Efficiency vs Resilience: Optimal Collateral With Delegated Validation in Proof of Stake Protocols

Swaminathan Balasubramaniam,[*] Jorge Sabat,[†] and Luana Zaccaria[‡]

First prepared: 2025-08-14

This version: 2026-04-18

Abstract

Proof of stake networks implement incentive-compatible transaction verification by requiring network participants, called validators, to post collateral (stake) that is forfeited upon failure. A growing share of staking, however, is delegated to professional operators. Motivated by the September 2025 Kiln security incident, where professional – but not solo – validators experienced a sharp decline in effectiveness, we develop a model of delegated validation. Validators post collateral (stake) that is forfeited upon failure. Delegation expands access by lowering the effective collateral and operational burden borne by individual validators and can improve efficiency through scale. At the same time, delegation concentrates operational infrastructure, increasing exposure to common shocks and "operational contagion". The protocol therefore faces a tradeoff familiar to the banking literature: tighter collateral requirements strengthen discipline but shift activity toward intermediated providers, raising concentration and correlated losses in stress states. The optimal minimum stake is interior, and collateral requirements and anti-correlation penalties are substitutes: protocols that penalize correlated failures more aggressively can lower their collateral requirements without sacrificing performance.

Keywords: Collateral, Staking, Contagion **JEL:** G10, G12, G41, C55.

[*]NEOMA Business School, Rouen, France; s.balasubramaniam@neoma-bs.fr

[†]Universidad Andrés Bello, Santiago, Chile; jorge.sabat@unab.cl

[‡]LUISS University and EIEF, Rome, Italy; lzaccaria@luiss.it

1 Introduction

Proof of stake blockchains such as Ethereum process financial transactions by requiring network participants, called validators, to post stake and verify the correctness of the transaction record. Validators earn gross protocol revenue when they perform this task accurately and in a timely manner, and they face penalties when they fail. This institutional arrangement admits a natural finance interpretation. Stake functions as pledged collateral, validation is an effort choice akin to independent auditing, and failures impose losses on those who supply the collateral. In its idealized form, proof of stake relies on a large set of independent validators, so that discipline is generated by many decentralized decision makers rather than concentrated in a few entities. In practice, however, validation has become increasingly intermediated. A substantial share of stake is delegated to professional operators and liquid staking arrangements that pool operational expertise, standardize infrastructure, and offer a convenient claim on staking returns.[1] This institutionalization can expand access and improve routine performance, but it also changes the nature of risk by concentrating operational dependencies.

We begin with an empirical observation. On September 9 and 10, 2025, Kiln, a professional staking operator managing roughly 15 percent of Ethereum's professional validation activity, took a large portion of its validator fleet offline after a security breach on its Solana operations. Using daily performance data from Rated Network for 16 professional operators and 322 solo validators, we document a sharp asymmetry in the incident's effects. Professional operators experienced large, statistically unusual drops in effectiveness and revenue, with the most severe losses at Kiln and meaningful declines across other institutional and pooled operators. Solo validators, individuals running independent hardware with no shared infrastructure, exhibited no abnormal performance.[2] The exit queue for validators spiked in the aftermath, concentrated among affected professional operators. These findings establish that a security event external to Ethereum propagated into correlated validator failures through shared professional infrastructure, while decentralized solo

[1] Intermediation is quantitatively important in Ethereum staking. For example, Lido is the largest staking provider and has been around one quarter of staked ETH in recent years, alongside large shares held by centralized exchanges such as Coinbase and Binance; see, for example, CoinDesk (August 14, 2025) and Lido (August 20, 2025). Earlier evidence already pointed to substantial concentration, with Lido and major exchanges together accounting for a majority of staked ETH; see TechCrunch (September 14, 2022, citing Nansen).

[2] The losses were economically significant: Kiln's foregone revenue exceeded $16.7 million over two days, and even the least affected professional category experienced losses exceeding three standard deviations from baseline.

validation continued to function normally. The incident is consistent with a model in which delegated operators face a positive systemic failure probability and manage large blocs of validators with shared infrastructure, so that a single shock causes many positions to fail simultaneously, while solo validators with independent infrastructure are unaffected.

This evidence motivates a model of protocol collateral policy under intermediated validation. The central question is how the minimum stake requirement interacts with the composition of validators when solo operation requires heterogeneous operational skill. Agents differ in solo operational skill, which captures the effective cost of running reliable independent infrastructure, reflecting heterogeneity in technical expertise, time availability, and operational sophistication. Delegation pools operational expertise and standardizes infrastructure, so it does not depend on any individual depositor's skill level. For given protocol parameters, there is a skill cutoff such that agents with sufficiently low skill cost (i.e. higher skilled agents) operate solo and those with higher skill cost delegate, with the margin determined endogenously by the relative payoffs of each mode. Delegation expands access by allowing low-wealth and low skilled agents to participate via pooled operators: while solo operation requires posting the full minimum stake, delegation accepts fractional deposits from any agent with positive wealth, making the staking ecosystem available to a broader population.

Raising the minimum stake strengthens incentives by increasing the cost of failure. At the same time, higher requirements tighten the wealth constraint for solo operation and can shift composition toward pooled operators by reducing the mass of agents who can profitably operate solo, especially among those with high effective solo investment costs. Because delegated validators share infrastructure, software clients, and operational procedures, they are exposed to common shocks that can cause many positions to fail at once. The protocol penalizes such correlated failures through anti-correlation slashing that increases in the concurrent failure share: when a larger fraction of validators fail simultaneously, each failing validator faces a harsher slashing penalty. This penalty acts as a Pigouvian instrument that prices the systemic tail risk of concentration, and the protocol therefore faces a tradeoff between stronger individual incentives and greater exposure to correlated failures through the equilibrium composition of the validator set.

The analysis builds on the corporate finance logic of incentives under limited pledgeability, adapted to a proof of stake environment (Tirole, 2006). Agents are heterogeneous in both wealth

and solo operational skill. Running a validator as a solo operator requires posting the minimum stake and incurring an investment cost that is scaled by the agent's skill type; agents for whom solo operation is costly delegate to an intermediary that pools operational expertise and monitoring across many positions. Delegation enjoys a scale advantage and can raise average performance relative to dispersed solo operation, but it concentrates operational choices and dependencies, raising the likelihood that many validators fail together. The protocol's collateral policy shapes both individual incentives and the allocation of stake across solo and delegated validation through its interaction with the skill distribution and the wealth distribution.

The model delivers four main results. First, in the solo sector, stronger collateral at risk improves incentives and increases the reliability investment chosen by a validator conditional on operating. The mechanism is the standard skin-in-the-game channel: a validator who stands to lose more collateral upon failure internalizes a larger share of the social cost of underperformance and therefore spends more on redundancy, monitoring, and uptime.

Second, delegation can improve routine performance because the intermediary spreads the fixed costs of operational infrastructure across many validator positions. This scale economy is the primary efficiency rationale for intermediation in staking, analogous to the diversification benefit in Diamond (1984): a single well-resourced monitor can achieve what many dispersed small operators cannot. When the effective pool scale, discounted for the systemic failure probability, is large enough relative to the inverse of the lowest solo skill cost, the per-validator cost of investment falls sufficiently that the monitor optimally chooses a level of reliability that strictly exceeds what even the most skilled individual solo operator would select.

Third, the two forces generate a tradeoff that the protocol must navigate through its collateral policy. Raising the minimum stake strengthens the skin-in-the-game channel, but it simultaneously tightens the wealth constraint for solo operation. Because delegated validators are exposed to common shocks that cause many positions to fail at once, the concentration of pooled validation raises the expected loss from systemic events. Furthermore, collateral is a blunt instrument: private agents choose between solo and delegated operation based on their own payoffs, without internalizing the protocol-level effect of changing network composition, so the equilibrium skill cutoff can be too low and the delegated sector too large relative to the social optimum. Under conditions ensuring that liveness strictly exceeds overhead costs at some interior collateral level,

the protocol objective may therefore favor an interior minimum stake requirement. Unlike traditional banking, where contagion typically propagates through interbank lending or deposit linkages (Allen and Gale, 2000), the source of correlation in this setting is fundamentally operational. The standardization of liquid staking infrastructure generates what we term operational contagion even in the absence of direct financial claims between operators.

Fourth, collateral and anti-correlation penalties are substitutes in incentive provision, holding composition fixed. When the skill cutoff is treated as invariant in the penalty strength parameter, stronger anti-correlation penalties raise the cost of failure for both solo and delegated operators, improving investment incentives without distorting the allocation between modes. In that benchmark, less collateral is needed to achieve a given reliability level, so the optimal minimum stake is decreasing in penalty strength. This substitutability has a direct policy implication: protocols that adopt anti-correlation penalties can lower their collateral requirements without sacrificing validator performance, potentially broadening access to solo validation. When composition responds endogenously, penalties reduce delegated attractiveness through two distinct channels — a common-slashing channel that operates even in non-systemic states, and a systemic channel that prices correlated failure risk directly — shifting marginal agents toward solo operation. If this composition shift is large enough to outweigh the fixed-composition incentive gain, the jointly optimal penalty strength is interior.

The paper contributes a framework for thinking about protocol collateral policy when validation is intermediated and agents differ in the cost of operating independently. The key implication is that collateral policy does not only discipline individual operators; it also determines who validates and with what composition of solo versus pooled positions. That composition effect, mediated by the skill distribution and the wealth distribution, can amplify or mitigate correlated failures. More broadly, proof of stake provides a new empirical setting in which collateral requirements, operational heterogeneity, and correlated default risk interact, and in which familiar corporate finance and banking tradeoffs can be studied with high-frequency operational data.

1.1 Related literature

Our paper contributes to the burgeoning literature on the economic underpinnings and systemic risks of Proof-of-Stake (PoS) ecosystems. We build on the foundational work of Lehar and Parlour (2022), who document the systemic fragility inherent in decentralized lending markets due to clustered liquidations and arbitrage limits. While they focus on DeFi lending platforms like Aave and Compound, we extend this logic to the staking layer, where the protocol's collateral requirement serves as a primary defense against operator failure. Our finding that an interior optimum may exist for minimum stake requirements under conditions complements Cong et al. (2025), who provide a continuous-time framework for tokenomics where the aggregate staking ratio shapes platform productivity and price dynamics. While Cong et al. (2025) emphasize the asset pricing implications of staking, our model highlights the underlying trade-off between individual operator discipline and the convex systemic losses that arise from correlated failures in delegated environments.

Furthermore, we relate to recent studies on the principal-agent problems and reputational incentives in liquid staking and restaking. Tzinas and Zindros (2023) first formalized the principal-agent dilemma in liquid staking, noting that the fungibility of delegated stake can misalign the interests of delegators and validators. Our model provides a structural basis for these misalignments by showing how heterogeneous operational skill and pooled monitoring shape the composition of the validator set, echoing the classical logic of Tirole (2006). In the context of liquid restaking, Mu et al. (2025) argue that reputation acts as a critical safeguard when physical collateral is reused across multiple networks. We complement this by showing that even with reputation-based monitoring, the standardization of liquid staking tokens (LSTs) increases correlation parameters, thereby exacerbating coordination risk. This perspective also complements Carré and Gabriel (2024), who caution that the "capital efficiency" of liquid staking can worsen blockchain security if DeFi lending platforms are not properly designed to handle the resulting shifts in validator incentives.

Our theoretical framework also builds upon the classic delegated monitoring literature, most notably Diamond (1984), which establishes that financial intermediaries can minimize the costs of auditing borrowers by diversifying across independent risks. We adapt this logic to the Proof-of-Stake context, where the protocol acts as a principal delegating the verification of state transitions to

validators. Intermediation pools operational expertise and monitoring; the key friction is heterogeneous operational capability and correlated operational shocks. While Diamond (1984) emphasizes the "diversification benefit" of a large portfolio, we highlight a counteracting force: the standardization of liquid staking tokens and the shared reliance on common software clients introduce a high degree of correlation. This relates to the survey by Fleckinger et al. (2024) on the trade-offs between relative performance evaluation and collective incentives. In their taxonomy, the desirability of competition versus cooperation hinges on the correlation of agents' outputs. We extend this by incorporating the insights of Kvaløy and Olsen (2019), who demonstrate that when agents' outputs are correlated, relational contracts and shared monitoring systems can facilitate joint failure rather than independent discipline. In our model, this correlation is the critical parameter that creates convex systemic losses.

Finally, we contribute to the extensive literature on contagion and systemic risk among financial intermediaries. Seminal work by Allen and Gale (2000) explores how interbank markets create channels for contagion through direct financial linkages such as loans and deposit claims. More recently, Goldsmith-Pinkham and Yorulmazer (2010) and Choi et al. (2024) document how informational externalities and liquidity dry-ups lead to correlated failures during episodes like the Northern Rock and Silicon Valley Bank runs. In the context of Ethereum, Mu et al. (2025) highlight contagion risks due to restaking. While these papers focus on balance sheet exposure, the source of risk in our framework is fundamentally different. Our model identifies correlation stemming from shared infrastructure and monitoring systems rather than direct credit linkages. This operational contagion implies that the concentration of staking power in a few standardized LSTs or professional operators increases the probability of simultaneous validator failures, necessitating higher minimum stake than would be required with independent, heterogeneous operators.

