Strategic Validator Liquidation in Proof-of-Stake Tokenomics

Lingxiao Zhao[1][0009−0005−5190−5365] ⋆, Madalina Sas[1][0009−0005−0323−6101], and Carmine Ventre[2][0000−0003−1464−1215]⋆⋆

[1] Input Output, United Kingdom
lingxiao.zhao@iohk.io, madalina.sas@iohk.io
[2] King's College London, London, United Kingdom
carmine.ventre@kcl.ac.uk

Abstract. Proof-of-Stake token economies contain a feedback loop between validator liquidation, token price, user demand, participation, and security. We study this loop using a validator-user simulation and empirical game-theoretic analysis. Continuous liquidation policies are reduced to three interval-based meta-strategies; simulation payoffs form a heuristic payoff table; and, Alpha-Rank identifies equilibrium liquidation intervals. Across an 81-regime market-protocol sweep, market environment is the main organising axis: bear and sideways regimes select the High liquidation interval, while bull regimes select the Medium interval. A supplementary drift sweep localises the High-to-Medium transition. The results show that within the simulated grid, equilibrium adaptation is primarily scenario-driven, while bonded security remains close to saturation. However, high bonded security does not eliminate fragility, which can still arise through price-demand-participation feedback.

Keywords: Proof of Stake · Tokenomics · Validator Incentives · Empirical Game-Theoretic Analysis · Agent-Based Modelling

1 Introduction

In recent years, the growth of Proof-of-Stake (PoS) token economies has led to increasingly complex interactions between protocol incentives, validator behaviour, token-market dynamics, and user activity. Validator rewards are designed to sustain participation and security, but once distributed, they remain within the token economy. Validators may retain or liquidate these rewards [16], and their collective liquidation behaviour can generate sell pressure, affect token prices, change the value of future rewards, and reshape incentives for continued participation. When user demand depends on perceived protocol security, these effects can feed back into protocol usage itself [1, 21].

⋆ Corresponding author.
⋆⋆ Carmine Ventre was partially supported by the Engineering and Physical Sciences Research Council (EPSRC) (Grant numbers EP/W034042/1 and UKRI4151).

Existing models of token economies and staking systems capture selected components of this feedback structure, including staking incentives, validator allocation decisions, reward compounding, and wealth concentration [3, 4]. However, many approaches focus on individual components or rely on analytical simplifications, representative agents, continuous-time approximations, or exogenous price processes. Although these assumptions provide tractability, they can obscure the discrete, stochastic, and heterogeneous dynamics of blockchain systems, where epoch-level participation, reward distribution, liquidity pressure, and user activity evolve jointly. In particular, treating validator reward liquidation as an exogenous parameter overlooks the possibility that liquidation behaviour itself changes across market and protocol regimes.

This paper develops a stochastic closed-loop validator-user model for a PoS token economy. The model links user-value growth, perceived security, fee demand, validator participation, bonding, committee selection, reward liquidation, token price formation, and belief-driven participation updates. The underlying tokenomics has the main ingredients shared by many PoS systems. Validator liquidation is modelled as a strategic decision rather than a fixed stress input. Because analytical evaluation is intractable, we use empirical game-theoretic analysis (EGTA) [18] to construct a finite meta-game over liquidation policies and Alpha-Rank [13] to identify stable population compositions.

Our analysis makes three contributions. First, we develop a closed-loop validator-user simulation in which liquidation, price, participation, bonded security, and user demand co-evolve. Second, we formulate validator liquidation as an empirical game over interval-based meta-strategies. Rather than assuming a small number of fixed liquidation rates, validators draw concrete liquidation intensities from Low, Medium, and High intervals, and we use EGTA with Alpha-Rank to identify the liquidation intervals that are stable under the induced population game. Third, we show that market conditions organise equilibrium adaptation more strongly than the swept protocol parameters: bear and sideways regimes select High liquidation, while bull regimes select Medium liquidation. The resulting fragility is not a bonded-security collapse, but a price-demand-participation feedback that can persist even when bonded security is near saturation.

2 Related Work

Research on PoS economics has examined how protocol incentives, token valuation, participation, and security interact. Kogan et al. [10] develop a valuation framework linking staking requirements to token holdings, validator cash flows, transaction activity, and network security. Jermann [7] proposes a dynamic stochastic general equilibrium model of PoS Ethereum, where token value depends on discounted future fee income per unit staked, while user demand drives fee revenue, staking returns, and capital allocation to consensus participation.

A related literature studies blockchain tokenomics as a dynamic control problem. Akcin et al. [1] formulate infrastructure-centric blockchain tokenomics as a feedback-control system, representing circulating supply, token price, and network infrastructure as a continuous dynamical process. Zhang et al. [21, 22] similarly model blockchain economies as stochastic dynamical systems in which token prices, participant behaviour, and protocol reward policies jointly determine system trajectories. These approaches motivate closed-loop tokenomic modelling, although their aggregate representations do not directly capture strategic reward liquidation by heterogeneous validators.

Other work focuses on participation, reward design, and decentralisation. Kiayias et al. [9] study stake-pool payment schemes, characterising trade-offs between participation, decentralisation, and protocol expenditure in a Bayesian game. Ferraioli et al. [6] examine algorithmic monetary policies in repeated participation games, showing how reward allocation affects performance and long-run decentralisation. These studies model strategic participation responses to protocol incentives, but not validator reward liquidation or its interaction with token markets and user demand.

Meanwhile, agent-based modelling provides an alternative approach for studying heterogeneous blockchain behaviour. Rosa et al. [15] develop an agent-based blockchain simulation framework designed to represent large populations and consensus-level interactions. Karra et al. [8] propose an agent-based framework for utility-based cryptoeconomies in which autonomous decisions affect system metrics across successive time steps, producing a closed feedback loop. Li and Wang [12] combine agent-based simulation, a co-evolutionary framework and EGTA to study role selection under Proposer–Builder Separation. Their work demonstrates the value of empirical games for analysing adaptive blockchain participants, but focuses on builder–searcher role selection and bidding rather than validator reward liquidation.

