Liquid Staking and the Control-Exposure Wedge

Marc Dordal i Carreras Gloria Christina Heesen Kohei Kawaguchi

Hong Kong University of Science and Technology

April 2026

Traditional Proof-of-Stake (PoS)

Traditional Proof-of-Stake (PoS)

- Validators lock Ether (ETH) as stake
- Voting power comes from active stake
- Misbehavior triggers slashing
- Attacks can also depress ETH value

Traditional PoS keeps control and loss-bearing aligned.

Liquid Staking

Liquid staking

- Operators originates and issue liquid staking tokens (LST)
- Operators run validators for the pool
- Operators bond part of their stake as collateral
- Users deposit ETH and receive LST as claims

Liquid staking can separate control from loss-bearing.

Control-Exposure Wedge

Definition

An operator can obtain voting power through the pool and sell the liquid claim before deciding whether to attack.

- Collateral absorbs slashing losses first
- Remaining slashing losses are borne by current LST holders
- The operator keeps control but sheds part of the downside

Research Questions

Security

- Does liquid staking weaken deterrence?
- Can rational LST pricing deter attack?
- Which losses can the LST price internalize?

Design

- Does a fee-charging protocol prefer no attack?
- Can bonding collateral alone implement that outcome?
- When are extra tools needed?

Baseline Setup

- Ethereum is the running example, but the mechanism is broader
- Date 1: operators stake through the protocol and sell LST claims
- Date 2: a malicious operator decides whether to attack
- The LST is a claim on future backing, denominated in ETH

This discount model makes pricing, participation, and deterrence transparent.

Main Results

- Liquid staking can weaken deterrence
- Competitive LST pricing partly offsets the wedge
- ETH-wide depreciation remains outside the LST price
- The protocol prefers no attack when feasible
- Collateral alone need not make no attack unique because it has both positive negative effects on security through the endogenous scale of the protocol

Exogenous-Model Environment

- Time has two dates
- At date 1, operators originate LST claims and sell them in competitive secondary markets
- At date 2, a strategic operator that acquired voting power through the pool chooses whether to attack
- Honest operators supply an exogenous pooled stake, denoted by $M^O > 0$

The first model studies how pool loss allocation feeds into attack incentives.

Loss Allocation Inside the Pool

- An attack requires stake $S > 0$, and the attacker operates $M^A = S$
- The bond ratio is $x \in [0, 1]$, so the attacker keeps xS locked and sells $(1 - x)S$ LST claims
- Slashing severity is γ, and ETH depreciation is d
- Honest and attacking stake together give total pooled stake $M^L = M^O + S$

The residual slashing loss borne by the pool is $U = \max\{0, (\gamma - x)S\}$.

Attacker's Problem

Traditional staking benchmark

$$\Delta\Pi_{\text{traditional}} = a + S(1-\gamma)(1-d) - (1+\lambda)S$$

Liquid-staking strategy

$$\Delta\Pi_{\text{liquid}}(p) = a + \Big[(1-d)\max\{0, x-\gamma\} + p(1-x) - 1 - (1-\tau)\lambda\Big]S$$

Liquid staking changes the payoff because the attacker sells the liquid claim before choosing whether to attack.

Profitability Cutoff

Attack condition

Attack is privately optimal when the date-1 LST price is high enough.

$$p_{\text{profitable}} = \frac{1 + (1-\tau)\lambda - (1-d)\max\{0, x-\gamma\} - a/S}{1-x}$$

Attack if and only if $p \geq p_{\text{profitable}}$.

The Backing Value of the LST

- Pre-attack LST supply is $L = (1 - x)M^L$
- Collateral absorbs losses first
- When $x \geq \gamma$, backing stays at par
- When $x < \gamma$, uncovered loss is spread across the non-collateral claim base

$$b = 1 - \max\{0, \gamma - x\} \frac{1}{1 - x} \frac{S}{M^L}$$

Equilibrium Under Competitive Pricing

Pricing

Risk-neutral users price the token at expected backing:

$$p = (1 - q) \cdot 1 + q \cdot b$$

Equilibrium logic

- If $p_{\text{profitable}} \geq 1$, no attack is self-enforcing
- If $b \geq p_{\text{profitable}}$, attack is self-enforcing
- If $1 > p_{\text{profitable}} > b$, the attacker mixes

Exogenous-Scale Equilibrium Regions

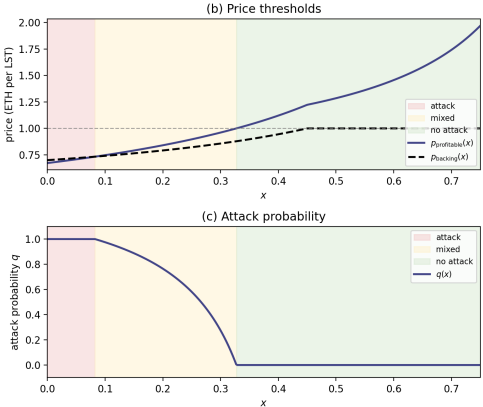

The comparison between backing and profitability yields the no-attack, mixed and attack regions.