The remainder of the paper is organized as follows. Section 2 presents a brief background on validation in Proof of Stake protocols. Section 3 presents empirical motivation for modeling correlated failures. Section 4 describes the model and characterizes equilibrium behavior under solo and delegated validation, along with the protocol choice of minimum stake. Section 5 describes anti correlation penalties. Section 6 concludes.

2 Background

2.1 Proof of Stake, Collateral, and Delegation

Proof of stake blockchains such as Ethereum settle transactions through a decentralized verification process. Network participants, called *validators*, lock a quantity of the network's native token (the "stake") as a form of pledged collateral and then run software that checks the accuracy of proposed transactions. When a validator performs this task correctly and in a timely manner, the protocol pays a reward. When a validator fails, whether through downtime, software errors, or misconduct, the protocol destroys a fraction of the posted collateral through a penalty known as *slashing*. The economics are straightforward: collateral at risk gives each validator a financial incentive to invest in reliable operations, much as margin requirements discipline a leveraged trader or capital requirements discipline a bank.

In its idealized form, proof of stake envisions a large and diverse set of independent validators, each running their own hardware and making their own operational decisions. This dispersion is valuable because it limits the damage any single point of failure can inflict on the network. In practice, however, the economics of validation have pushed the sector toward intermediation. Ethereum currently requires each validator position to carry a minimum stake of 32 ETH (roughly $100,000 at recent prices), and the fixed costs of maintaining reliable infrastructure, including redundant hardware, monitoring systems, secure key management, and continuous software updates, create meaningful barriers to entry for individuals.

The result has been the rapid growth of *delegated staking*. In a delegated arrangement, token holders deposit their ETH with a professional operator or a pooling protocol (such as Lido, Rocket Pool, or a centralized exchange like Coinbase), which bundles these deposits, finances the required collateral positions, and manages the underlying validator infrastructure on behalf of depositors. In return, depositors receive either a liquid claim on the pooled stake (a "liquid staking token") or a proportional share of staking returns net of fees. The intermediary bears the operational burden and, in many cases, the slashing risk. Delegation is now quantitatively significant: Lido alone accounts for roughly a quarter of all staked ETH, and the combined share of Lido, Coinbase, Binance, and a handful of institutional operators exceeds two thirds of the network's validator set.[3]

[3] See, for example, CoinDesk (August 14, 2025) and Lido (August 20, 2025). Earlier evidence already pointed

This concentration introduces a tension that is familiar from banking and financial intermediation more broadly. Delegation offers clear benefits. It lowers the effective wealth threshold for participation, allowing small token holders to earn staking returns without purchasing and maintaining their own infrastructure. It also allows professional operators to spread fixed costs across thousands of validator positions and to invest in operational quality at a level that would be uneconomical for an individual running one or two validators from home. During normal periods, these scale and monitoring advantages tend to produce higher average performance among professional operators than among solo stakers.

The cost of these benefits is a form of concentration risk. When a professional operator manages tens of thousands of validators, those validators typically share the same software clients, the same cloud infrastructure, the same monitoring and alerting systems, and the same incident response procedures. A disruption that would be purely idiosyncratic if it struck a single home staker can become a correlated failure when it strikes an operator whose infrastructure underpins a large share of the network. The potential for such clustered failures is precisely the systemic risk channel that our theoretical framework formalizes. The model in Section 4 shows that the protocol's collateral policy shapes not only individual incentives to invest in reliability, but also the equilibrium composition of the validator set, and that this composition effect creates a tradeoff between the discipline benefits of higher collateral and the concentration costs of pushing more stake toward intermediated providers.

3 Empirical Motivation

This section documents an episode in Ethereum staking that motivates the theoretical framework developed in Section 4. We describe the incident, present the data and identification strategy, and report the main empirical findings. The results establish several stylized facts about the relationship between operator concentration and correlated validator failures that the model is designed to explain.

to substantial concentration, with Lido and major exchanges together accounting for a majority of staked ETH; see TechCrunch (September 14, 2022, citing Nansen).

3.1 The September 2025 Infrastructure Incident

The September 2025 security breach on Kiln's Solana operations serves as an exogenous shock to Ethereum's validation infrastructure. This event is particularly instructive because the disruption originated on a separate blockchain, meaning it was unanticipated by Ethereum participants and unrelated to the protocol's internal mechanics or prevailing market conditions. We exploit this natural experiment to study the relationship between operator concentration and correlated failures. As discussed below, the incident created a sharp divergence between professional and solo validators. While professional operators experienced statistically unusual drops in effectiveness, solo validators, who operate independent hardware with no shared infrastructure dependencies, maintained stable performance throughout the window. The subsequent spike in the validator exit queue, concentrated among affected professional operators, further suggests that market participants viewed the disruption as economically significant.

Figures 1 through 3 document the incident. Figure 1 plots daily validator effectiveness (a composite of attestation accuracy, inclusion speed, and uptime, scored 0 to 100 percent) for three categories of professional operators: Kiln, other institutional operators (Kraken, Everstake, Figment), and Lido sub-operators (P2P, Chorus One, and others). Kiln's effectiveness fell sharply on September 9 and continued to decline on September 10. Other professional categories also experienced losses, though smaller in magnitude and quicker to reverse.

Figure 2 introduces the key comparison. Solo validators, individuals operating independent hardware from home or small data centers, maintained stable performance throughout the incident. Because they share no infrastructure or operational dependencies with professional firms, they were insulated from the disruption. Figure 3 shows a sharp spike in the Ethereum validator exit queue immediately following the incident, concentrated among validators associated with affected professional operators. The exit response suggests that market participants viewed the disruption as economically meaningful.

3.2 Data and Measurement

We obtain daily validator effectiveness data from Rated Network, which computes standardized performance metrics aggregating attestation accuracy, inclusion delay, and uptime.[4] Our sample covers 16 professional node operators and 322 solo validators during August and September 2025. Professional operators are observed for 60 days (August 1 through September 30), yielding 960 operator-day observations across three categories: institutional operators (Kiln, Kraken, Everstake, Figment, Staked.us) managing 154,451 validators; Lido sub-operators (P2P, Chorus One, Attestant, and seven others) managing 97,340 validators; and centralized exchange operators (Coinbase Cloud) managing 10,476 validators. Solo validators are drawn via stratified random sampling from deposit addresses controlling 1 to 100 validators with no affiliation to known staking services, observed for 11 days (September 5 through 15), yielding 3,542 validator-day observations. Table 1 presents sample composition.

Our primary outcome variable is the effectiveness drop, defined as the deviation of incident-period performance from baseline:

$$\text{Drop}_{it} = \text{Effectiveness}_{i,\text{baseline}} - \text{Effectiveness}_{it}, \tag{1}$$

where i indexes validators or operators and t indexes days. Validator effectiveness is a composite score (0 to 100%) combining attestation correctness, inclusion speed, and uptime.[5] We use September 8, 2025, the last normal day before the incident, as the baseline. We measure effectiveness drops over three temporal windows: the total drop (September 8 to September 10), the first-day drop (September 8 to September 9), and the second-day drop (September 9 to September 10). The decomposition allows us to separate common shocks from operator-specific recovery dynamics.

Economic losses are measured analogously for consensus-layer priority fees, execution-layer MEV rewards, and total validator rewards. All revenue measures are per validator in Gwei and standardized by baseline standard deviation.[6]

[4] See https://www.rated.network for methodology.
[5] See Rated Network documentation for precise weighting methodology.
[6] 1 ETH = 10^9 Gwei.

3.3 Identification

Our identification strategy relies on the fact that solo validators provide a clean control group. Because they share no infrastructure with professional firms, their performance serves as a counterfactual for what would have occurred in the absence of shared infrastructure dependence. Under a parallel trends assumption, which we support by documenting stable and comparable performance for both groups prior to the incident, the difference in effectiveness drops between professional and solo validators identifies the causal impact of operational concentration. This approach allows us to separate protocol-level shocks from failures specific to the professional sector.

For professional operators, we estimate weighted cross-sectional regressions:

$$\text{Drop}_i = \beta_0 + \beta_1 \,\text{Institutional}_{ex\text{-}Kiln,i} + \beta_2 \,\text{Lido}_i + \beta_3 \,\text{CEX}_i + \varepsilon_i, \tag{2}$$

where i indexes operators and Kiln is the omitted category. Observations are weighted by validator count; standard errors are heteroskedasticity-robust. To compare professional and solo validators, we estimate a validator-level panel:

$$\text{Drop}_{it} = \alpha + \sum_c \delta_c \mathbb{1}[\text{Category}_i = c] + \gamma \mathbb{1}[\text{Date}_t = \text{Sept 10}] + \epsilon_{it}, \tag{3}$$

where i indexes 322 solo plus 16 professional validators and $t \in \{\text{Sept 9, Sept 10}\}$. Solo validators are the reference category; δ_c measures the differential impact for professional category c.

Three features of the incident support a causal interpretation. First, the shock was unanticipated and unrelated to validator behavior or ETH market conditions. Second, solo validators provide a natural control group, since they share no infrastructure with professional operators. Third, the effect appears discretely on September 9 and 10 with recovery beginning September 11, consistent with a sharp shock rather than gradual deterioration. Under a parallel trends assumption, differences in effectiveness drops between professional and solo validators identify the causal effect of shared infrastructure dependence. We provide supporting evidence by showing that both groups exhibited stable, comparable performance from August 1 through September 8.

3.4 Heterogeneous Losses Across Professional Operators

Table 2 presents the primary performance analysis. To make losses comparable across categories with different baseline volatility, we express incident-period drops as z-scores relative to each category's standard deviation over the 58 non-incident days.

Kiln experienced the most severe degradation: a 1.66 percentage point effectiveness loss, or 3.19 standard deviations from baseline ($p < 0.003$). Under normal conditions such a disruption would occur less than 0.3 percent of the time. Lido sub-operators lost 1.06 percentage points (2.41 standard deviations, $p < 0.013$), other institutional operators lost 0.91 percentage points (1.90 standard deviations, $p < 0.046$), and CEX validators lost 0.65 percentage points (1.07 standard deviations). Panel B confirms that Kiln's losses significantly exceeded those of every other professional category: institutional operators lost 0.75 percentage points less ($p < 0.013$), Lido 0.60 points less ($p < 0.046$), and CEX 1.01 points less ($p < 0.013$). For an operator managing 50,000 validators, a 0.75 percentage point effectiveness gap translates to roughly 375 missed attestations per day.

Table 3 translates these losses into economic terms. Kiln validators lost 90.9 Gwei per validator in total rewards, comprising 52.3 Gwei in priority fees and 38.6 Gwei in MEV rewards (6.35 standard deviations from baseline). Aggregated across 57,344 validators, total losses exceeded 5,200 ETH, or roughly $16.7 million at the prevailing price. Other professional categories lost 58 to 72 Gwei per validator, all exceeding 3 standard deviations. Kiln's differential losses (29.2 Gwei more than institutional operators, 33.2 Gwei more than CEX, both $p < 0.01$) point to operator-specific factors compounding the common shock. Notably, losses were comparable across consensus-layer and execution-layer revenue streams, which operate through distinct technical pathways. Their simultaneous deterioration indicates that the incident impaired core validator operations rather than only specialized functions.

Table 4 decomposes the impact by day. On September 9, all categories experienced initial losses (0.43 to 2.28 percentage points, 21 to 42 Gwei), consistent with a common shock propagating through professional infrastructure. On September 10, paths diverged sharply. Institutional operators stabilized (0.32 percentage points, 18.1 Gwei), CEX validators recovered (effectiveness improved 1.63 percentage points), while Kiln's losses accelerated (1.23 percentage points, 68.4 Gwei). Kiln's day-2 differential relative to other institutional operators (0.91 percentage

points, 50.3 Gwei, both $p < 0.001$) indicates that an operator-specific problem compounded the initial common disruption. The pattern of uniform initial impact followed by heterogeneous recovery suggests two distinct forces at work: a shared vulnerability across the professional sector and variation in firm-level resilience.

3.5 Solo Validators and the Evidence for Correlated Failure

The analysis above establishes that losses were large and heterogeneous across professional operators. Whether these losses reflect a problem specific to professional infrastructure or a broader network-wide disruption is the question that solo validators can help answer. Solo validators operate independently, with no shared infrastructure, software dependencies, or monitoring systems linking them to professional firms. If the incident reflected a protocol-level or network-wide failure, solo validators should have been affected as well.

A measurement challenge is that the two groups are observed over different windows: 11 days for solo validators versus 60 for professionals. We address this by expressing each group's incident losses as z-scores relative to its own baseline standard deviation, so that the comparison captures how anomalous the incident was for each group given its typical variation.

Table 5 presents the results. Solo validators' incident losses (0.03 percentage points) represent 0.07 standard deviations from their 9-day baseline ($p = 0.94$), with a 95 percent confidence interval of $[-0.86, +0.92]$ percentage points. The tight baseline distribution (SD = 0.45 percentage points) and near-zero incident losses indicate genuine operational stability rather than limited statistical power. All professional categories, by contrast, exceeded 1.90 standard deviations. Kiln's losses exceeded solo validators' by 1.63 percentage points (Cohen's $d = 3.50$); even institutional operators, the least affected professional category, exceeded solo validators by 0.88 percentage points ($d = 1.89$).