Our endogenous price mechanism is also related to research on automated market making. Angeris and Chitra [2] analyse constant-function market makers and their price and reserve properties, while Park [14] examines structural limitations of decentralised automated market making, including its dependence on arbitrage and external price discovery. Unlike these studies, we do not model the complete reserve mechanics of a constant-function market maker. Instead, we use a reduced-form liquidity layer in which user fee demand and validator liquidation jointly determine price impact.

Taken together, existing work analyses PoS valuation, reward design, participation, and token-market mechanisms, but gives less attention to validator reward liquidation as an endogenous strategic behaviour. We address this gap using a closed-loop validator-user simulation and an empirical meta-game over liquidation policies.

3 Closed-Loop Validator-User Model

The model represents a proof-of-stake token economy as a closed-loop dynamic system. Each epoch links perceived security, user demand, reserve and fee rewards, validator participation, bonding, committee selection, liquidation, token price, token inventories, and participation beliefs. This structure captures feedback channels that are lost when validator behaviour or token price is treated as exogenous.

3.1 Epoch Timing and State

Time is divided into epochs indexed by t. In the main implementation, each epoch contains a single block; participation, bonding, and committee selection are evaluated once per epoch. Epoch t begins with the inherited reference price P_t^{day}, validator token inventories $TK_{v,t-1}$, reserve balance $B_{\text{reserve},t-1}$, participation beliefs $\hat{u}_{v,t-1}$, and bonded stake carried from the previous epoch. It ends with a settlement price P_t^{night}, inventories $TK_{v,t}$, reserve balance $B_{\text{reserve},t}$, and beliefs $\hat{u}_{v,t}$.

To make within-epoch timing explicit, the model distinguishes a day reference price from a night settlement price, mimicking the well-known Lagos-Wright framework [11]. These labels represent two sequential stages of the same token market rather than two separate markets. At the start of epoch t, validators and users observe the inherited reference price P_t^{day}. User fee demand and validator liquidation quantities are then formed using this common reference price. Their net order flow determines the end-of-epoch settlement price P_t^{night}, at which validator sales are realised. This settlement price becomes the reference price for the next epoch: $P_{t+1}^{day} = P_t^{night}$.

We write P_0 for the initial token price and set $P_0 = P_0^{day} = P_0^{night}$ at initialisation. Within each epoch, security and demand are evaluated first, followed by reward-pool formation, participation, bonding, committee selection, reward allocation, liquidation, price settlement, inventory update, and belief update. This ordering separates formation of the aggregate reward pool from its allocation to validators, while liquidation and market settlement occur once per epoch.

3.2 User Demand and Security Feedback

Let $b_{v,t}$ denote the tokens bonded by validator v in epoch t, as defined in the participation and bonding rule below. The perceived security at the start of epoch t depends on bonded stake carried from epoch $t-1$:

$$B_{\text{bond},t-1} = \sum_v b_{v,t-1}, \qquad \gamma_t = 1 - \exp\left(-K_{\text{sec}} \frac{B_{\text{bond},t-1}}{N_{\text{supply}}}\right), \tag{1}$$

where K_{sec} controls how quickly perceived security saturates as bonded stake rises, and N_{supply} is the fixed total token supply used to normalise bonded stake. Meanwhile, aggregate user value follows a log-growth process

$$\log V_t = \log V_{t-1} + g_V + \sigma_V \epsilon_t, \qquad \epsilon_t \sim \mathcal{N}(0, 1). \tag{2}$$

Here V_t represents latent aggregate value for using the protocol service, rather than direct token-purchase demand. Perceived security converts this latent value into realised protocol usage:

$$D_t = V_t \cdot \gamma_t, \qquad F_t^{USD} = \eta \cdot D_t. \tag{3}$$

Here D_t is security-adjusted protocol usage value, η is the protocol fee rate, and F_t^{USD} is the USD-denominated fee demand generated by user activity. Since fees are settled through the token economy, F_t^{USD} is the user-side buy pressure that enters the reward-pool and price-formation equations. This security-gated channel allows participation and bonding in one epoch to affect protocol usage and token demand in the next.

3.3 Reward-Pool Formation

Validator rewards combine reserve emission and fee-derived tokens. Let θ_B and θ_T denote the fractions of the fixed total token supply N_{supply} initially allocated to the reserve and treasury, respectively. The initial reserve balance is therefore $B_{\text{reserve},0} = \theta_B \cdot N_{\text{supply}}$. Both θ_B and θ_T are fixed allocation parameters and are distinct from the time-varying reserve balance $B_{\text{reserve},t}$. The treasury is treated as inactive over the simulation horizon: it neither funds validator rewards nor participates in token-market transactions.

Given the *annual* inflation parameter π (loosely the percentage of tokens invested as rewards), the corresponding annual reserve-emission rate is

$$R_{\text{annual}} = \pi \frac{1 - \theta_B - \theta_T}{\theta_B}. \tag{4}$$

Let E_{year} denote the number of simulation epochs in one year. We set each epoch to 12 hours as a modelling convention, providing sufficient temporal resolution for validator participation, reward allocation, liquidation, and price adjustment without incurring the computational cost of substantially shorter intervals. Accordingly, the model uses $E_{\text{year}} = 2 \times 365 = 730$. The corresponding per-epoch reserve-emission rate and base reserve reward are

$$R_{\text{epoch}} = \frac{R_{\text{annual}}}{E_{\text{year}}}, \qquad R_{\text{base},t} = B_{\text{reserve},t-1} \cdot R_{\text{epoch}}. \tag{5}$$

Thus, π specifies annual token issuance relative to the initial token allocation outside the reserve and treasury, while R_{annual} expresses the corresponding annual emission as a fraction of the reserve balance.

