Proceedings of the 3rd ACM Conference on Advances in Financial Technologies (AFT '21)*. P. 29–42. https://doi.org/10.1145/3479722.3480988. 6. Saad M., Qin Z., Njilla L., Kamhoua C. A., Nyang D., Mohaisen D., Song D., Taduri S. (2021) e-PoS: Making Proof-of-Stake Decentralized and Fair. *IEEE Transactions on Parallel and*

Distributed Systems. Vol. 32, No. 8. P. 1961–1973. https://doi.org/10.1109/TPDS.2020.3048853. 7. Leporati A., Rovida L. (2024) Looking for Stability in Proof-of-Stake Based Consensus Mechanisms. *Blockchain: Research and Applications*. Vol. 5, No. 4. Art. 100222. https://doi.org/10.1016/j.bcra.2024.100222. 8. Windiatmaja J. H., Hanggoro D., Salman M., Sari R. F. (2023) PoIR: A Node Selection Mechanism in Reputation-Based Blockchain Consensus Using Bidirectional LSTM Regression Model. *Computers, Materials & Continua*. Vol. 77, No. 2. P. 2309–2339. https://doi.org/10.32604/CMC.2023.041152. 9. Zhang M., Liu M., Ding X., Wang Y., Li G. (2025) GPE-PoS: A Fair and Sybil-Resistant Proof of Stake Consensus. *Journal of Internet Technology*. P. 463–470. https://doi.org/10.70003/160792642025072604005. 10. Solomka I. R., Liubinskyi B. B., Peniak B. O. (2025) Mixed-Weight Committee Selection in Proof-of-Stake: Tunable Stake-Baseline Mixing with Exponential Tail Guarantees and Incentive Compatibility. *Mathematical Modeling and Computing*. Vol. 12, No. 4. P. 1320–1332. https://doi.org/10.23939/mmc2025.04.1320. 11. Solomka I., Liubinskyi B. (2026) Devising a Method for Dynamic Stabilization of PoS Consensus Using Adaptive Weight Mixing and a Time Factor. *Eastern-European Journal of Enterprise Technologies*. Vol. 2, No. 9 (140). P. 19–28. https://doi.org/10.15587/1729-4061.2026.355853.12. Gaži P., Kiayias A., Russell A. (2023) Fait Accompli Committee Selection: Improving the Size-Security Tradeoff of Stake-Based Committees. *Proceedings of the 2023 ACM SIGSAC Conference on Computer and Communications Security (CCS '23)*. P. 845–858. https://doi.org/10.1145/3576915.3623194.13. Melnikov G., Müller S., Polyanskii N. (2024) Deterministic Bounds in Committee Selection: Enhancing Decentralization and Scalability in Distributed Ledgers. *arXiv preprint arXiv:2409.10727*. https://doi.org/10.48550/arXiv.2409.10727.

REFERENCES:

1. Saleh F. (2021) Blockchain without Waste: Proof-of-Stake. *The Review of Financial Studies*. Vol. 34, No. 3. P. 1156–1190. https://doi.org/10.1093/rfs/hhaa075. 2. Motepalli S., Jacobsen H.-A. (2024) How Does Stake Distribution Influence Consensus? Analyzing Blockchain Decentralization. *2024 IEEE International Conference on Blockchain and Cryptocurrency (ICBC)*. P. 343–352. https://doi.org/10.1109/ICBC59979.2024.10634400. 3. Motepalli S., Jacobsen H.-A. (2025) Decentralization in PoS Blockchain Consensus: Quantification and Advancement. *IEEE Transactions on Network and Service Management*. Vol. 22, No. 4. P. 2930–2943. https://doi.org/10.1109/TNSM.2025.3561098. 4. Roşu I., Saleh F. (2021) Evolution of Shares in a Proof-of-Stake Cryptocurrency. *Management Science*. Vol. 67, No. 2. P. 661–672.