Table 6 confirms this pattern in economic terms. Solo validators lost 1.3 Gwei per validator (0.07 standard deviations, $p = 0.94$), stable across all revenue components. Professional categories lost 56 to 90 Gwei more per validator ($d = 3.40$ to 5.50). Kiln's 89.6 Gwei differential, aggregated across 57,344 validators, amounts to roughly 5,140 ETH ($16.4 million) in excess losses. Even institutional operators' smaller differential (60.4 Gwei) implies aggregate excess losses of approximately 5,810 ETH ($18.6 million) across 96,107 validators. These figures exclude opportunity

costs and reputational effects.

Table 7 decomposes losses by day. Solo validators maintained z-scores below 0.05 on both days, ruling out network-wide explanations. Professional operators displayed the same two-stage pattern: common initial disruption on September 9, divergent recovery on September 10. Kiln's day-2 differential relative to solo validators (2.33 effectiveness standard deviations, 4.74 reward standard deviations) quantifies the combined effect of persistent Kiln-specific problems and the broader common shock across professional infrastructure.

3.6 Takeaways for the Model

The empirical analysis establishes four facts that jointly motivate the theoretical framework in Section 4.

Fact 1: The disruption was specific to professional infrastructure. Solo validators experienced no abnormal performance during the incident, ruling out a network-wide consensus problem. The losses were concentrated entirely within the professional staking sector, where operators share common software clients, cloud providers, monitoring systems, and incident response procedures.

Fact 2: Correlated failure extends beyond the directly affected operator. Although Kiln was the epicenter, all professional operator categories experienced statistically unusual losses. Shared operational dependencies transmitted the shock across firms, generating clustered underperformance well beyond Kiln itself.

Fact 3: Scale and specialization create a tradeoff. Professional operators deliver higher routine performance through scale economies and specialized infrastructure. Yet the same concentration of operational choices that enables this performance advantage also concentrates risk. Solo validators, individually less efficient under normal conditions, proved collectively more resilient during the stress episode.

Fact 4: The economic magnitudes are large enough to matter for protocol design. Two-day losses of 60 to 90 Gwei per validator represent approximately 0.15 to 0.22 percent of annual staking yields (assuming 32 ETH stake at a 4 percent return). While modest for a single event, incidents of this scale, if they recur, would materially affect the calculus of how much collateral the protocol should require and how it should penalize correlated failures.

These facts point to a tradeoff that individual-level incentive analysis alone cannot capture. The protocol's collateral requirement shapes not only how much each validator invests in reliability, but also the equilibrium split between solo and delegated validation, and hence the network's exposure to correlated failures. The model in Section 4 formalizes this tradeoff and characterizes the resulting optimal collateral policy.

4 Model

4.1 Environment

We study a proof-of-stake blockchain in which transactions are processed by a set of *validators*. A validator is a network participant that runs software to verify and propose blocks, and earns gross protocol revenue $R > 0$ when it performs successfully in a period and 0 otherwise. Each validator position requires posting the protocol's minimum collateral (stake) C. Because delegated validator count equals total delegated stake divided by the collateral requirement, we restrict attention to $C \in [\underline{C}, \bar{w}]$ with $\underline{C} > 0$, the smallest technologically meaningful stake unit. Reliable operation also demands investment in infrastructure quality, which we model as a non-monetary investment decision.

There is a unit mass of agents indexed by i, each characterized by two independently drawn parameters.[7] *Staking wealth* w_i is drawn i.i.d. from the uniform distribution on $[0, \bar{w}]$, so $H(C) = C/\bar{w}$ and $1 - H(C) = (\bar{w} - C)/\bar{w}$. Uniformity yields tractable closed-form expressions for equilibrium composition; the qualitative results do not depend on it. *Solo skill cost* $s_i \in [\underline{s}, \bar{s}]$ is drawn i.i.d. from the uniform distribution on $[\underline{s}, \bar{s}]$, with $0 < \underline{s} < \bar{s}$, so $F(s) = (s - \underline{s})/(\bar{s} - \underline{s})$, independently of w_i. For reference, the truncated wealth moments are

$$\mathbb{E}[w\,\mathbf{1}\{w \geq C\}] = \frac{\bar{w}^2 - C^2}{2\bar{w}}, \qquad \mathbb{E}[w\,\mathbf{1}\{w < C\}] = \frac{C^2}{2\bar{w}}, \qquad \mathbb{E}[w] = \frac{\bar{w}}{2}.$$

Agents maximize expected payoffs and may supply validation services either by operating a validator directly (*solo*) or by delegating their stake to a pooled provider. Solo participation requires

[7] We work with a continuum of agents. "Validator counts" are therefore validator densities (positions per unit mass of agents), and all staked-ETH aggregates are per-unit-mass expectations.

$w_i \geq C$: the agent must post the full collateral. Because the skill cost $s_i I$ reflects time and expertise rather than a cash outlay, it does not require additional wealth beyond the collateral. Delegation is accessible to any agent with $w_i > 0$: the pool aggregates fractional deposits, so each depositor contributes w_i and holds a proportional claim on pooled positions, with no minimum deposit beyond a negligible entry threshold. Agents with $w_i \in (0, C)$ can delegate but cannot operate solo.

This section studies baseline investment incentives, equilibrium composition, protocol choice, and welfare when delegated validation is exposed to a reduced-form common-shock probability $\rho_D \in (0, 1)$. Anti-correlation penalties are introduced as a policy instrument in Section 5.

4.2 Solo and delegated validation

Solo validation. A solo validator with skill cost s_i chooses an investment level $I \geq 0$ in operational reliability. To obtain closed-form investment rules, we assume

$$p(x) = 1 - e^{-ax}, \qquad a > 0. \tag{4}$$

This exponential specification is increasing, strictly concave, satisfies $p(0) = 0$ and $\lim_{x\to\infty} p(x) = 1$, and yields an explicit first-order condition. If the validator fails, the protocol imposes a baseline slashing fraction $\underline{\phi} \in (0, 1)$, so the slashing loss is $\underline{\phi}\, C$. The solo validator's expected payoff from investment I is

$$U_S(I; C, s) = p(I)(R + \underline{\phi}\, C) - \underline{\phi}\, C - sI. \tag{5}$$

Proposition 1 (Optimal solo investment)**.** *Under* (4)*, for any* (C, s) *with* $s > 0$*, the solo problem* (5) *has a unique optimum* $I_S^\star(C, s) \geq 0$*. Let* $\mu(C) \equiv a(R + \underline{\phi}\, C)$*.*

1. *If* $\mu(C) > s$*, the optimum is interior:*

$$I_S^\star(C, s) = \frac{1}{a} \log\left(\frac{\mu(C)}{s}\right), \tag{6}$$

with equilibrium success probability $p(I_S^\star) = 1 - s/\mu(C)$ *and maximized payoff*

$$U_S^\star(C, s) = R - \frac{s}{a}\left(1 + \log\left(\frac{\mu(C)}{s}\right)\right). \tag{7}$$

2. *If* $\mu(C) \le s$, *then* $I_S^\star = 0$ *and* $U_S^\star(C, s) = -\underline{\phi}\, C$.

3. *On the interior region,* $I_S^\star$ *and* $p(I_S^\star)$ *are strictly increasing in* C *and strictly decreasing in* s.

Proof is in Appendix A.1.

When interior, the first-order condition $ae^{-aI}(R + \underline{\phi}C) = s$ gives the log formula. The success probability $1 - s/\mu(C)$ is increasing in C and decreasing in s: higher collateral raises the skin-in-the-game incentive; higher skill cost compresses investment. On the interior region, the maximized payoff $U_S^\star = R - (s/a)(1 + \log(\mu(C)/s))$ is increasing in C and decreasing in s. At the corner, the agent does not invest and the payoff equals $-\underline{\phi}\, C$, which is decreasing in C.

Delegated validation. Delegated validation is organized by a monitor who operates a pool of validators funded by delegated deposits. The monitor chooses a pool-level investment $J \ge 0$ that benefits all delegated validators, with the cost J shared across N_D positions, so the average cost per position is J/N_D. The delegated technology does not depend on any individual depositor's skill. In a systemic period, which occurs with probability $\rho_D \in (0, 1)$, a common infrastructure shock causes the entire pool to fail with certainty; in a non-systemic period, occurring with probability $1 - \rho_D$, each delegated validator succeeds with probability $p(J)$.[8] The reduced-form common-shock probability ρ_D captures correlated operational risk inherent to pooled validation; it is treated as exogenous to the level of delegation. The expected payoff per delegated validator position is

$$U_D(J; C, N_D) = \big(1 - \rho_D\big)\Big(p(J)(R + \underline{\phi}\, C) - \underline{\phi}\, C\Big) - \rho_D\, \underline{\phi}\, C - \frac{J}{N_D}. \tag{8}$$

Proposition 2 (Delegated investment and the scale economy)**.** *Under* (4)*, for any* (C, N_D) *with* $N_D > 0$*, the monitor's problem* (8) *has a unique optimum* $J^\star(C, N_D) \ge 0$. *Let* $\nu(C, N_D) \equiv a(1 - \rho_D)N_D(R + \underline{\phi}\, C)$.

1. *If* $\nu(C, N_D) > 1$*, the optimum is interior:*

$$J^\star(C, N_D) = \frac{1}{a} \log\big(\nu(C, N_D)\big), \tag{9}$$

[8] We abstract away from delegated pool fees.

with success probability $p(J^\star) = 1 - 1/\nu(C, N_D)$ *and maximized payoff*

$$U_D^\star(C, N_D) = (1 - \rho_D)R - \rho_D \underline{\phi}\, C - \frac{1}{aN_D}\big(1 + \log\big(\nu(C, N_D)\big)\big)\,. \tag{10}$$

2. *If* $\nu(C, N_D) \le 1$*, then* $J^\star = 0$ *and* $U_D^\star(C, N_D) = -\underline{\phi}\, C$*.*

3. *On the interior region,* $J^\star$ *is strictly increasing in* N_D *and strictly increasing in* C*.*

4. *If both* $J^\star$ *and* $I_S^\star(C, \underline{s})$ *are interior, then*

$$p(J^\star) > p\big(I_S^\star(C, \underline{s})\big) \quad \Longleftrightarrow \quad (1 - \rho_D)\, N_D > \frac{1}{\underline{s}}. \tag{11}$$

Proof is in Appendix A.2.

The composite $\nu(C, N_D) = a(1 - \rho_D)N_D(R + \underline{\phi}C)$ is the effective scale-adjusted marginal return to pool investment. Larger pools spread fixed monitoring costs over more positions, raising ν and hence $J^\star$: this is the scale economy. The factor $1 - \rho_D$ discounts delegated returns by the probability of a systemic failure period, but this discount is exogenous and does not depend on the level of delegation. Part (iv) gives the exact threshold for delegated investment to outperform even the best solo operator: risk-adjusted pool scale $(1 - \rho_D)N_D$ must exceed $1/\underline{s}$.

4.3 Equilibrium composition

Denote maximized payoffs by $U_S^\star(C, s)$ from (7) and $U_D^\star(C, N_D)$ from (10). On the interior solo region, $U_S^\star$ is strictly decreasing in s (envelope theorem: $\partial U_S^\star/\partial s = -I_S^\star < 0$); at the corner it equals $-\underline{\phi}C$ and is independent of s. Since $U_D^\star$ does not depend on s, the choice between solo and delegation for agents with $w_i \ge C$ is characterized by a skill cutoff s^* satisfying

$$U_S^\star(C, s^*) = U_D^\star(C, N_D). \tag{12}$$

This cutoff is unique whenever $U_D^\star$ lies strictly between $U_S^\star(C, \underline{s})$ and $U_S^\star(C, \bar{s})$; corner cutoffs apply otherwise. Agents with $s_i \le s^*$ operate solo; those with $s_i > s^*$ delegate. We focus on interior equilibria where $s^* \in (\underline{s}, \bar{s})$ and $U_D^\star \ge 0$ so that agents with $w_i \in (0, C)$ delegate rather than abstain.

For a given collateral level $C \in [\underline{C}, \bar{w})$ and candidate delegated scale $n \geq 0$, the induced cutoff $s^*(C, n)$ solves (12) with $N_D = n$. The corresponding masses and stake are

$$N_S(C, n) = \frac{\bar{w} - C}{\bar{w}} \cdot \frac{s^*(C, n) - \underline{s}}{\bar{s} - \underline{s}}, \tag{13}$$

$$S_D(C, n) = \left(1 - \frac{s^*(C, n) - \underline{s}}{\bar{s} - \underline{s}}\right) \frac{\bar{w}^2 - C^2}{2\bar{w}} + \frac{C^2}{2\bar{w}}, \tag{14}$$

$$\Phi_C(n) = \frac{S_D(C, n)}{C}. \tag{15}$$

An equilibrium is a fixed point $n^* = \Phi_C(n^*)$, with $N(C) = N_S(C) + N_D(C)$.

Proposition 3 (Equilibrium composition)**.** *Fix $C \in [\underline{C}, \bar{w})$. The map Φ_C defined by (14)–(15) maps $[0, \bar{w}/(2C)]$ into itself and is continuous. Hence there exists an equilibrium delegated scale $N_D^*(C) \in [0, \bar{w}/(2C)]$. Once selected, all equilibrium objects $s^*(C)$, $N_S(C)$, $N_D(C)$, and $N(C)$ are well defined as functions of C alone.*

Proof is in Appendix A.3.

Higher collateral affects composition through two channels: it directly tightens the wealth threshold for solo participation, and it shifts the payoff comparison between solo operation and delegation through the indifference condition (12).