The reference-price accounting quantity associated with fee demand is $Q_{\text{fee},t} = F_t^{USD}/P_t^{day}$. The parameter $\phi \in [0, 1]$ is the fraction of these fee-derived tokens retained by the reserve, while the remaining fraction enters the validator reward pool:

$$R_t = R_{\text{base},t} + (1 - \phi) Q_{\text{fee},t}. \tag{6}$$

The reserve balance therefore evolves as

$$B_{\text{reserve},t} = B_{\text{reserve},t-1} - R_{\text{base},t} + \phi Q_{\text{fee},t}. \tag{7}$$

Thus, a larger ϕ replenishes the reserve more strongly and passes a smaller fraction of fee-derived tokens to validators.

3.4 Participation and Bonding

At the start of epoch t, validator v observes the smoothed utility belief inherited from the previous epoch, $\hat{u}_{v,t-1}$. Participation follows a logistic response to this belief:

$$\Pr(S_{v,t}=1) = \mathbf{1}\{TK_{v,t-1} > TK_{\min}\}\sigma\left(\frac{\hat{u}_{v,t-1}}{\tau}\right),\ \sigma(z) = \frac{1}{1+\exp(-z)}. \quad (8)$$

Here $S_{v,t} \in \{0,1\}$ is the epoch-level participation indicator, $TK_{\min}$ is the minimum inventory required to participate, $N_{\text{part},t} = \sum_v S_{v,t}$ is the number of participating validators, and τ controls the stochasticity of the response. Conditional on participation, validator v bonds a noisy fraction of its beginning-of-epoch inventory. Define

$$x_t = \frac{P_t^{day}}{P_0}, \qquad \bar{x}_t = \frac{x_t}{1+x_t}, \qquad \widetilde{\rho}_{v,t} = (k_b\bar{x}_t + (1-k_b)\alpha)\,\varepsilon_{v,t}, \quad (9)$$

where $\varepsilon_{v,t} \sim \mathcal{N}(1, \sigma^2_{\text{bond}})$ is multiplicative bonding noise. The effective bonding fraction is bounded between zero and one:

$$\rho_{v,t} = \min\Big\{\max\{\widetilde{\rho}_{v,t}, 0\}, 1\Big\}. \quad (10)$$

The bonded amount is $b_{v,t} = S_{v,t} \cdot \rho_{v,t} \cdot TK_{v,t-1}$. Thus, only participating validators bond tokens, and bonded stake is non-negative and cannot exceed the validator's inventory. Bonded tokens remain owned by the validator and determine security and reward eligibility rather than being deducted from inventory.

3.5 Committee Selection

Committee selection is restricted to validators that participate and bond a positive number of tokens. Define the sampling weight

$$w_{v,t} = TK_{v,t-1} \cdot \mathbf{1}\{b_{v,t} > 0\}. \quad (11)$$

Let $N_{\text{active},t} = \sum_v \mathbf{1}\{b_{v,t} > 0\}$ denote the number of eligible validators in epoch t. If $N_{\text{active},t} = 0$, the committee is empty. Otherwise, the committee size is drawn as $Q_t \sim \text{UniformInt}(1, N_{\text{active},t})$. Here, conditional on Q_t, the committee $\mathcal{C}_t$ is sampled without replacement from eligible validators using $w_{v,t}$ as the sampling weights.

The committee-membership indicator is $C_{v,t} = \mathbf{1}\{v \in \mathcal{C}_t\}$. Thus, positive bonded stake determines committee eligibility, while a larger token inventory increases a validator's chance of selection. The sampled committee contains exactly Q_t validators.

3.6 Reward Distribution

Validator v receives reward

$$r_{v,t} = R_t \left[(1-\alpha) \frac{S_{v,t}}{N_{\text{part},t}} + \alpha \frac{b_{v,t} \cdot C_{v,t}}{\sum_u b_{u,t} \cdot C_{u,t}} \right]. \tag{12}$$

The first component distributes a fraction $1-\alpha$ of the reward pool equally among participating validators. The second distributes the remaining fraction α among committee members in proportion to their bonded stake. Hence, $\alpha \in [0,1]$ controls the trade-off between broad participation rewards and stake-weighted committee rewards. The same α also appears in the reduced-form bonding rule, capturing the assumption that stronger committee-weighted reward emphasis is associated with stronger bonding incentives.

3.7 Liquidation, Price Formation, and Token Inventories

Validator liquidation converts allocated rewards into market sell pressure. After rewards are distributed, each validator first sells reward tokens up to the amount required to cover operating costs, and then sells a strategy-dependent fraction of any remaining reward surplus:

$$\text{sell}_{\text{ops},v,t} = \min\left(\frac{c_{\text{ops}}}{P_t^{day}}, r_{v,t} \right), \tag{13}$$

$$\text{sell}_{\text{take},v,t} = \beta_{\text{take},v} \max\left(r_{v,t} - \frac{c_{\text{ops}}}{P_t^{day}}, 0 \right). \tag{14}$$

The liquidation intensity $\beta_{\text{take},v} \in [0,1]$ is the strategic variable. A value of zero means that validator v sells only enough reward tokens to cover operating costs, while a value of one means that all rewards remaining after operating-cost sales are liquidated. If validator v is inactive in epoch t, then $r_{v,t} = 0$, so both operating-cost and discretionary sales are zero without an additional participation multiplier. In this way, total sales are

$$\text{sell}_{v,t} = \text{sell}_{\text{ops},v,t} + \text{sell}_{\text{take},v,t}. \tag{15}$$

User-side fee settlement generates buy pressure in a reduced-form token market, while validator liquidation generates sell pressure into the same market. These flows are not modelled as bilateral trades between users and validators. Instead, they interact through an aggregate liquidity layer, so any net order imbalance is absorbed by price impact. Let $L_t = \sum_v \text{sell}_{v,t}$ denote aggregate validator sales. Reference-price excess demand is

$$Z_t = F_t^{USD} - P_t^{day} L_t. \tag{16}$$

The night settlement price is determined by the reduced-form price-impact rule [5]

$$P_t^{night} = P_t^{day} \cdot \exp\left(\lambda_{\text{mkt}} \frac{Z_t}{P_t^{day}} \right). \tag{17}$$

Here λ_{mkt} is scaled by market liquidity so that the exponent is dimensionless. Positive excess demand raises the settlement price, while negative excess demand lowers it. This price formation is a tractable approximation that preserves endogenous feedback from aggregate order imbalance without modelling the full market microstructure.