ISSN 2306-5478

https://doi.org/10.1287/mnsc.2020.3791. 5. Cevallos A., Stewart A. (2021) A Verifiably Secure and Proportional Committee Election Rule. *Proceedings of the 3rd ACM Conference on Advances in Financial Technologies (AFT '21)*. P. 29–42. https://doi.org/10.1145/3479722.3480988. 6. Saad M., Qin Z., Njilla L., Kamhoua C. A., Nyang D., Mohaisen D., Song D., Taduri S. (2021) e-PoS: Making Proof-of-Stake Decentralized and Fair. *IEEE Transactions on Parallel and Distributed Systems*. Vol. 32, No. 8. P. 1961–1973. https://doi.org/10.1109/TPDS.2020.3048853. 7. Leporati A., Rovida L. (2024) Looking for Stability in Proof-of-Stake Based Consensus Mechanisms. *Blockchain: Research and Applications*. Vol. 5, No. 4. Art. 100222. https://doi.org/10.1016/j.bcra.2024.100222. 8. Windiatmaja J. H., Hanggoro D., Salman M., Sari R. F. (2023) PoIR: A Node Selection Mechanism in Reputation-Based Blockchain Consensus Using Bidirectional LSTM Regression Model. *Computers, Materials & Continua*. Vol. 77, No. 2. P. 2309–2339. https://doi.org/10.32604/CMC.2023.041152. 9. Zhang M., Liu M., Ding X., Wang Y., Li G. (2025) GPE-PoS: A Fair and Sybil-Resistant Proof of Stake Consensus. *Journal of Internet Technology*. P. 463–470. https://doi.org/10.70003/160792642025072604005. 10. Solomka I. R., Liubinskyi B. B., Peniak B. O. (2025) Mixed-Weight Committee Selection in Proof-of-Stake: Tunable Stake-Baseline Mixing with Exponential Tail Guarantees and Incentive Compatibility. *Mathematical Modeling and Computing*. Vol. 12, No. 4. P. 1320–1332. https://doi.org/10.23939/mmc2025.04.1320. 11. Solomka I., Liubinskyi B. (2026) Devising a Method for Dynamic Stabilization of PoS Consensus Using Adaptive Weight Mixing and a Time Factor. *Eastern-European Journal of Enterprise Technologies*. Vol. 2, No. 9 (140). P. 19–28. https://doi.org/10.15587/1729-4061.2026.355853.12. Gaži P., Kiayias A., Russell A. (2023) Fait Accompli Committee Selection: Improving the Size-Security Tradeoff of Stake-Based Committees. *Proceedings of the 2023 ACM SIGSAC Conference on Computer and Communications Security (CCS '23)*. P. 845–858. https://doi.org/10.1145/3576915.3623194.13. Melnikov G., Müller S., Polyanskii N. (2024) Deterministic Bounds in Committee Selection: Enhancing Decentralization and Scalability in Distributed Ledgers. *arXiv preprint arXiv:2409.10727*. https://doi.org/10.48550/arXiv.2409.10727.

Соломка І. Р.[1] ([1:ORCID: 0009-0004-3705-8801])
аспірант
Любінський Б. Б.[1] ([1:ORCID: 0000-0002-0715-8068])
к.ф.-м.н., доцент

[1]Кафедра прикладної математики, Інститут прикладної математики та фундаментальних наук, Національний університет «Львівська політехніка»

МЕТОД ОЦІНЮВАННЯ РИЗИКУ ЗАХОПЛЕННЯ КОМІТЕТУ ПРИ КОНЦЕНТРАЦІЇ СТЕЙКУ В МЕРЕЖАХ PROOF-OF-STAKE

У статті досліджено відбір комітету в мережах Proof-of-Stake **комітетного типу при зростанні концентрації стейку. Розглянуто проблему оцінювання ризику захоплення комітету та ймовірності того, що відстежувана коаліція отримає частку місць, достатню для контролю над семпльованим комітетом. На практиці протокол не має простої аудитованої величини, яку можна обчислити зі стану ланцюга і за якою можна зрозуміти, чи залишаються поточні параметри комітету безпечними. Через це розмір комітету, поріг захоплення та допустиме коригувальне змішування часто вибирають евристично, а момент, коли заданий допустимий рівень ризику вже не забезпечується, можна пропустити. Для розв'язання цієї проблеми запропоновано метод оцінювання ризику захоплення комітету за концентрації стейку. Метод поєднує точний сертифікат ризику для відстежуваної коаліції з межею здійсненності/насичення, яка показує, до якого рівня концентрації ризик ще можна втримати в межах заданого бюджету, а після якого це вже неможливо за обмеженого втручання. Метод реалізовано у відтворюваному** Proof-of-Concept **для мережі** Proof-of-Stake. **Отримані результати показують, що запропонований сертифікат дає інформативну оцінку ризику захоплення комітету в робочому діапазоні параметрів. Також показано, що обмежене втручання залишається можливим лише до певного порога концентрації, після якого нездійсненність виявляється явно. У** Proof-of-Concept **сценарії, де концентрація поступово зміщується на користь відстежуваної** top-k **коаліції, метод піднімає червоний прапорець у момент, коли бюджет ризику порушується або стає недосяжним, а фактичні результати жеребкування комітету**

ISSN 2306-5478

демонструють той самий тренд погіршення. Практичне значення роботи полягає в тому, що запропонований метод може використовуватися як інструмент моніторингу, підтримки врядування та аудиту параметрів у Proof-of-Stake **системах.**
Ключові слова: Proof-of-Stake; відбір комітету; захоплення комітету; концентрація стейку; оцінка ризику захоплення комітету; допустимий рівень ризику; межа здійсненності; можливість аудиту зі стану ланцюга; математичне моделювання.

Отримано: 31 березня 2026 року
Прорецензовано: 29 квітня 2026 року
Прийнято до друку: 29 травня 2026 року