4.4 Protocol problem

The protocol designer chooses $C \in [\underline{C}, \bar{w}]$ to maximize expected network performance net of overhead costs. Performance is measured by *effective liveness*: the expected mass of validators that perform successfully.

Under the uniform skill distribution and the interior success probability $p(I_S^\star) = 1 - s/\mu(C)$ from Proposition 1, the average solo success probability among equilibrium solo operators is

$$\bar{p}_S(C) = 1 - \frac{\underline{s} + s^*(C)}{2\mu(C)}. \tag{16}$$

Effective liveness is

$$\mathcal{L}(C) \equiv N_S(C)\, \bar{p}_S(C) + (1 - \rho_D)\, N_D(C)\, p\big(J^\star(C)\big), \tag{17}$$

where $J^\star(C) \equiv J^\star(C, N_D(C))$. To capture increasing protocol congestion and communication costs, we assume a quadratic overhead cost

$$g(N) = \frac{\chi}{2}N^2, \qquad \chi > 0. \tag{18}$$

The protocol's problem is

$$\max_{C \in [\underline{C}, \bar{w}]} \quad \mathcal{V}(C) \equiv \mathcal{L}(C) - \frac{\chi}{2}N(C)^2. \tag{19}$$

Raising C improves investment incentives for any fixed operator, but it also tightens the solo participation constraint and can shift activity toward the delegated sector. Because delegated validation is exposed to a reduced-form common-shock probability $\rho_D > 0$, reallocation toward delegation may reduce effective liveness even when delegated monitoring benefits from scale.

Proposition 4 (Interior optimal collateral)**.** *Suppose the equilibrium selection is continuous on* $[\underline{C}, \bar{w}]$. *If:*

1. $\mathcal{V}'(\underline{C}) > 0$,
2. *there exists* $\bar{C} \in (\underline{C}, \bar{w})$ *such that* $\mathcal{V}'(C) < 0$ *for all* $C \in [\bar{C}, \bar{w}]$,

then the protocol problem admits at least one interior maximizer $C^\star \in (\underline{C}, \bar{w})$.

Proof is in Appendix A.4.

Remark 1 (Economically interpretable sufficient conditions for interiority)**.** *The following are economically interpretable sufficient conditions that imply (A1)–(A2) of Proposition 4:*

1. $\mu(\underline{C}) > \bar{s}$, *so that all skill types have interior investment at* $\underline{C}$ *and the direct incentive effect of collateral is positive there; and*
2. *on some interval* $[\bar{C}, \bar{w}]$, *the participation loss and the overhead effect dominate the remaining incentive gain, so that* $\mathcal{V}'(C) < 0$ *throughout that interval.*

A sufficient primitive way to obtain (P2) is to have a sufficiently large overhead parameter χ, *or sufficiently low delegated productivity in the high-collateral region due to the fixed discount factor* $1 - \rho_D$.

At $C = \underline{C}$, a small increase in collateral raises $\mu(C) = a(R + \underline{\phi}C)$, lifting investment incentives for all interior solo and delegated types. Near $\bar{w}$, solo eligibility collapses as $(\bar{w} - C)/\bar{w} \to 0$, and the quadratic overhead cost continues to weigh on welfare. The interior optimum balances the incentive gain against the participation and overhead effects.

4.5 Welfare in the baseline model

The protocol objective $\mathcal{V}(C) = \mathcal{L}(C) - \chi N(C)^2/2$ captures welfare in the baseline model. A first structural observation is that delegated validation is exposed to a common shock with probability $\rho_D \in (0, 1)$, while solo validation is not, so the delegated contribution to effective liveness is discounted by the factor $1 - \rho_D$ relative to the solo contribution. For a given delegated success probability and delegated mass, this means that reallocation of validation activity toward delegation reduces effective liveness through the common-shock channel, even holding delegated investment fixed. The details are recorded in Appendix A.5.

Turning to welfare, write $\mathcal{L}(C) = \mathcal{L}(\mu(C), s^*(C), N_D(C))$ where the three state variables capture the direct incentive level, the solo–delegated composition, and the delegated pool scale respectively. Under the uniform benchmark, $N_S(C) = \frac{\bar{w}-C}{\bar{w}} \cdot \frac{s^*(C)-\underline{s}}{\bar{s}-\underline{s}}$ is pinned down by $(C, s^*(C))$, and $J^\star(C) = J^\star(C, N_D(C))$ is pinned down by $(C, N_D(C))$, so this representation is complete. Under a continuous equilibrium selection the derivative of protocol welfare decomposes as

$$\mathcal{V}'(C) = \underbrace{\frac{\partial \mathcal{L}}{\partial \mu}\mu'(C)}_{\text{direct incentive effect } >0} + \underbrace{\frac{\partial \mathcal{L}}{\partial s^*}s^{*\prime}(C)}_{\text{cutoff effect}} + \underbrace{\frac{\partial \mathcal{L}}{\partial N_D}N_D'(C)}_{\text{delegated-scale effect}} - \chi N(C)N'(C), \tag{20}$$

where $\mu(C) = a(R + \underline{\phi}C)$ and $\mu'(C) = a\underline{\phi} > 0$. The direct incentive effect is strictly positive because higher collateral raises $\mu(C)$, lifting both $I_S^\star$ and $J^\star$ on the interior region. The cutoff effect captures how s^* responds to changing collateral, shifting agents between solo and delegated operation. The delegated-scale effect operates through how N_D responds, affecting pool investment through the scale economy while also altering exposure to ρ_D. The overhead term $-\chi N(C)N'(C)$ is negative when the validator count rises. The signs of $s^{*\prime}(C)$, $N_D'(C)$, and $N'(C)$ are in general endogenous and ambiguous, so the net sign of $\mathcal{V}'(C)$ is ambiguous. Details of the decomposition are relegated to Appendix A.6.

A related observation concerns the planner's problem of choosing the skill cutoff. Fix $C \in [\underline{C}, \bar{w})$ and let $\mathcal{V}(C, s) \equiv \mathcal{L}(C, s) - \chi N(C, s)^2/2$ denote protocol welfare when the cutoff is s, holding C fixed. Define

$$\Xi(C, s) \equiv \frac{\bar{w} - C}{\bar{w}(\bar{s} - \underline{s})} > 0 \tag{21}$$

as the density of wealth-eligible agents at cutoff s, and define $\Lambda(C, s)$ residually by

$$\Lambda(C, s) \equiv \frac{\partial \mathcal{V}(C, s)}{\partial s} - \Xi(C, s)\big(U_S^\star(C, s) - U_D^\star(C, N_D(C, s))\big)\,. \tag{22}$$

By construction,

$$\frac{\partial \mathcal{V}(C, s)}{\partial s} = \Xi(C, s)\big(U_S^\star(C, s) - U_D^\star(C, N_D(C, s))\big) + \Lambda(C, s). \tag{23}$$

The first term is proportional to the private payoff difference at the cutoff; $\Lambda(C, s)$ collects all protocol-level composition effects not captured by private payoffs. At the equilibrium cutoff s_{eq}, private indifference gives $U_S^\star - U_D^\star = 0$, so $\partial \mathcal{V}/\partial s|_{s=s_{\text{eq}}} = \Lambda(C, s_{\text{eq}})$. When $\Lambda(C, s_{\text{eq}}) > 0$, a marginal increase in the skill cutoff — which shifts agents from delegation toward solo — raises protocol welfare, implying the equilibrium cutoff is too low and the delegated sector is too large relative to what a planner would choose. Private agents compare only their own payoffs, while the planner also values the protocol-level composition effects embedded in Λ. Details are in Appendix A.7.

Collateral is a blunt policy instrument: it affects investment incentives, participation, and composition simultaneously. This motivates the introduction of anti-correlation penalties as a second tool aimed more directly at the pricing of correlated failure risk.

5 Anti-correlation penalties

5.1 Penalty structure

The baseline model established that delegation benefits from scale economies in monitoring but is exposed to correlated failure risk through the reduced-form common-shock probability $\rho_D \in (0, 1)$. Anti-correlation penalties do not create this exposure; it is already present in the baseline. Instead,

the penalties price correlated failure more sharply by making slashing depend on the expected size of the concurrent failure set, creating incentives for validators to separate operationally.[9]

We model systemic risk through the *expected concurrent failure share* α_i, the expected fraction of all validators that fail simultaneously when validator i fails. Solo validators fail independently, so

$$\alpha_S(C,\kappa) \equiv \frac{1}{N(C,\kappa)}, \tag{24}$$

where the dependence on (C,κ) is induced through the equilibrium object $N(C,\kappa)$. Delegated validators fail together with probability ρ_D and independently otherwise, giving

$$\alpha_D(C,\kappa) \equiv \rho_D \frac{N_D(C,\kappa)}{N(C,\kappa)} + (1-\rho_D)\frac{1}{N(C,\kappa)}, \tag{25}$$

where again the dependence on (C,κ) operates through the equilibrium objects $N_D(C,\kappa)$ and $N(C,\kappa)$. We have $\alpha_D \geq \alpha_S$, with a strict gap whenever delegation is a nontrivial share of the validator set. Two objects play distinct roles: ρ_D is the probability that a systemic failure event occurs, while α_D is the expected share of the network that fails simultaneously conditional on a given delegated validator failing. Penalties are indexed to α_D because they are designed to respond to the realized size of the concurrent failure set, not merely to whether a systemic event occurred.

To keep the concentration-sensitive penalty schedule transparent, we assume it is linear in the concurrent failure share,

$$\Gamma(\alpha) = \alpha, \tag{26}$$

so that the marginal slashing cost is proportional to the share of the network that fails simultaneously. Let $\kappa \in [0,1]$ control penalty strength. When a validator fails and the concurrent failure share is α, the slashing fraction is

$$\phi(\alpha,\kappa) \equiv \underline{\phi}(1+\kappa\,\alpha). \tag{27}$$

Applied to the two validator types:

$$\phi_S(C,\kappa) = \underline{\phi}\left(1+\frac{\kappa}{N(C,\kappa)}\right), \qquad \phi_D(C,\kappa) = \underline{\phi}\Big(1+\kappa\,\alpha_D(C,\kappa)\Big). \tag{28}$$

[9] See Vitalik Buterin's discussion of anti-correlation incentives: https://ethresear.ch/t/supporting-decentralized-staking-through-more-anti-correlation-incentives/19116.

In large networks, $\alpha_S = 1/N \approx 0$, so $\phi_S \approx \underline{\phi}$: penalties barely affect solo validators. By contrast, ϕ_D is sensitive to κ and the delegation share through α_D, so penalties fall disproportionately on pooled operators.

5.2 Investment under penalties

Under anti-correlation penalties, a solo validator with skill cost s faces

$$U_S(I; C, \kappa, s) = p(I)(R + \phi_S(C, \kappa)C) - \phi_S(C, \kappa)C - sI, \tag{29}$$

and the delegated payoff per position is

$$U_D(J; C, \kappa, N_D) = (1 - \rho_D)\Big(p(J)(R + \phi_S(C, \kappa)C) - \phi_S(C, \kappa)C\Big) - \rho_D \phi_D(C, \kappa)C - \frac{J}{N_D}. \tag{30}$$

Note that ϕ_D enters (30) only through the level term $\rho_D \phi_D C$, which is independent of J: the delegated FOC depends on ϕ_S but not ϕ_D. The payoff specification uses ϕ_S in the non-systemic delegated states because failures in those states are idiosyncratic (each delegated validator fails independently with probability $1 - p(J)$), so the concurrent failure share is the same as for solo validators. The higher concentration-sensitive penalty ϕ_D applies only in the systemic state, where the entire pool fails jointly. Define the penalized composites

$$\mu_\kappa(C) \equiv a(R + \phi_S(C, \kappa)C), \qquad \nu_\kappa(C, N_D) \equiv a(1 - \rho_D)N_D(R + \phi_S(C, \kappa)C). \tag{31}$$

Proposition 5 (Solo investment under anti-correlation penalties)**.** *Under* (4)*, for a given penalty environment $\phi_S(C, \kappa)$ (treated as fixed in the agent's best-response problem), the solo problem has a unique optimum $I_S^\star(C, \kappa, s) \geq 0$.*

1. *If $\mu_\kappa(C) > s$, the optimum is interior:*

$$I_S^\star(C, \kappa, s) = \frac{1}{a} \log\left(\frac{\mu_\kappa(C)}{s}\right), \tag{32}$$

with success probability $p(I_S^\star) = 1 - s/\mu_\kappa(C)$ *and maximized payoff*

$$U_S^\star(C, \kappa, s) = R - \frac{s}{a}\left(1 + \log\left(\frac{\mu_\kappa(C)}{s}\right)\right). \tag{33}$$

2. *If* $\mu_\kappa(C) \leq s$*, then* $I_S^\star = 0$ *and* $U_S^\star(C, \kappa, s) = -\phi_S(C, \kappa)C$*.*

3. *On the interior region,* $I_S^\star$ *is strictly increasing in* ϕ_S *(and hence weakly increasing in* κ *whenever* ϕ_S *is weakly increasing in* κ *along the equilibrium branch), strictly increasing in* C*, and strictly decreasing in* s*.*

Proof is in Appendix A.8.