For reporting, realised user token acquisition is $B_t^{\mathrm{buy,realised}} = F_t^{\mathrm{USD}}/P_t^{\mathrm{night}}$. This is the token quantity implied by user-side fee settlement at the settlement price, and is distinct from the reference-price accounting quantity $Q_{\mathrm{fee},t}$ used in reward-pool formation. Then, $TK_{v,t}$ denotes the total token inventory owned by validator v. Let N_{val} denote the total number of validators. Validators receive equal shares of the initial validator allocation:

$$TK_{v,0} = \frac{A_{\mathrm{drop}} \cdot N_{\mathrm{supply}}}{N_{\mathrm{val}}}, \qquad v \in \{1, \ldots, N_{\mathrm{val}}\}, \tag{18}$$

where $A_{\mathrm{drop}} \in [0, 1]$ is the fraction of the initial token supply allocated equally among validators and is already included in, rather than additional to, N_{supply}. The validator inventory is updated after liquidation:

$$TK_{v,t} = \max\left\{TK_{v,t-1} + r_{v,t} - \mathrm{sell}_{v,t}, 0\right\}. \tag{19}$$

Bonded tokens remain within this inventory, while reward tokens that are sold are removed at settlement.

3.8 Utility Signal and Belief Update

After epoch t is realised, validator v observes the behavioural payoff proxy

$$\widetilde{u}_{v,t} = \mathrm{sell}_{v,t} \cdot P_t^{night} - c_{\mathrm{ops}} \cdot S_{v,t}. \tag{20}$$

This signal measures realised liquidation revenue net of the epoch-level operating cost paid by participating validators. It governs subsequent participation decisions but is not the payoff used to evaluate strategies in the empirical game. Beliefs are then updated through exponential smoothing:

$$\hat{u}_{v,t} = (1 - \delta)\hat{u}_{v,t-1} + \delta\widetilde{u}_{v,t}, \tag{21}$$

Here $\delta \in [0, 1]$ determines how strongly the latest realised signal affects the validator's belief. A larger δ places more weight on the current epoch, whereas a smaller δ produces more persistent beliefs. The resulting belief $\hat{u}_{v,t}$ enters the participation probability in epoch $t + 1$. This behavioural signal is distinct from the full-horizon wealth payoff used for EGTA evaluation in Section 4.

4 Empirical Game-Theoretic Analysis

The simulator induces a strategic game among validators because the payoff of a liquidation policy depends on the liquidation behaviour of the rest of the population. The underlying action space is continuous, making direct evaluation of all joint policies infeasible. We therefore construct a meta-game whose finite meta-strategies represent regions of the continuous liquidation space, and estimate its payoffs using empirical game-theoretic analysis (EGTA).

4.1 Continuous Validator Game

Let $\mathcal{V} = \{1, \ldots, N\}$ denote the strategic validators embedded in the full simulated validator population. Each strategic validator v chooses a liquidation intensity

$$\beta_v \equiv \beta_{\text{take},v} \in \mathcal{B} = [0,1], \tag{22}$$

which is held fixed over the simulated horizon. The resulting continuous simulation-induced game is

$$\mathcal{G} = \left(\mathcal{V}, (\mathcal{B}_v)_{v\in\mathcal{V}}, (U_v)_{v\in\mathcal{V}}\right), \qquad \mathcal{B}_v = [0,1], \tag{23}$$

where each payoff function U_v is induced by the closed-loop simulator. In the main experiments, $N = 100$ validators are strategic; the remaining validators in the 3,000-validator population use the fixed background intensity $\beta_{\text{take}} = 0.5$. A continuous strategy profile is

$$\boldsymbol{\beta} = (\beta_v)_{v\in\mathcal{V}} \in [0,1]^N. \tag{24}$$

For each market and protocol regime, the initial state Ω_0 is fixed. Simulation replications vary the stochastic simulator draws ξ and, once meta-strategies are introduced below, the random assignment and within-interval liquidation realisations used to estimate empirical payoffs. The simulator then maps a strategy profile into a trajectory

$$\mathcal{T}(\boldsymbol{\beta}; \Omega_0, \xi) = \left\{P_t^{day}, P_t^{night}, TK_{v,t}, S_{v,t}, r_{v,t}, \text{sell}_{v,t}, V_t, \gamma_t, b_{v,t}, C_{v,t}\right\}_{t=1}^{T}. \tag{25}$$

Recall that $S_{v,t}$ is epoch-level participation, $b_{v,t}$ is bonded stake, and $C_{v,t}$ is committee membership under the single-block-per-epoch model. The profile enters the trajectory through validator liquidation:

$$\text{sell}_{v,t} = \min\left(\frac{c_{\text{ops}}}{P_t^{day}}, r_{v,t}\right) + \beta_v \max\left(r_{v,t} - \frac{c_{\text{ops}}}{P_t^{day}}, 0\right). \tag{26}$$