Security Implications in the Exogenous Model

Deterrence comparison

$$a^*_{\text{traditional}} = S\Big(\lambda + d + (1-d)\gamma\Big)$$

$$a^*_{\text{liquid}}(x) = S\Big((1-\tau)\lambda + dx + (1-d)\min(\gamma, x)\Big)$$

- Liquid staking weakly lowers the deterrence threshold
- Competitive pricing partly offsets the wedge by lowering resale proceeds
- The price still does not internalize ETH-wide depreciation borne outside the pool

Exogenous-Scale Deterrence Thresholds

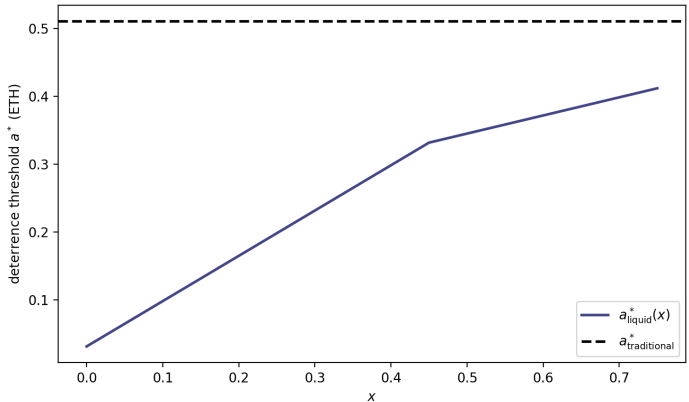

Liquid staking lowers the deterrence threshold relative to traditional staking.

Par-Pricing Benchmark

Benchmark

Fix the date-1 price at the no-attack level $p = 1$.

- The mixed region disappears except at a knife-edge equality
- If $p_{\text{profitable}} \geq 1$, no attack survives
- If $p_{\text{profitable}} < 1$, the outcome collapses to attack

The benchmark isolates the role of price adjustment in the exogenous model.

Why Endogenize Operator Participation?

- If honest operator stakes were exogenous, attack could be prevented by increasing the bond ratio x
- What if honest operator stakes enter endogenously?
- Security affects price
- Price affects participation
- Participation feeds back into scale and therefore into security

In the endogenous scale model, M^O are jointly determined simultaneously with the attack probability q and the date-1 price p.

Endogenous-Model Environment

Demand

Users require an outside-option return μ, so competitive pricing becomes

$$p = \frac{(1-q)\cdot 1 + qb}{1+\mu}$$

Rewards

Ethereum rewards fall with total staked ETH:

$$\lambda = \frac{k}{\sqrt{M^S}}$$

Traditional staking has fixed scale α, while liquid staking scale depends on honest entry.

Honest-Operator Participation

Expected profit

$$\Pi^O = \Big[(1-\tau)\lambda - (1-p)(1-x) - qd\big(x + (1-\tau)\lambda\big)\Big] M^O$$

Equilibrium condition

$$\Pi^O = \mu\big[x + (1-p)(1-x)\big] M^O$$

Honest operators enter until expected profit equals the outside option.

Attacker's Problem with Endogenous Scale

- An attack requires a fraction $f \in (0, 1)$ of total staked ETH
- Traditional stake is fixed at $M^T = \alpha$
- The attacker still uses the liquid-staking strategy from the exogenous model

$$M^S = \frac{1}{1-f}\left[\alpha + M^O\right], \qquad S = \frac{f}{1-f}\left[\alpha + M^O\right]$$

Higher honest participation raises the attack stake required for a successful attack.

Endogenous-Scale Equilibrium Conditions

Unknowns

$$(q, M^O, p)$$

Conditions

- No-arbtrage of LST price: $p = \frac{(1-q)+qb}{1+\mu}$
- Zero-profit of operators: honest operators satisfy the break-even condition
- Indifference of attack: in interior cases, $p = p_{\text{profitable}}$

Interpretation

- Pricing maps attack risk into the LST discount
- Participation maps expected returns into honest stake
- Incentives map the price into attack behavior

Endogenous-Scale Illustration

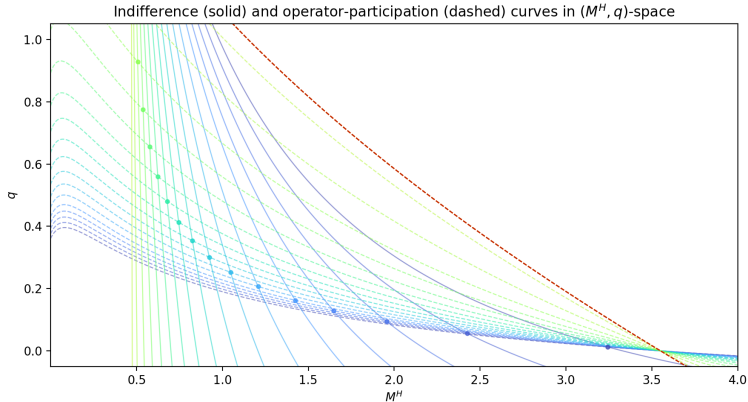

Intersections of the indifference and participation loci pin down operating equilibria.