The proposition treats ϕ_S as given in the agent's problem, since ϕ_S depends on aggregate composition rather than any individual's choice. Under fixed composition, $\phi_S(C, \kappa) = \underline{\phi}(1 + \kappa/N)$ is weakly increasing in κ, so $I_S^\star$ is weakly increasing in κ on the interior region. In large networks, $1/N \approx 0$, so the effect is small.

Proposition 6 (Delegated investment under anti-correlation penalties)**.** *Under* (4)*, for a given* $\phi_S(C, \kappa)$ *(treated as fixed), the monitor's problem has a unique optimum* $J^\star(C, \kappa, N_D) \geq 0$*.*

1. *If* $\nu_\kappa(C, N_D) > 1$*, the optimum is interior:*

$$J^\star(C, \kappa, N_D) = \frac{1}{a}\log\big(\nu_\kappa(C, N_D)\big), \tag{34}$$

with success probability $p(J^\star) = 1 - 1/\nu_\kappa$ *and maximized payoff*

$$U_D^\star(C, \kappa, N_D) = (1 - \rho_D)R - \rho_D\phi_D(C, \kappa)C - \frac{1}{aN_D}\big(1 + \log\big(\nu_\kappa(C, N_D)\big)\big). \tag{35}$$

2. *If* $\nu_\kappa(C, N_D) \leq 1$*, then* $J^\star = 0$ *and*

$$U_D^\star(C, \kappa, N_D) = -(1 - \rho_D)\phi_S(C, \kappa)C - \rho_D\phi_D(C, \kappa)C.$$

3. *On the interior region,* $J^\star$ *is strictly increasing in* ϕ_S *(and hence weakly increasing in* κ *under fixed composition), strictly increasing in* N_D*, and strictly increasing in* C *(holding* ϕ_S *fixed).*

4. *$\phi_D(C, \kappa)$ does not appear in the first-order condition; it enters the payoff only through the level term $-\rho_D \phi_D C$.*

Proof is in Appendix A.9.

Higher ϕ_S, which rises with κ under fixed composition, raises μ_κ and hence ν_κ, improving both investment and the success probability. The effect of κ through ϕ_D is exclusively a level pricing of correlated failure risk in delegated payoffs.

5.3 Composition and protocol choice with penalties

With κ as an additional instrument, the protocol objective becomes

$$\mathcal{V}(C, \kappa) \equiv \mathcal{L}(C, \kappa) - \frac{\chi}{2} N(C, \kappa)^2, \tag{36}$$

where

$$\mathcal{L}(C, \kappa) = N_S(C, \kappa)\, \bar{p}_S(C, \kappa) + (1 - \rho_D)\, N_D(C, \kappa)\, p(J^\star(C, \kappa)).$$

The equilibrium skill cutoff $s^*(C, \kappa, N_D)$ solves $U_S^\star(C, \kappa, s^*) = U_D^\star(C, \kappa, N_D)$, with κ entering both sides. The joint policy domain is $C \in [\underline{C}, \bar{w}]$, $\kappa \in [0, 1]$.

Proposition 7 (Collateral and penalty strength are substitutes)**.** *Under the parametric benchmark, holding composition (N_S, N_D, s^*) fixed so that N is constant and $\phi_S(C, \kappa) = \underline{\phi}(1 + \kappa/N)$, suppose*

$$\mu_\kappa(C) > 2a\underline{\phi}\, C\left(1 + \frac{\kappa}{N}\right) \tag{37}$$

at the relevant interior optimum. Then

$$\frac{\partial^2 \mathcal{V}}{\partial C\, \partial \kappa} > 0.$$

Hence, if additionally $\partial^2 \mathcal{V}/\partial C^2 < 0$ at a locally unique interior optimum, then

$$\frac{dC^\star}{d\kappa} = -\frac{\partial^2 \mathcal{V}/\partial C\, \partial \kappa}{\partial^2 \mathcal{V}/\partial C^2} < 0. \tag{38}$$

Proof is in Appendix A.11.

Under fixed composition, liveness from both sectors depends on (C, κ) only through $\mu_\kappa(C)$. When condition (37) holds, the mixed partial $\partial^2 \mathcal{L}/\partial C\, \partial \kappa > 0$: C and κ are *complements in the objective* in the sense that raising κ increases the marginal return to collateral. This complementarity in the objective, combined with strict concavity in C, produces *policy substitutability* through the implicit function theorem (38): the optimal collateral $C^\star$ falls when κ rises. Condition (37) is equivalent to $R > \underline{\phi} C(1 + \kappa/N)$; in large networks where κ/N is small, this is close to the requirement $R > \underline{\phi} C$, which holds under standard parameters. Higher penalty strength thus allows the protocol to achieve a given liveness target with a lower collateral requirement.

Proposition 8 (Penalties and delegated attractiveness)**.** *Define the equilibrium delegated payoff*

$$\pi_D(C, \kappa) \equiv U_D^\star(C, \kappa, N_D(C, \kappa)). \tag{39}$$

Holding composition fixed,

$$\frac{\partial \pi_D}{\partial \kappa} = - \underbrace{(1 - \rho_D) \frac{C \underline{\phi}}{N \nu_\kappa}}_{\text{common-slashing channel}} \underbrace{- \rho_D \underline{\phi}\, \alpha_D(C, \kappa)\, C}_{\text{systemic channel}} < 0. \tag{40}$$

Thus stronger anti-correlation penalties reduce delegated attractiveness under fixed composition.

Proof is in Appendix A.12.

Higher κ affects delegated attractiveness through two distinct penalty channels. First, it raises the common slashing exposure ϕ_S, which enters the delegated problem even in non-systemic states and lowers the optimized delegated payoff by the envelope theorem. Second, it raises the concentration-specific slashing term ϕ_D, which lowers delegated payoff directly through the systemic-failure state. Both effects reduce delegated attractiveness under fixed composition.

5.4 Extension: endogenizing penalty strength

We now consider the joint design problem

$$\max_{C \in [\underline{C}, \bar{w}],\, \kappa \in [0,1]} \quad \mathcal{V}(C, \kappa). \tag{41}$$

To isolate the incentive channel, we first analyze a fixed-composition benchmark in which the skill cutoff s^* and the induced validator sets are held constant as κ varies.

Proposition 9 (Weakly improving κ under fixed composition)**.** *Under fixed composition (N_S, N_D independent of κ), the effective liveness $\mathcal{L}(C, \kappa)$ is weakly increasing in κ for all $C \in [\underline{C}, \bar{w}]$. If at least one sector with positive mass is on the interior investment region, then $\partial \mathcal{L}/\partial \kappa > 0$ and the fixed-composition optimal penalty strength is $\kappa^\star = 1$.*

Proof is in Appendix A.13.

Under fixed composition, N is constant, so $\phi_S(C, \kappa) = \underline{\phi}(1 + \kappa/N)$ is strictly increasing in κ, which raises both $I_S^\star$ and $J^\star$ on the interior region whenever $C \geq \underline{C} > 0$ and the relevant optima are interior.

Proposition 10 (Interior optimal penalty strength)**.** *Fix $C \in (\underline{C}, \bar{w})$. Suppose $\mathcal{V}(C, \kappa)$ is continuous on $[0, 1]$ and differentiable on $(0, 1)$. If*

$$\left.\frac{\partial \mathcal{V}(C, \kappa)}{\partial \kappa}\right|_{\kappa=0^+} > 0 \qquad \textit{and} \qquad \left.\frac{\partial \mathcal{V}(C, \kappa)}{\partial \kappa}\right|_{\kappa=1} < 0, \tag{42}$$

then there exists an interior optimizer $\kappa^\star(C) \in (0, 1)$.

Proof is in Appendix A.14.

The endpoint conditions (42) can be verified from the preceding propositions. At $\kappa = 0^+$, the fixed-composition incentive gain is strictly positive and dominates composition responses that are second-order at $\kappa = 0$, so the first inequality holds. At $\kappa = 1$, by Proposition 8, both penalty channels reduce delegated attractiveness. If the resulting contraction of the delegated sector and scale economy is large enough to outweigh the remaining fixed-composition incentive gain, the second inequality holds. An interior $\kappa^\star$ then balances the incentive-improvement and composition-shift effects.

6 Conclusion

Proof of stake protocols rely on staked collateral to discipline validators, but the staking sector has become increasingly intermediated through professional operators and liquid staking arrangements.

This paper provides a simple framework in which collateral policy affects not only individual incentives to invest in operational reliability, but also the equilibrium composition of validators. Delegation can raise average performance through scale and monitoring, yet it can also concentrate operational dependencies and increase exposure to common shocks. As a result, tighter collateral requirements may strengthen discipline while simultaneously increasing systemic fragility through correlated failures, implying an interior optimal minimum stake.

We provide evidence consistent with this mechanism using the September 2025 Kiln infrastructure incident in Ethereum and daily validator performance from Rated Network. Professional operators experienced sharp, economically meaningful declines in effectiveness and revenues during the incident, while solo validators exhibited little to no abnormal disruption. The evidence supports a delegated monitoring channel in which intermediation improves routine efficiency but amplifies tail risk.

These findings have implications for protocol design and for the broader economics of intermediated infrastructure. Policies that focus narrowly on average validator performance or operational efficiency may overlook composition effects that shape systemic risk. More generally, proof of stake provides a new setting in which collateral requirements, intermediation, and correlated default risk interact, and in which familiar corporate finance and banking trade offs can be studied with high frequency operational data.

References

Allen, F. and Gale, D. (2000). Financial contagion. *Journal of Political Economy*, 108(1):1–33.

Carré, S. and Gabriel, F. (2024). Liquid staking: When does it help? deploying liquidity on well-designed defi lending platforms. *Working Paper*.

Choi, D. B., Goldsmith-Pinkham, P., and Yorulmazer, T. (2024). Contagion effects of the silicon valley bank run. NBER Working Paper 31772.

Cong, L. W., He, Z., and Tang, K. (2025). The tokenomics of staking. *Working Paper*.

Diamond, D. W. (1984). Financial intermediation and delegated monitoring. *The Review of Economic Studies*, 51(3):393–414.

Fleckinger, P., Martimort, D., and Roux, N. (2024). Should they compete or should they cooperate? the view of agency theory. *Journal of Economic Literature*. Forthcoming.

Goldsmith-Pinkham, P. and Yorulmazer, T. (2010). Liquidity, bank runs, and bailouts: Spillover effects during the northern rock episode. *Journal of Financial Services Research*, 37(2-3):83–98.

Kvaløy, O. and Olsen, T. E. (2019). Relational contracts, multiple agents, and correlated outputs. *Management Science*, 65(11):5360–5370.

Lehar, A. and Parlour, C. A. (2022). Systemic fragility in decentralized markets. *BIS Working Papers*, (1062).

Mu, B., Prat, J., and Tovanich, N. (2025). Incentives and reputation in liquid restaking protocols. *CBER CtCe Working Paper*.

Tirole, J. (2006). *The Theory of Corporate Finance*. Princeton University Press.

Tzinas, A. and Zindros, D. (2023). The principal-agent problem in liquid staking. *Working Paper*.

A Appendix: Tables and Figures

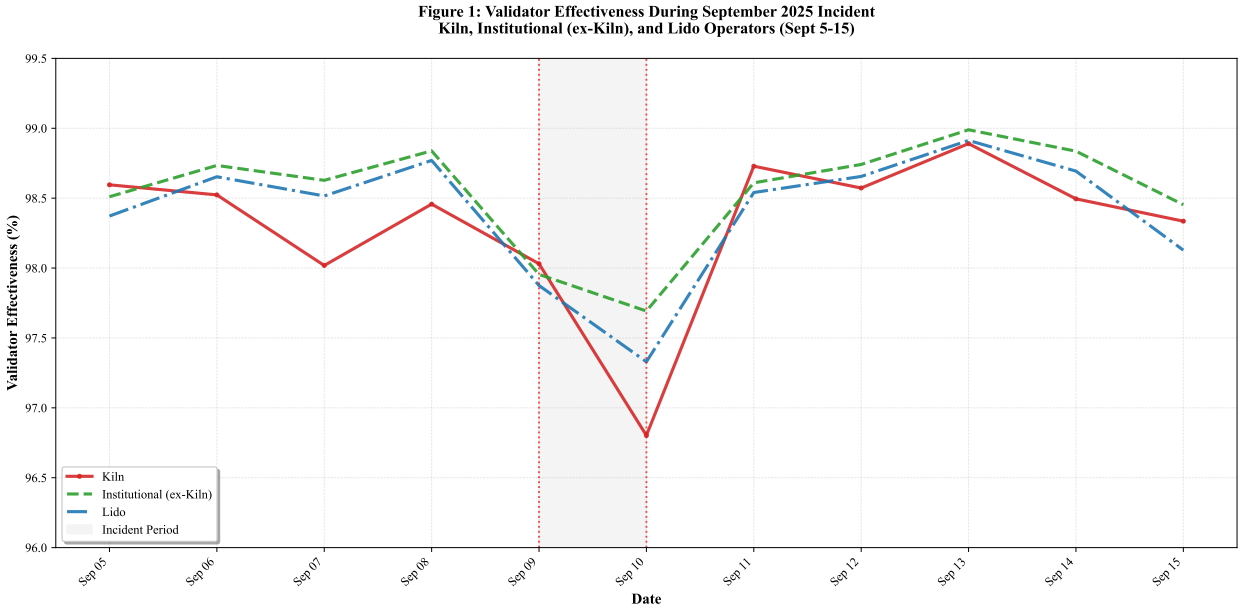

Figure 1: Validator Effectiveness During September 2025 Incident. This figure plots daily validator effectiveness for three operator categories: Kiln (solid red), Institutional operators excluding Kiln (dashed green), and Lido sub-operators (dash-dot blue). The shaded area indicates the September 9-10 incident period. Effectiveness is weighted by validator count within each category.