There is no separate multiplication by the participation indicator in the sale equation: if $S_{v,t} = 0$, then $r_{v,t} = 0$ and therefore $\text{sell}_{v,t} = 0$. Thus, the strategy profile affects aggregate selling pressure, price, rewards, participation, bonding, and ultimately validator payoff. Then, the realised full-horizon payoff of validator v under profile $\boldsymbol{\beta}$ is

$$u_v(\boldsymbol{\beta}; \Omega_0, \xi) = \sum_{t=1}^{T}\left(\text{sell}_{v,t} P_t^{night} - c_{\text{ops}} S_{v,t}\right) + TK_{v,T} P_T^{night} - TK_{v,0} P_0. \tag{27}$$

The first term measures realised cashflow net of operating costs, while the terminal term accounts for the change in token inventory value over the simulation horizon. For a fixed initial state, the expected continuous-game payoff is

$$U_v(\boldsymbol{\beta} \mid \Omega_0) = \mathbb{E}_\xi\left[u_v(\boldsymbol{\beta}; \Omega_0, \xi)\right]. \tag{28}$$

Because $\boldsymbol{\beta}$ affects the endogenous price path and participation process, U_v is a profile-dependent payoff, not a payoff of a liquidation strategy in isolation.

4.2 Meta-Game Construction

The continuous profile space $[0,1]^N$ cannot be exhaustively evaluated. We therefore follow the meta-game approach in which a finite meta-strategy summarises a region of a richer underlying strategy space [20]. Let

$$\mathcal{M} = \{m_L, m_M, m_H\}, \tag{29}$$

denote the Low, Medium, and High liquidation meta-strategies. Their associated regions are

$$\mathcal{I}_L = [0, 1/3], \qquad \mathcal{I}_M = (1/3, 2/3], \qquad \mathcal{I}_H = (2/3, 1]. \tag{30}$$

A meta-strategy is therefore not identified with a single liquidation intensity. Instead, it induces a distribution over concrete liquidation policies:

$$\beta_v \mid m_v = m_\ell \sim \mathrm{Unif}(\mathcal{I}_\ell), \qquad \ell \in \{L, M, H\}. \tag{31}$$

Endpoint conventions make the intervals disjoint; because the realisations are continuous, they do not affect the simulated expectations.

For each simulation replication, every strategic validator is first assigned a meta-strategy and then independently draws one β_v from the corresponding interval. This realised intensity remains fixed for all T epochs in that replication. Consequently, variation within an interval is integrated into the meta-strategy payoff rather than suppressed by selecting an arbitrary representative point. This broadens the behavioural coverage relative to the previous fixed-action specification while retaining a tractable three-strategy empirical game.

For a meta-strategy profile $\boldsymbol{m} = (m_v)_{v\in\mathcal{V}} \in \mathcal{M}^N$, define the realisation kernel

$$Q(d\boldsymbol{\beta} \mid \boldsymbol{m}) = \prod_{v\in\mathcal{V}} \mathrm{Unif}\left(d\beta_v; \mathcal{I}_{m_v}\right). \tag{32}$$

The induced meta-game is

$$\widetilde{\mathcal{G}} = \Big(\mathcal{V}, (\mathcal{M}_v)_{v\in\mathcal{V}}, (\widetilde{U}_v)_{v\in\mathcal{V}}\Big), \qquad \mathcal{M}_v = \mathcal{M}, \tag{33}$$

with payoff

$$\widetilde{U}_v(\boldsymbol{m} \mid \Omega_0) = \mathbb{E}_{\boldsymbol{\beta}\sim Q(\cdot\mid\boldsymbol{m}),\,\xi}\left[u_v(\boldsymbol{\beta}; \Omega_0, \xi)\right]. \tag{34}$$

Thus, the meta-game compares distributions of liquidation policies. The expectation covers both within-interval policy realisations and all stochastic elements of the validator-user simulation.

4.3 Symmetric Heuristic Payoff Table

Explicitly evaluating Eq. (34) for every labelled meta-profile would require 3^N profiles. EGTA reduces this burden using an anonymous strategy-count representation [17, 19]. We construct a symmetrised empirical game by representing each population profile as a count vector and randomly assigning meta-strategy labels to validators for each count vector and replication. A count vector is

$$\boldsymbol{n}^k = (n_L^k, n_M^k, n_H^k), \qquad \sum_{\ell \in \{L,M,H\}} n_\ell^k = N, \tag{35}$$

where n_ℓ^k is the number of validators choosing meta-strategy m_ℓ. The number of reduced population profiles is

$$\binom{N + |\mathcal{M}| - 1}{|\mathcal{M}| - 1} = \binom{N+2}{2}, \tag{36}$$

which is much smaller than 3^N. With $N = 100$, this gives 5,151 population profiles.

For count vector $\boldsymbol{n}^k$, let $\boldsymbol{m}^{k,r}$ denote the random assignment of meta-strategy labels to validators in replication r, and let $\boldsymbol{\beta}^{k,r} \sim Q(\cdot \mid \boldsymbol{m}^{k,r})$. Randomising labels prevents a validator's array position from being systematically associated with a meta-strategy. Let

$$\mathcal{V}_\ell^{k,r} = \{v \in \mathcal{V} : m_v^{k,r} = m_\ell\}, \qquad |\mathcal{V}_\ell^{k,r}| = n_\ell^k. \tag{37}$$

The simulated average payoff for meta-strategy m_ℓ in replication r is

$$p_\ell^{(r)}(\boldsymbol{n}^k) = \frac{1}{n_\ell^k} \sum_{v \in \mathcal{V}_\ell^{k,r}} u_v(\boldsymbol{\beta}^{k,r}; \Omega_0, \xi^r), \qquad n_\ell^k > 0. \tag{38}$$

With R independent replications, the empirical payoff entry is

$$\hat{U}_{\ell,k} = \frac{1}{R} \sum_{r=1}^{R} p_\ell^{(r)}(\boldsymbol{n}^k), \qquad n_\ell^k > 0. \tag{39}$$

In the main experiments, $R = 100$ replications are used for each meta-strategy-count configuration.