Endogenous-Scale Equilibrium Classification

- No attack: $q = 0$ and $p = \frac{1}{1+\mu}$
- Attack: $q = 1$ and $p = \frac{b}{1+\mu}$
- Mixed: indifference and participation hold together
- Shutdown: no operating pair (q, M^O) exists

Positive Attack Risk Lowers Scale

Equilibrium property

Let M_0^O denote the honest stake in the no-attack benchmark equilibrium. Any operating equilibrium with $q > 0$ satisfies

$$M^O \leq M_0^O$$

- Attack risk lowers expected operator surplus
- Lower backing and lower price matter when $x < \gamma$
- Restoring break-even requires a smaller operating scale

Bonding Has Competing Effects on Scale

Participation condition along a positive-attack branch

$$(1-\tau)\lambda - \frac{\mu}{1-qd} - \frac{qdx}{1-qd} - \frac{q(1-x)(1-b)}{1-qd} = 0$$

- Higher x lowers socialized slashing
- Higher x increases retained depreciation exposure
- When $x < \gamma$, it also changes the price-based deterrence channel

The net effect of bonding on honest scale is not signed.

Why Multiplicity Can Survive

High-entry branch

- Larger M^O
- Higher attack threshold S
- Lower attack risk

Low-entry branch

- Smaller M^O
- Lower attack threshold S
- Higher attack risk

The same bond ratio can support more than one self-consistent operating branch.

Multiplicity Feedback

$$M^O \uparrow \Rightarrow S \uparrow \Rightarrow a/S \downarrow \Rightarrow q \downarrow$$

This feedback explains why equilibrium selection is difficult when participation is endogenous.

The Protocol's Problem

Objective

The protocol chooses the bond requirement x and fee share τ to maximize fee revenue.

$$\Pi^P = \tau\lambda\left(M^O + \mathbf{1}_{\{A=0\}} M^A\right)$$

- More honest stake raises the fee base
- Attack can also remove fee revenue from attacking stake

Protocol Preference

Main design result

If a no-attack operating equilibrium exists, the protocol prefers it.

$$q > 0 \Rightarrow M^O \leq M_0^O \Rightarrow \Pi^P(q > 0) \leq \Pi^P(q = 0)$$

The key issue is implementation, not preference.

Endogenous-Scale Outcomes

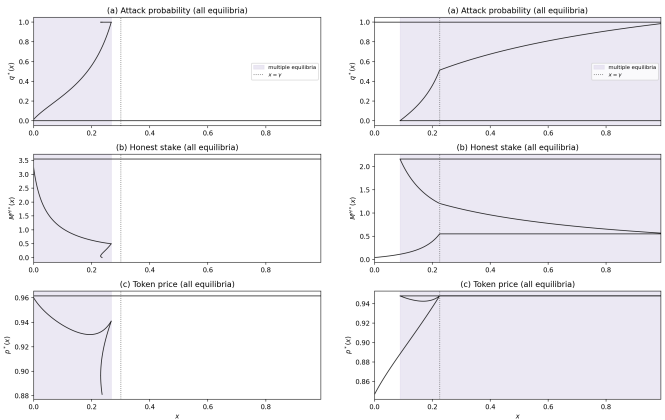

The same bond ratio can support attack, mixing, no attack, or more than one operating branch.

Implementation Problem

- Bonding can shrink the set of attack-feasible equilibria
- It need not select a unique no-attack branch
- Multiplicity can survive even when $x \geq \gamma$
- Robust implementation can require complementary instruments

Examples

Permissioning, screening, delegated-stake caps, reserve funds and external insurance.

Takeaways

1. Liquid staking creates a control-exposure wedge
2. Market pricing softens, but does not solve, the problem
3. Security and scale are jointly determined
4. More collateral is not always better
5. Preference and implementation are different questions

Contribution to the Literature

- We build on the proof-of-stake security literature that studies security as an equilibrium outcome: Saleh (2021), John et al. (2021), Jermann (2024) and Cong et al. (2025).
- The closest paper is Tzinas et al. (2023). They study liquid staking through a principal-agent lens. We instead model a control-exposure wedge and tie competitive LST pricing to attack incentives.
- We complement work on staking pools and liquid staking, such as Gersbach et al. (2022), Jeong et al. (2020), Irresberger et al. (2023), Tang et al. (2024), Gogol et al. (2024a,b) and Carre et al. (2024). We also connect the mechanism to finance settings with separated control and exposure, such as Pennacchi (1988), Bebchuk et al. (2000) and Chemla et al. (2014).