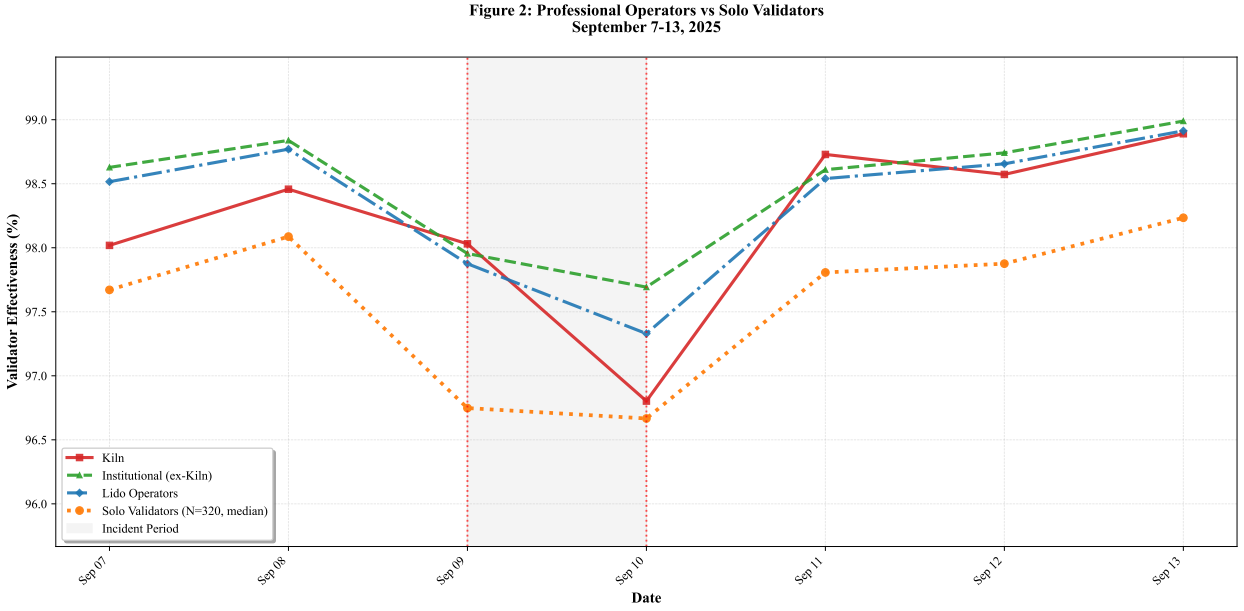

Figure 2: Institutional validator Effectiveness During September 2025 Incident compared to solo stakers

Table 1: Sample Composition

Category	N Entities	N Validators	Time Period (2025)	Total Obs.
Panel A: Professional Operators				
Institutional Operators	5	154,451	Aug 1 - Sep 30	300
Kiln	1	57,344	(60 days)	60
Other Institutional	4	97,107		240
Lido Sub-Operators (P2P, Chorus One, Attestant, others)	10	97,340	Aug 1 - Sep 30 (60 days)	600
Centralized Exchange (Coinbase Cloud)	1	10,476	Aug 1 - Sep 30 (60 days)	60
Professional Total	**16**	**262,267**	**60 days**	**960**
Panel B: Solo Validators				
True Solo (1 validator)	250	250	Sep 5 - Sep 15	2,750
Small Solo (2-5 validators)	150	450	(11 days)	1,650
Medium Solo (6-32 validators)	80	1,200		880
Large Solo (33-100 validators)	42	2,600		462
Solo Total	**322**	**4,500**	**11 days**	**3,542**
Panel C: Liquid Staking Protocol Aggregates				
Lido	1	268,920	Aug 1 - Sep 30	60
Stader	1	28,450	(60 days)	60
Frax	1	12,340		60
Swell	1	8,120		60
Protocol Total	**4**	**317,830**	**60 days**	**240**

Notes: This table presents the composition of our sample. Panel A shows professional node operators aggregated into three categories: institutional operators (including Kiln separately), Lido sub-operators, and centralized exchange validators. Professional operators are observed for 60 days from August 1 to September 30, 2025, yielding 960 operator-day observations (16 entities $\times$ 60 days). Panel B shows solo validators sampled across four size categories based on number of validators controlled by each deposit address. Solo validators are observed for 11 days from September 5-15, 2025, yielding 3,542 validator-day observations (322 validators $\times$ 11 days). Panel C shows liquid staking protocol-level aggregates for comparison with sub-operator performance. Validator counts reflect active validators as of September 8, 2025. Total observations equal entities multiplied by days, with all observations included regardless of incident status.

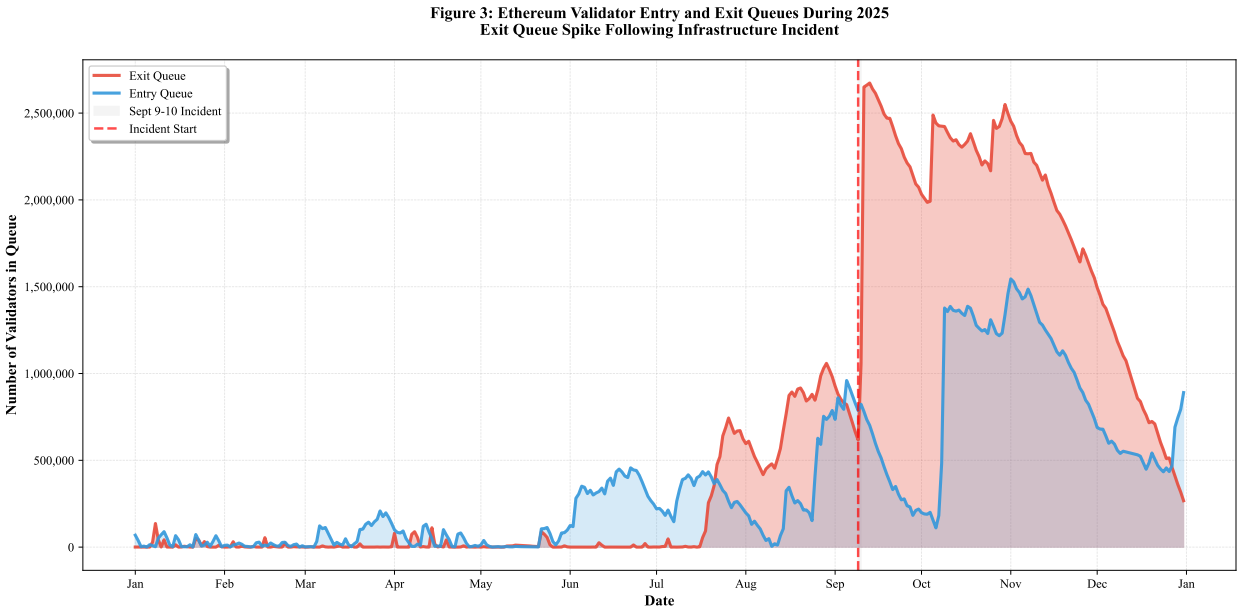

Figure 3: Validator Exit Spike During the September 2025 Incident

Table 2: Performance Impact Among Professional Operators

Category	Baseline Mean (%)	SD (pp)	Incident Loss (pp)	Standardized Loss (z-score)	N Validators
Kiln	98.46	0.52	1.66***	3.19	57,344
			(0.18)		
Institutional (ex-Kiln)	98.58	0.48	0.91*	1.90	96,107
			(0.16)		
Lido	98.22	0.44	1.06**	2.41	97,340
			(0.15)		
CEX	98.12	0.61	0.65	1.07	10,476
			(0.21)		
Differential Impact (Relative to Kiln)					
Institutional - Kiln			−0.75**	−1.44	
			(0.24)		
Lido - Kiln			−0.60*	−1.15	
			(0.23)		
CEX - Kiln			−1.01**	−1.66	
			(0.28)		

Notes: This table reports effectiveness drops during September 9-10, 2025 for professional operators. Baseline statistics computed over 58 non-incident days (Aug 1 - Sep 30, excluding Sep 9-10). Losses calculated as validator-weighted category means. Standard errors in parentheses account for weighting by validator count. Standardized losses (z-scores) computed as (incident loss) / (baseline SD). Differential impacts report category differences using pooled standard errors. Significance based on normal approximation: *** —z— > 3.0 ($p < 0.003$), ** —z— > 2.5 ($p < 0.013$), * —z— > 2.0 ($p < 0.046$).

Table 3: Economic Impact Among Professional Operators

	Loss per Validator (Gwei)			
Category	Priority Fees	MEV Rewards	Total Rewards	Z-score
Kiln	52.3***	38.6***	90.9***	6.35
	(6.0)	(5.2)	(9.9)	
Institutional (ex-Kiln)	35.6***	26.1***	61.7***	4.26
	(5.8)	(4.9)	(9.5)	
Lido	41.8***	30.3***	72.1***	4.87
	(5.7)	(4.8)	(9.3)	
CEX	33.2***	24.5**	57.7***	3.87
	(7.8)	(6.6)	(12.7)	
Differential Impact (Relative to Kiln)				
Institutional - Kiln	−16.7**	−12.5*	−29.2**	−2.04
	(6.5)	(5.5)	(10.6)	
Lido - Kiln	−10.5·	−8.3	−18.8*	−1.31
	(6.3)	(5.4)	(10.4)	
CEX - Kiln	−19.1**	−14.1*	−33.2**	−2.32
	(7.5)	(6.3)	(12.2)	
Baseline SD (Gwei)	14.3	12.1	14.3	

Notes: This table reports economic losses per validator during September 9-10, 2025. Losses computed as (baseline mean - incident period mean) for each revenue category, normalized by validator count. Baseline calculated over 58 non-incident days. Standard errors in parentheses reflect validator-weighted category-level variation. Z-scores for total rewards computed using baseline SD of 14.3 Gwei. Significance levels: ***p<0.001, **p<0.01, *p<0.05, ·p<0.10.

Table 4: Temporal Dynamics Among Professional Operators

	Effectiveness Loss (pp)		Total Rewards Loss (Gwei)	
Category	Sept 8→9	Sept 9→10	Sept 8→9	Sept 9→10
Kiln	0.43*	1.23**	22.5***	68.4***
	(0.18)	(0.18)	(5.4)	(8.8)
Institutional	0.59***	0.32	23.6***	18.1*
	(0.16)	(0.16)	(5.2)	(8.5)
Lido	0.45**	0.61***	21.8***	30.3***
	(0.15)	(0.15)	(5.1)	(8.3)
CEX	2.28***	−1.63**	42.0***	−14.3
	(0.21)	(0.21)	(7.0)	(11.5)
Day 2 Differential (relative to Kiln)				
Institutional - Kiln		−0.91***		−50.3***
		(0.24)		(11.2)
Lido - Kiln		−0.62**		−38.1***
		(0.23)		(10.9)
CEX - Kiln		−2.86***		−82.7***
		(0.28)		(14.3)

Notes: This table decomposes incident impacts by day. September 8 used as pre-incident baseline. Day 1 losses (Sept 8→9) capture initial impact; Day 2 losses (Sept 9→10) capture evolution. Standard errors in parentheses. All operators experienced initial losses on September 9. However, while other operators stabilized or recovered on September 10, Kiln's losses accelerated substantially. Significance levels: ***p<0.001, **p<0.01, *p<0.05.

Table 5: Effectiveness Impact: Solo Validators vs. Professional Operators

Category	N Validators	Baseline Days	Baseline SD (pp)	Incident Loss (pp)	Standardized Loss (z)	Interpretation
Solo	322	9	0.45	0.03 (0.45)	0.07	Normal variation
Kiln	57,344	58	0.52	1.66*** (0.18)	3.19	Highly anomalous
Institutional	96,107	58	0.48	0.91* (0.16)	1.90	Significant
Lido	97,340	58	0.44	1.06** (0.15)	2.41	Significant
CEX	10,476	58	0.61	0.65 (0.21)	1.07	Marginally significant
Effect Sizes (Cohen's d, pooled SD)						
Kiln vs. Solo				1.63pp	3.50***	Very large
Institutional vs. Solo				0.88pp	1.89*	Large
Lido vs. Solo				1.03pp	2.21**	Large
CEX vs. Solo				0.62pp	1.33	Medium

Notes: This table compares effectiveness losses between solo validators and professional operators. Solo baseline computed over 9 non-incident days (Sep 5-15 excluding Sep 9-10); professional baselines over 58 days (Aug 1 - Sep 30 excluding Sep 9-10). Standard errors for solo validators (in parentheses) computed from 644 individual validator-day observations with clustering by validator. Professional operators show category-level weighted means. Standardized losses (z-scores) computed as (incident loss) / (baseline SD) for each group independently. Effect sizes (Cohen's d) use pooled standard deviations. Significance: *** —z— > 3.0 ($p < 0.003$), ** —z— > 2.5 ($p < 0.013$), * —z— > 2.0 ($p < 0.046$).