The heuristic payoff table is $\mathcal{H} = (\mathcal{N}, \hat{U})$, where $\mathcal{N}$ contains all feasible count vectors and $\hat{U}$ contains the corresponding empirical meta-strategy payoffs. If $n_\ell^k = 0$, no validator is observed using meta-strategy m_ℓ under count profile $\boldsymbol{n}^k$; such entries are placeholders and are not interpreted as realised payoffs.

4.4 Alpha-Rank Selection

We apply the Alpha-Rank algorithm [13] to the symmetric HPT. Alpha-Rank constructs an evolutionary Markov chain over feasible strategy-count profiles and assigns each profile $\boldsymbol{n} \in \mathcal{N}$ a stationary mass $\mu(\boldsymbol{n})$. Profiles with larger mass are interpreted as more evolutionarily stable compositions of the Low, Medium, and High liquidation meta-strategies. We report the highest-mass profile as the top Alpha-Rank state. For each meta-strategy m_ℓ, its marginal Alpha-Rank mass is

$$\mu_\ell^{marg} = \sum_{\boldsymbol{n} \in \mathcal{N}} \mu(\boldsymbol{n}) \frac{n_\ell}{N}. \tag{40}$$

For aggregate system outcomes, we use the smallest ranked set $\mathcal{N}_{80}$ whose cumulative stationary mass reaches 80%, as a representative space of highly likely outcomes. For each count profile $\boldsymbol{n}$, let $x_{\boldsymbol{n},t}$ denote the empirical mean trajectory of outcome x across simulation replications. The Alpha-Rank-weighted trajectory is

$$\bar{x}_t^{80} = \sum_{\boldsymbol{n} \in \mathcal{N}_{80}} \frac{\mu(\boldsymbol{n})}{\sum_{\boldsymbol{n}' \in \mathcal{N}_{80}} \mu(\boldsymbol{n}')} x_{\boldsymbol{n},t}. \tag{41}$$

We use this trajectory to report participation, price, demand, security, and realised user-token outcomes.

4.5 Experimental Design

We evaluate equilibrium liquidation behaviour across 81 regimes formed by three market scenarios and a $3 \times 3 \times 3$ protocol grid. The market scenarios vary the drift g_V in the aggregate user-value process in Eq. (2), representing bear, sideways, and bull conditions. The protocol grid varies the committee-weighted reward share, the reserve-emission parameter, and the reserve fee-retention share:

$$\alpha \in \{0.10, 0.20, 0.30\}, \ \pi \in \{0.03, 0.05, 0.07\}, \ \phi \in \{0.01, 0.05, 0.10\}. \tag{42}$$

For each market–protocol regime, we estimate a symmetric HPT over the three liquidation meta-strategies, compute Alpha-Rank on the induced population game, and use the resulting ranking to identify stable liquidation intervals. We then generate diagnostic system trajectories from the selected equilibrium region, aggregating them using the top-80% weighted rule in Eq. (41).

We report outcomes that track both strategic selection and closed-loop system performance: Alpha-Rank liquidation selection, marginal strategy masses, validator participation, price–demand dynamics, bonded security, realised user token acquisition, and fragility classification. These measures allow us to distinguish three questions: which liquidation interval is stable, whether the selected equilibrium sustains validator activity and bonded security, and whether the surrounding token economy remains healthy under the induced buy–sell pressure. Because liquidation intervals are selected endogenously, cross-scenario comparisons should be interpreted as equilibrium-regime outcomes rather than isolated causal effects of market drift alone.

5 Results

In this section we analyse how the equilibrium liquidation behaviour identified by Alpha-Rank propagates through the closed-loop token economy. The presentation follows the model-implied sequence of feedback: we first identify which liquidation intervals are selected under different market and protocol conditions; we then examine how these equilibrium selections are associated with validator participation, price and demand dynamics, bonded security and fragility.

A central pattern emerges throughout the section: under equilibrium adaptation, market-growth conditions organise the dominant liquidation interval and most aggregate outcomes, whereas protocol parameters mainly affect narrower margins, especially realised token acquisition and residual security shortfall.

5.1 Equilibrium Liquidation Behaviour

We first examine whether equilibrium liquidation behaviour is primarily shaped by protocol parameters or by the broader market environment. Across the 81-regime grid, the top Alpha-Rank state is sharply separated by market scenario: all bear and sideways regimes select the High liquidation interval, $\beta_{\text{take}} \in (2/3, 1]$, whereas all bull regimes select the Medium interval, $\beta_{\text{take}} \in (1/3, 2/3]$. Varying α, π, and ϕ does not induce an additional top-strategy transition within the main grid. The selected interval is strongly dominant in each configuration, but this sharp separation should be interpreted at the resolution of the three-interval meta-game rather than as evidence that the continuous liquidation space has no nearby transition.

5.2 Validator Participation

Having established which liquidation behaviours are selected, we next examine whether equilibrium participation differs across market conditions. Table 1 shows a clear ordering: participation is lowest in bear regimes, intermediate in sideways regimes, and highest in bull regimes. Mean participation is approximately 0.692, 0.770, and 0.858, respectively, while the within-scenario ranges across protocol configurations are narrow. This pattern is consistent with the closed-loop mechanism in the model: stronger user-value growth supports stronger price paths and validator cashflows, which improves incentives for continued participation. Conversely, adverse market conditions weaken participation, although the bear-regime values remain close to 0.7 and therefore do not indicate a collapse of the validator base.

Table 1. Equilibrium participation by market scenario. Values summarise the 27 protocol configurations within each scenario.

Scenario	Mean participation	Range
Bear	0.692	0.688–0.694
Sideways	0.770	0.767–0.774
Bull	0.858	0.849–0.864

Protocol parameters produce only modest within-scenario variation. The largest visible pattern is a small positive association between π and participation, especially in bull regimes, but this effect is much smaller than the separation across market scenarios. We therefore treat participation as primarily scenario-driven under the present calibration, and analyse residual protocol effects in Section 5.5.