Table 6: Economic Impact: Solo Validators vs. Professional Operators

Category	Loss per Validator (Gwei)			Baseline	Standardized	
	Priority Fees	MEV Rewards	Total Rewards	SD (Gwei)	Loss (z-score)	Interpretation
Solo	0.8 (1.0)	0.5 (0.8)	1.3 (1.3)	18.2 (9 days)	0.07	Normal variation
Kiln	52.3	38.6	90.9	14.3 (58 days)	6.35	Highly anomalous
Institutional	35.6	26.1	61.7	14.5 (58 days)	4.26	Highly anomalous
Lido	41.8	30.3	72.1	14.8 (58 days)	4.87	Highly anomalous
CEX	33.2	24.5	57.7	14.9 (58 days)	3.87	Highly anomalous
Effect Sizes (Professional - Solo, pooled SD)						
Kiln - Solo	51.5	38.1	89.6	16.3	5.50***	Very large
Institutional - Solo	34.8	25.6	60.4	16.4	3.68***	Very large
Lido - Solo	41.0	29.8	70.8	16.5	4.29***	Very large
CEX - Solo	32.4	24.0	56.4	16.6	3.40***	Very large

Notes: This table compares economic losses between solo validators and professional operators. Solo baseline computed over 9 non-incident days; professional baselines over 58 days. Standard errors for solo validators (in parentheses) from 644 individual observations. Professional operators show weighted category means. Baseline SDs calculated separately for each group over their respective non-incident periods. Standardized losses (z-scores) measure deviation from each group's own baseline. Effect sizes use pooled SDs accounting for different baseline periods. All professional categories exhibit losses exceeding 3.87 standard deviations while solo validators remained within normal variation. Significance: *** —z— > 3.0 ($p < 0.003$).

Table 7: Temporal Dynamics: Solo vs. Professional Operators

Category	Effectiveness Loss (pp)		Total Rewards Loss (Gwei)	
	Sept 8→9 (z-score)	Sept 9→10 (z-score)	Sept 8→9 (z-score)	Sept 9→10 (z-score)
Solo	0.01	0.02	0.5	0.8
	(0.02)	(0.04)	(0.03)	(0.04)
Kiln	0.43	1.23	22.5	68.4
	(0.83)	(2.37)**	(1.57)	(4.78)***
Institutional	0.59	0.32	23.6	18.1
	(1.23)	(0.67)	(1.63)	(1.25)
Lido	0.45	0.61	21.8	30.3
	(1.02)	(1.39)	(1.47)	(2.05)**
CEX	2.28	−1.63	42.0	−14.3
	(3.74)***	(−2.67)**	(2.82)**	(−0.96)
Day 2 Differential (Professional - Solo, z-scores)				
Kiln - Solo		2.33**		4.74***
Institutional - Solo		0.63		1.21
Lido - Solo		1.35		2.01**
CEX - Solo		−2.71**		−1.00

Notes: This table decomposes incident impacts by day. Z-scores computed using each category's baseline SD. Solo validators show consistent stability across both days ($z < 0.1$). Professional operators exhibited varying patterns: initial disruption on September 9 affected all categories, but recovery diverged substantially on September 10. Kiln's losses accelerated ($z = 2.37$ effectiveness, 4.78 rewards on day 2) while other operators stabilized or recovered. CEX validators recovered on day 2, yielding negative differentials vs. solo. Significance: *** —z— > 3.0, ** —z— > 2.5, * —z— > 2.0.

A Appendix: Proofs

A.1 Proof of Proposition 1

Proof. Strict concavity and unique optimum. Under $p(x) = 1 - e^{-ax}$, the objective is $U_S(I; C, s) = (1 - e^{-aI})(R + \underline{\phi}C) - \underline{\phi}C - sI$. The second derivative in I is $-a^2 e^{-aI}(R + \underline{\phi}C) < 0$, so U_S is strictly concave and the optimum is unique.

Interior characterization and value function. The first-order condition is $ae^{-aI}(R + \underline{\phi}C) = s$, giving $e^{-aI_S^\star} = s/\mu(C)$ and $I_S^\star = \log(\mu(C)/s)/a > 0$ iff $\mu(C) > s$. Substituting:

$$\begin{aligned} U_S^\star(C, s) &= (1 - s/\mu(C))(R + \underline{\phi}C) - \underline{\phi}C - \frac{s}{a}\log\left(\frac{\mu(C)}{s}\right) \\ &= R - \frac{s}{a} - \frac{s}{a}\log\left(\frac{\mu(C)}{s}\right) = R - \frac{s}{a}\left(1 + \log\left(\frac{\mu(C)}{s}\right)\right), \end{aligned}$$

using $(R + \underline{\phi}C)s/\mu(C) = s/a$. When $\mu(C) \le s$, the derivative at $I = 0$ is non-positive, so $I_S^\star = 0$ and $U_S^\star = -\underline{\phi}C$.

Monotonicity. $dI_S^\star/dC = (1/a)\mu'(C)/\mu(C) = \underline{\phi}/\mu(C) > 0$. $dI_S^\star/ds = -1/(as) < 0$.

$$\frac{\partial U_S^\star}{\partial C} = \frac{s\,\underline{\phi}}{\mu(C)} > 0,$$

obtained by differentiating $R - (s/a)(1 + \log\mu(C) - \log s)$ with respect to C: $(s/a)\cdot\mu'(C)/\mu(C) = s\underline{\phi}/\mu(C) > 0$. $\partial U_S^\star/\partial s = -I_S^\star < 0$. □

A.2 Proof of Proposition 2

Proof. Strict concavity and unique optimum. The objective is $U_D = (1 - \rho_D)(1 - e^{-aJ})(R + \underline{\phi}C) - (1 - \rho_D)\underline{\phi}C - \rho_D\underline{\phi}C - J/N_D$. The second derivative in J is $-(1 - \rho_D)a^2 e^{-aJ}(R + \underline{\phi}C) < 0$, so U_D is strictly concave with a unique optimum.

Interior characterization and value function. The first-order condition is $a(1 - \rho_D)e^{-aJ}(R + \underline{\phi}C) = 1/N_D$, giving $e^{-aJ^\star} = 1/\nu$ and $J^\star = \log(\nu)/a > 0$ iff $\nu > 1$, where $\nu = \nu(C, N_D) =$

$a(1-\rho_D)N_D(R+\underline{\phi}C)$. Substituting:

$$U_D^\star = (1-\rho_D)(1-1/\nu)(R+\underline{\phi}C) - (1-\rho_D)\underline{\phi}C - \rho_D\underline{\phi}C - \frac{\log\nu}{aN_D}$$
$$= (1-\rho_D)(R+\underline{\phi}C) - \frac{(1-\rho_D)(R+\underline{\phi}C)}{\nu} - \underline{\phi}C - \frac{\log\nu}{aN_D}.$$

Since $(1-\rho_D)(R+\underline{\phi}C)/\nu = 1/(aN_D)$:

$$U_D^\star = (1-\rho_D)R - \rho_D\underline{\phi}C - \frac{1}{aN_D}(1+\log\nu),$$

which is (10). When $\nu \le 1$, $J^\star = 0$ and $U_D^\star(0;C,N_D) = (1-\rho_D)(0-\underline{\phi}C) - \rho_D\underline{\phi}C = -(1-\rho_D)\underline{\phi}C - \rho_D\underline{\phi}C = -\underline{\phi}C$.

Monotonicity. $\partial J^\star/\partial N_D = (1/a)(1/N_D) > 0$. $\partial J^\star/\partial C = (1/a)\underline{\phi}/(R+\underline{\phi}C) > 0$.

Part (iv). $p(J^\star) = 1-1/\nu$ and $p(I_S^\star(C,\underline{s})) = 1-\underline{s}/\mu(C)$. Using $\nu = (1-\rho_D)N_D\mu(C)$:

$$p(J^\star) > p(I_S^\star(C,\underline{s})) \iff \frac{1}{\nu} < \frac{\underline{s}}{\mu(C)} \iff (1-\rho_D)N_D > \frac{1}{\underline{s}}. \qquad \square$$

A.3 Fixed-point existence (Proof of Proposition 3)

Proof. Fix $C \in [\underline{C},\bar{w})$ and let $\bar{N} = \bar{w}/(2C)$.

Well-definedness of $s^*(C,n)$. For given (C,n), the cutoff $s^*(C,n)$ solves $U_S^\star(C,s) = U_D^\star(C,n)$ where both value functions are given by (7) and (10). On the interior region, $U_S^\star(C,s)$ is strictly decreasing in s (envelope theorem: $\partial U_S^\star/\partial s = -I_S^\star < 0$). The cutoff is unique whenever $U_D^\star$ lies strictly between $U_S^\star(C,\bar{s})$ and $U_S^\star(C,\underline{s})$; otherwise a corner cutoff applies.

Φ_C maps $[0,\bar{N}]$ into $[0,\bar{N}]$. Since $[1-F(s^*)] \in [0,1]$:

$$S_D(C,n) = (1-F(s^*))\,\frac{\bar{w}^2-C^2}{2\bar{w}} + \frac{C^2}{2\bar{w}} \le \frac{\bar{w}^2-C^2}{2\bar{w}} + \frac{C^2}{2\bar{w}} = \frac{\bar{w}}{2} = \mathbb{E}[w].$$

Also $S_D(C,n) \ge 0$. Hence $\Phi_C(n) = S_D(C,n)/C \in [0,\bar{N}]$.

Continuity. The map $n \mapsto U_D^\star(C,n)$ is continuous by the explicit formula (10) and continuity of $\nu(C,n)$ in n. The map $U_D^\star \mapsto s^*$ is continuous on the interior region by strict monotonicity of $U_S^\star$ in s. The uniform CDF F is continuous. Hence Φ_C is continuous.

Existence by the intermediate value theorem. Define $f(n) = \Phi_C(n) - n$ on $[0, \bar{N}]$. Since $\Phi_C(n) \geq 0$ for all n, we have $f(0) = \Phi_C(0) \geq 0$. Since $\Phi_C(n) \leq \bar{N}$ for all n, we have $f(\bar{N}) = \Phi_C(\bar{N}) - \bar{N} \leq 0$. The function f is continuous. By the intermediate value theorem, there exists $N_D^* \in [0, \bar{N}]$ with $f(N_D^*) = 0$, i.e. $\Phi_C(N_D^*) = N_D^*$. □

A.4 Proof of Proposition 4

Proof. By hypothesis, $\mathcal{V}$ is continuous on $[\underline{C}, \bar{w}]$. By Weierstrass, a maximizer $C^\star \in [\underline{C}, \bar{w}]$ exists. By (A1), $\mathcal{V}'(\underline{C}) > 0$, so the maximizer cannot be $\underline{C}$. By (A2), $\mathcal{V}'(C) < 0$ for $C \in [\bar{C}, \bar{w}]$, so the maximizer cannot lie in that region. Hence $C^\star \in (\underline{C}, \bar{C}) \subset (\underline{C}, \bar{w})$. □

Discussion of Remark 1. Under (P1), $\mu(\underline{C}) > \bar{s}$ ensures all skill types are interior at $\underline{C}$, so the direct incentive derivative of liveness is strictly positive there. The cutoff and overhead responses at $\underline{C}$ are bounded, making (A1) plausible but not globally proved from primitives. Under (P2), solo eligibility has contracted on $[\bar{C}, \bar{w}]$, making the participation and overhead effects likely to dominate, but this depends on the specific parameter configuration. These are interpretive conditions, not tight closed-form implications. □

A.5 Proof of Proposition ??

Proof. By construction, (17) multiplies the delegated liveness contribution $N_D(C)p(J^\star(C))$ by $(1 - \rho_D)$, while the solo contribution $N_S(C)\bar{p}_S(C)$ carries no such discount. The asymmetry is therefore a direct consequence of the payoff specification (8). □

A.6 Proof of Proposition ??

Proof. Differentiate $\mathcal{V}(C) = \mathcal{L}(C) - \chi N(C)^2/2$ along the equilibrium. Under the uniform benchmark,

$$N_S(C) = \frac{\bar{w} - C}{\bar{w}} \cdot \frac{s^*(C) - \underline{s}}{\bar{s} - \underline{s}},$$

so the solo mass is determined by $(C, s^*(C))$, and delegated investment $J^\star(C) = J^\star(C, N_D(C))$ is determined by $(C, N_D(C))$. Hence liveness can be written as a function $\mathcal{L}(\mu(C), s^*(C), N_D(C))$

of these three state variables, where $\mu(C) = a(R + \underline{\phi}C)$ governs investment incentives through both the solo and delegated composites. The chain rule then gives:

$$\mathcal{L}'(C) = \frac{\partial \mathcal{L}}{\partial \mu}\mu'(C) + \frac{\partial \mathcal{L}}{\partial s^*}s^{*\prime}(C) + \frac{\partial \mathcal{L}}{\partial N_D}N_D'(C).$$

Subtracting $\chi N(C)N'(C)$ gives (20).

Direct incentive effect. $\mu'(C) = a\underline{\phi} > 0$. From $\bar{p}_S = 1 - (\underline{s} + s^*)/(2\mu)$, $\partial \bar{p}_S/\partial \mu = (\underline{s} + s^*)/(2\mu^2) > 0$. Similarly, $\partial p(J^\star)/\partial \mu = (1/\nu^2) \cdot a(1-\rho_D)N_D > 0$ (since $\nu = a(1-\rho_D)N_D\mu$). Hence $(\partial \mathcal{L}/\partial \mu)\mu'(C) > 0$.

Remaining terms. $s^{*\prime}(C)$, $N_D'(C)$, and $N'(C)$ are determined by implicit differentiation of the equilibrium fixed-point system; their signs depend on the relative movements of $U_S^\star$ and $U_D^\star$ with C and are in general ambiguous. Hence the net sign of $\mathcal{V}'(C)$ is ambiguous. □

A.7 Proof of Proposition ??

Proof. Fix $C \in [\underline{C}, \bar{w})$. Define $\mathcal{V}(C, s)$ as protocol welfare when the cutoff is s, treating C as fixed:

$$\mathcal{V}(C,s) = N_S(C,s)\bar{p}_S(C,s) + (1-\rho_D)N_D(C,s)p(J^\star(C, N_D(C,s))) - \frac{\chi}{2}N(C,s)^2.$$

Under the uniform distributions:

$$N_S(C,s) = \frac{\bar{w}-C}{\bar{w}} \cdot \frac{s-\underline{s}}{\bar{s}-\underline{s}}, \quad S_D(C,s) = \left(1 - \frac{s-\underline{s}}{\bar{s}-\underline{s}}\right)\frac{\bar{w}^2 - C^2}{2\bar{w}} + \frac{C^2}{2\bar{w}}, \quad N_D(C,s) = \frac{S_D(C,s)}{C}.$$

Derivatives with respect to s. Let $\Delta = \bar{s} - \underline{s}$. $\partial N_S/\partial s = (\bar{w} - C)/(\bar{w}\Delta) > 0$. $\partial N_D/\partial s = -(\bar{w}^2 - C^2)/(2\bar{w}C\Delta) \leq 0$. $\partial N/\partial s = \partial N_S/\partial s + \partial N_D/\partial s$, with ambiguous sign.

Planner derivative. Differentiating $\mathcal{V}(C,s)$ with respect to s:

$$\frac{\partial \mathcal{V}}{\partial s} = \bar{p}_S\frac{\partial N_S}{\partial s} + N_S\frac{\partial \bar{p}_S}{\partial s} + (1-\rho_D)\frac{\partial [N_D p(J^\star)]}{\partial s} - \chi N\frac{\partial N}{\partial s}.$$

$\Xi(C,s) = (\bar{w} - C)/[\bar{w}(\bar{s} - \underline{s})] > 0$ is the density of wealth-eligible agents at the cutoff s, i.e. the mass of agents with $w_i \geq C$ and $s_i \in [s, s + ds]$. This equals $\partial N_S/\partial s$, which also equals

the marginal mass reassigned from delegation to solo when the cutoff rises. $\Lambda(C, s)$ is defined residually by (22), which makes (23) tautologically true. The content of the proposition is that at $s = s_{\text{eq}}$, where $U_S^\star - U_D^\star = 0$, the sign of $\partial\mathcal{V}/\partial s$ is determined entirely by Λ, which captures protocol-level composition effects not internalized by private agents.

Evaluation at s_{eq}. At the equilibrium cutoff s_{eq}, $U_S^\star(C, s_{\text{eq}}) = U_D^\star(C, N_D)$, so by (22):

$$\left.\frac{\partial \mathcal{V}}{\partial s}\right|_{s=s_{\text{eq}}} = \Lambda(C, s_{\text{eq}}).$$

If $\Lambda(C, s_{\text{eq}}) > 0$, then $\mathcal{V}$ is increasing in s at the equilibrium cutoff, implying $s_{\text{pl}}(C) > s_{\text{eq}}(C)$. □

A.8 Proof of Proposition 5

Proof. The proof is structurally identical to Proposition 1 with $\underline{\phi}$ replaced by $\phi_S(C, \kappa)$ and $\mu(C)$ replaced by $\mu_\kappa(C) = a(R + \phi_S(C, \kappa)C)$.

Value function. Substituting $I_S^\star = \log(\mu_\kappa/s)/a$: $U_S^\star(C, \kappa, s) = R - (s/a)(1 + \log(\mu_\kappa(C)/s))$ by the same algebra. When $\mu_\kappa \le s$: $I_S^\star = 0$ and $U_S^\star = -\phi_S(C, \kappa)C$.

Monotonicity in ϕ_S. $\partial\mu_\kappa/\partial\phi_S = aC > 0$. Hence $\partial I_S^\star/\partial\phi_S = C/\mu_\kappa > 0$ on the interior region. Under fixed composition, $\phi_S = \underline{\phi}(1 + \kappa/N)$ and $\partial\phi_S/\partial\kappa = \underline{\phi}/N \ge 0$, so $I_S^\star$ is weakly increasing in κ. □

A.9 Proof of Proposition 6

Proof. The proof follows Proposition 2 with $\underline{\phi}$ replaced by $\phi_S(C, \kappa)$ and $\nu(C, N_D)$ replaced by $\nu_\kappa(C, N_D) = a(1 - \rho_D)N_D(R + \phi_S(C, \kappa)C)$.

Interior value function. Substituting $J^\star = \log(\nu_\kappa)/a$:

$$U_D^\star(C, \kappa, N_D) = (1 - \rho_D)R - \rho_D\phi_D(C, \kappa)C - \frac{1}{aN_D}(1 + \log\nu_\kappa),$$

noting that ϕ_D (not ϕ_S) enters the level term because the systemic slashing uses ϕ_D.

Corner value function. When $\nu_\kappa \le 1$, $J^\star = 0$ and:

$$U_D(0; C, \kappa, N_D) = (1 - \rho_D)(0 - \phi_S C) - \rho_D\phi_D C = -(1 - \rho_D)\phi_S(C, \kappa)C - \rho_D\phi_D(C, \kappa)C.$$

Monotonicity in ϕ_S. $\partial \nu_\kappa / \partial \phi_S = a(1-\rho_D) N_D C > 0$, so $\partial J^\star / \partial \phi_S = (C/\nu_\kappa)(1-\rho_D) N_D > 0$. Under fixed composition, $\partial \phi_S / \partial \kappa = \underline{\phi}/N \geq 0$, so $J^\star$ is weakly increasing in κ.

ϕ_D excluded from FOC. The first-order condition $a(1-\rho_D) e^{-aJ}(R + \phi_S C) = 1/N_D$ depends on ϕ_S but not ϕ_D, since ϕ_D appears only in the level term $-\rho_D \phi_D C$.

Monotonicity in C (holding ϕ_S fixed). $\partial \nu_\kappa / \partial C = a(1-\rho_D) N_D \phi_S > 0$, so $J^\star$ is strictly increasing in C when ϕ_S is held fixed. □

A.10 Fixed-point existence under penalties

Proof. The argument of Appendix A.3 carries over directly. For given (C, κ) with $C \in [\underline{C}, \bar{w})$, replace $U_S^\star(C, s)$ with $U_S^\star(C, \kappa, s)$ from (33) and $U_D^\star(C, n)$ with $U_D^\star(C, \kappa, n)$ from (35) (with $N_D = n$). The bound $\Phi_{C,\kappa}(n) \in [0, \bar{w}/(2C)]$ holds by the same delegated-stake inequality. Continuity holds by continuity of ϕ_S, ν_κ, $J^\star$, and $U_S^\star$ in their arguments. The intermediate value theorem applied to $f_\kappa(n) = \Phi_{C,\kappa}(n) - n$ yields an equilibrium $N_D^*(C, \kappa) \in [0, \bar{w}/(2C)]$. □

A.11 Proof of Proposition 7

Proof. Under fixed composition, N is constant and $\phi_S(C, \kappa) = \underline{\phi}(1 + \kappa/N)$, so

$$\mu_\kappa(C) = a\Big(R + \underline{\phi} C\Big(1 + \frac{\kappa}{N}\Big)\Big), \quad \frac{\partial \mu_\kappa}{\partial C} = a\underline{\phi}\Big(1 + \frac{\kappa}{N}\Big), \quad \frac{\partial \mu_\kappa}{\partial \kappa} = \frac{a\underline{\phi} C}{N}, \quad \frac{\partial^2 \mu_\kappa}{\partial C\, \partial \kappa} = \frac{a\underline{\phi}}{N} > 0.$$

Under fixed composition, $\mathcal{L} = N_S \bar{p}_S + (1-\rho_D) N_D p(J^\star)$ where $\bar{p}_S = 1 - (\underline{s} + s^*)/(2\mu_\kappa)$ and $p(J^\star) = 1 - 1/\nu_\kappa$ with $\nu_\kappa = a(1-\rho_D) N_D \mu_\kappa$.

Computing the mixed partial:

$$\frac{\partial^2 \mathcal{L}}{\partial C\, \partial \kappa} = \frac{a\underline{\phi}}{N}\left[N_S \frac{\underline{s} + s^*}{2\mu_\kappa^2} + (1-\rho_D)^2 \frac{N_D^2}{\nu_\kappa^2}\right] - \frac{a^2 \underline{\phi}^2 C(1 + \kappa/N)}{N}\left[N_S \frac{\underline{s} + s^*}{\mu_\kappa^3} + (1-\rho_D)^2 \frac{2N_D^2}{\nu_\kappa^3}\right].$$

The first term is positive. The second is negative. The net is positive whenever $\mu_\kappa > 2a\underline{\phi} C(1+\kappa/N)$. Since $\phi_S = \underline{\phi}(1 + \kappa/N)$, this condition is equivalent to

$$R > \underline{\phi}\, C\Big(1 + \frac{\kappa}{N}\Big),$$

which is precisely condition (37) divided by a. In large networks κ/N is small, so this is close to the simpler condition $R > \underline{\phi}C$, which holds under standard parameters. This establishes $\partial^2\mathcal{V}/\partial C\,\partial\kappa > 0$ under fixed composition. The IFT then gives (38). □

A.12 Proof of Proposition 8

Proof. Hold composition (N_D) fixed. The interior delegated value function is

$$U_D^\star(C,\kappa,N_D) = (1-\rho_D)R - \rho_D\phi_D(C,\kappa)C - \frac{1}{aN_D}(1+\log\nu_\kappa),$$

where $\nu_\kappa = a(1-\rho_D)N_D(R+\phi_S(C,\kappa)C)$.

Common-slashing channel. $\partial U_D^\star/\partial\phi_S = -(1/(aN_D\nu_\kappa))\cdot a(1-\rho_D)N_DC = -(1-\rho_D)C/\nu_\kappa < 0$. This is confirmed by the envelope theorem: $\partial U_D/\partial\phi_S = (1-\rho_D)(p(J)C-C) = (1-\rho_D)C(p(J)-1)$; at $J=J^\star$, $p(J^\star)-1 = -1/\nu_\kappa$, so $\partial U_D^\star/\partial\phi_S = -(1-\rho_D)C/\nu_\kappa < 0$.

Since $\partial\phi_S/\partial\kappa = \underline{\phi}/N > 0$:

$$\frac{\partial U_D^\star}{\partial\phi_S}\cdot\frac{\partial\phi_S}{\partial\kappa} = -(1-\rho_D)\frac{C\underline{\phi}}{N\nu_\kappa} < 0.$$

Concentration level-cost channel. $\partial[-\rho_D\phi_DC]/\partial\kappa = -\rho_DC\,\partial\phi_D/\partial\kappa = -\rho_DC\underline{\phi}\alpha_D < 0$.

Combining.

$$\frac{\partial\pi_D}{\partial\kappa} = -(1-\rho_D)\frac{C\underline{\phi}}{N\nu_\kappa} - \rho_D\underline{\phi}\alpha_DC < 0.$$ □

A.13 Proof of Proposition 9

Proof. Under fixed composition, N, N_S, N_D, s^* are all constant in κ, so $g(N) = \chi N^2/2$ is constant and $\partial\mathcal{V}/\partial\kappa = \partial\mathcal{L}/\partial\kappa$.

Under fixed composition, $\phi_S(C,\kappa) = \underline{\phi}(1+\kappa/N)$ and $\mu_\kappa(C) = a(R+\phi_SC)$, $\nu_\kappa = a(1-\rho_D)N_D\mu_\kappa$. $\partial\mu_\kappa/\partial\kappa = a\underline{\phi}C/N \geq 0$, with strict inequality since $C \geq \underline{C} > 0$.

Solo: $\partial\bar{p}_S/\partial\kappa = [(\underline{s}+s^*)/(2\mu_\kappa^2)]\cdot a\underline{\phi}C/N > 0$. Delegated: $\partial p(J^\star)/\partial\kappa = (1/\nu_\kappa^2)\cdot a(1-\rho_D)N_D\cdot a\underline{\phi}C/N > 0$.

Hence $\partial\mathcal{L}/\partial\kappa = N_S(\partial\bar{p}_S/\partial\kappa) + (1-\rho_D)N_D(\partial p/\partial\kappa) > 0$ whenever at least one sector with positive mass is on the interior investment region (so that the relevant derivative is non-zero). □

A.14 Proof of Proposition 10

Proof. By continuity of $\mathcal{V}(C, \kappa)$ on $[0, 1]$, a maximizer exists. Because $\partial\mathcal{V}/\partial\kappa|_{\kappa=0^+} > 0$, the maximizer cannot be at $\kappa = 0$. Because $\partial\mathcal{V}/\partial\kappa|_{\kappa=1} < 0$, the maximizer cannot be at $\kappa = 1$. Hence any maximizer lies in $(0, 1)$. □