5.3 Price, Demand, and Liquidation Dynamics

The participation differences documented above are accompanied by systematic changes in price and demand. Equilibrium liquidation behaviour affects these variables through aggregate validator sell pressure, realised user purchase activity, and the security-mediated demand response. Figure 1 shows that the top-80% weighted equilibrium price path declines in the bear regime, remains close to its initial level in the sideways regime, and appreciates strongly in the bull regime. The median final price is approximately 0.402 in bear regimes, 0.961 in sideways regimes, and 2.279 in bull regimes. Relative to the initial price, these values correspond to median changes of approximately −59.9%, −4.1%, and +127.4%, respectively.

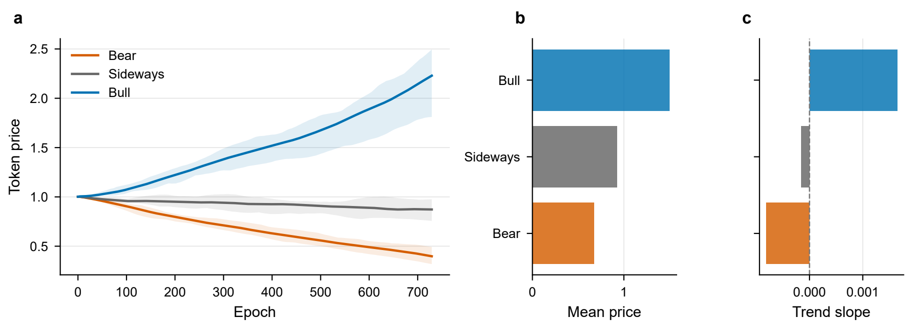

Fig. 1. Equilibrium price dynamics by market scenario under the top-80% Alpha-Rank weighted trajectory rule. Panel a shows the mean price path with 10th–90th percentile bands across trajectory replications. Panels b and c summarise the same trajectories by mean price and linear trend slope, respectively.

To examine outcomes around this transition, Fig. 2 reports the fixed-parameter drift sweep. Higher user-value growth is associated with higher final price, participation, and demand. Nevertheless, realised user token acquisition declines because $B_t^{buy,realised} = F_t^{USD}/P_t^{night}$: the higher settlement price reduces the number of tokens acquired per dollar of fee demand. These patterns coincide with the High-to-Medium transition but should not be interpreted as its isolated causal effect.

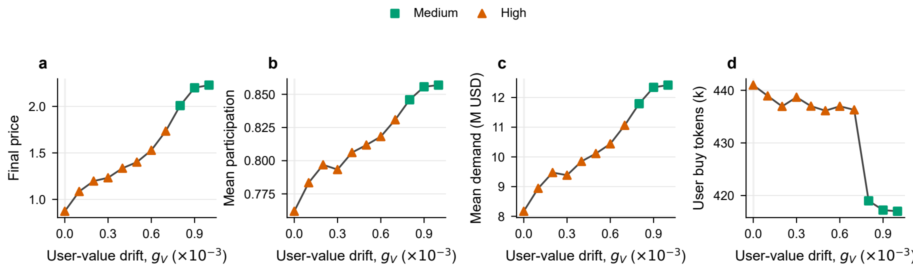

Fig. 2. Local strategy-outcome transition in the fixed-parameter g_V drift sweep at $(\alpha, \pi, \phi) = (0.20, 0.05, 0.05)$. Marker colour denotes the dominant liquidation interval, and the shaded band marks the High-to-Medium transition between $g_V = 0.0007$ and $g_V = 0.0008$.

5.4 Security and Closed-Loop Fragility

The preceding results show large differences in participation, price, and demand. Bonded security, however, remains close to saturation across the main sweep: the median mean-security values are approximately 0.9962 in bear regimes, 0.9995 in sideways regimes, and 0.9999 in bull regimes. Because the raw security measure is compressed near one, the log-scale security shortfall in Fig. 3 is more informative than mean security itself. It reveals small but systematic differences that would otherwise be visually obscured.

High bonded security does not rule out closed-loop fragility. The fragility classification combines price change, final participation, and security shortfall across the 81-regime main sweep. Under the adopted classification thresholds, all bear regimes are classified as fragile, whereas sideways and bull regimes are not. Because bonded security remains high, this classification reflects a price–demand–participation failure mode rather than an immediate collapse in the bonded stake.

5.5 Parameter Importance and Residual Protocol Effects

The preceding results show that market scenario organises the main equilibrium outcomes: bear and sideways regimes select High liquidation, bull regimes select Medium liquidation, and fragility remains confined to the bear scenario. We next use an OLS drop-one R^2 diagnostic to identify which swept dimensions explain variation across outcome measures.

Table 2 shows that in this descriptive diagnostic, the scenario explains most variation in participation, demand change, price change, and final price. Protocol parameters are more important for specific margins: π dominates variation in realised user token acquisition, while α dominates variation in security shortfall. Because scenario and equilibrium selection are coupled in the main sweep, the drop-one R^2 values should be read as explanatory diagnostics over realised equilibrium-regime outcomes, not as causal estimates.

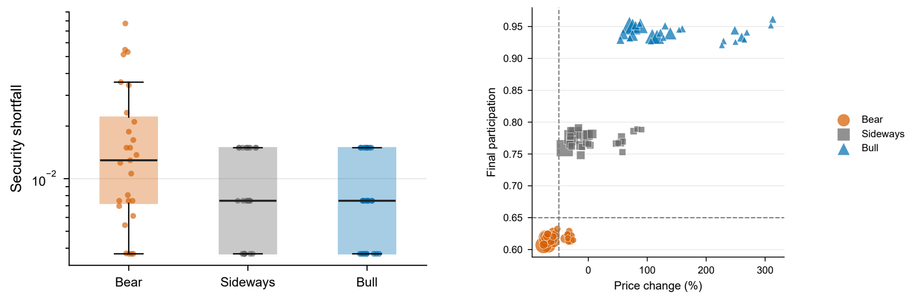

Fig. 3. Security shortfall and closed-loop fragility across the main sweep. The left panel reports security shortfall on a logarithmic scale, while the right panel reports fragility classification. Fragility is concentrated in bear regimes, where adverse price-demand feedback dominates the otherwise high bonded-security level.

Table 2. Dominant factors in the OLS drop-one R^2 diagnostic.

Outcome	Dominant factor	Contribution
Mean participation	Scenario	0.989
Demand change	Scenario	0.989
Price change	Scenario	0.727
Log final price	Scenario	0.805
Log realised user buy tokens	Inflation parameter π	0.954
Log security shortfall	Reward sharing parameter α	0.715

6 Discussion and Conclusion

This paper studies validator liquidation as an endogenous equilibrium response in a closed-loop token economy. Liquidation affects sell pressure, token price, participation, and security-gated user demand, making equilibrium behaviour regime-dependent. Under the current calibration, bear and sideways regimes select High liquidation, while bull regimes select Medium liquidation, with the transition occurring between $g_V = 0.0007$ and $g_V = 0.0008$. Participation, price, and demand follow the same ordering. Protocol parameters mainly affect secondary margins: π shapes user token acquisition, while α is associated with residual security shortfall. High bonded security therefore does not eliminate fragility, which can emerge through price-demand-participation feedback.

The analysis has two main limitations. The continuous liquidation space is approximated by three interval-based meta-strategies, so the results identify equilibrium regions rather than pointwise optimal policies. Moreover, the scenario comparisons do not causally separate market drift from strategy selection. Future work could use finer strategy partitions, broader parameter grids, and robustness checks around the High-to-Medium transition and security-shortfall effects.

References

1. Akcin, O., Streit, R.P., Oommen, B., Vishwanath, S., Chinchali, S.: A control theoretic approach to infrastructure-centric blockchain tokenomics. arXiv preprint arXiv:2210.12881 (2022)
2. Angeris, G., Chitra, T.: Improved price oracles: Constant function market makers. In: Proceedings of the 2nd ACM Conference on Adv. in Fin. Tech. pp. 80–91 (2020)
3. Chitra, T.: Competitive equilibria between staking and on-chain lending (2021)
4. Fanti, G., Kogan, L., Oh, S., Ruan, K., Viswanath, P., Wang, G.: Compounding of wealth in proof-of-stake cryptocurrencies. In: International conference on financial cryptography and data security. pp. 42–61. Springer (2019)
5. Farmer, J.D., Joshi, S.: The price dynamics of common trading strategies. Journal of Economic Behavior & Organization **49**(2), 149–171 (2002)
6. Ferraioli, D., Penna, P., Schneider, M., Ventre, C.: Algorithmic monetary policies for blockchain participation games. arXiv preprint arXiv:2512.16514 (2025)
7. Jermann, U.J.: A macro finance model for proof-of-stake Ethereum. Available at SSRN 4335835 (2023)
8. Karra, K., Mellan, T., Silva, M., Madrigal-Cianci, J.P., Cortes, A.C., Zhang, Z.: An agent-based model framework for utility-based cryptoeconomies. arXiv preprint (2023)
9. Kiayias, A., Koutsoupias, E., Marmolejo-Cossio, F., Stouka, A.P.: Balancing participation and decentralization in proof-of-stake cryptocurrencies. In: International Symposium on Algorithmic Game Theory. pp. 333–350. Springer (2024)
10. Kogan, L., Fanti, G., Viswanath, P.: Economics of proof-of-stake payment systems (2021)
11. Lagos, R., Wright, R.: A unified framework for monetary theory and policy analysis. Journal of political Economy **113**(3), 463–484 (2005)
12. Li, Y., Wang, Z.: Role-selection game in block production under proposer-builder separation. In: 2025 IEEE International Conference on Blockchain and Cryptocurrency (ICBC). pp. 1–9. IEEE (2025)
13. Omidshafiei, S., Papadimitriou, C., Piliouras, G., Tuyls, K., Rowland, M., Lespiau, J.B., Czarnecki, W.M., Lanctot, M., Perolat, J., Munos, R.: α-rank: Multi-agent evaluation by evolution. Scientific reports **9**(1), 9937 (2019)
14. Park, A.: The conceptual flaws of decentralized automated market making. Management Science **69**(11), 6731–6751 (2023)
15. Rosa, E., D'Angelo, G., Ferretti, S.: Agent-based simulation of blockchains. In: Asian Simulation Conference. pp. 115–126. Springer (2019)
16. Tang, W., Yao, D.D.: Trading under the proof-of-stake protocol–a continuous-time control approach. Mathematical Finance **33**(4), 979–1004 (2023)
17. Walsh, W.E., Das, R., Tesauro, G., Kephart, J.O.: Analyzing complex strategic interactions in multi-agent systems. In: AAAI-02 workshop on game-theoretic and decision-theoretic agents. pp. 109–118 (2002)
18. Wellman, M.P.: Methods for empirical game-theoretic analysis. In: Proceedings of the Twenty-First National Conference on Artificial Intelligence. pp. 1552–1556 (2006)
19. Wellman, M.P., Tuyls, K., Greenwald, A.: Empirical game theoretic analysis: A survey. Journal of artificial intelligence research **82**, 1017–1076 (2025)
20. Wu, F., Thiery, T., Leonardos, S., Ventre, C.: From competition to centralization: The oligopoly in ethereum block building auctions. arXiv preprint arXiv:2412.18074 (2024)

21. Zhang, Z.: Engineering token economy with system modeling. arXiv preprint arXiv:1907.00899 (2019)
22. Zhang, Z., Zargham, M., Preciado, V.M.: On modeling blockchain-enabled economic networks as stochastic dynamical systems. Applied Network Science **5**(1), 19 (2020)