Liquid staking*

Alfred Lehar†
Haskayne School of Business
University of Calgary

Christine A. Parlour‡
Haas School of Business
UC Berkeley

Kathy Yuan§
LSE

May 18, 2026

Preliminary and Incomplete
Comments Welcome

Abstract

Liquid staking allows agents to sell ownership of an illiquid claim to satisfy a liquidity need. We develop a model of liquid staking and characterize the effect of the secondary market on protocol stability. We establish that the liquid market has two effects: first, it allows agents to redeem illiquid assets and thus reduces run risk on the protocol, but also conveys information and can act as a coordination mechanism and increase market run risk. Using novel data on the Lido protocol, we present stylized facts on the staking market and relate our results to the design of digital deposits.

Keywords: Staking, Blockchain, Decentralized Finance, Tokenization

*We thank Samuel Ergando, Jialiang Lou, and Yichen Lou for research assistance and seminar participants at the Oxford Man Institute of Quantitative Finance for helpful comments.

†Corresponding author, email: alfred.lehar@haskayne.ucalgary.ca, Tel: (403) 220 4567. Alfred Lehar is grateful to SSHRC, the Ethereum Foundation and the Fintech Dauphine Chair for financial support.

‡email:parlour@berkeley.edu

§email:k.yuan@lse.ac.uk

1 Introduction

Banks' role in liquidity transformation has, since the seminal work of Diamond and Dybvig (1983), been tied to fragility. In a bank run, suppliers of capital faced with a liquidity shock, withdraw limited liquid capital, to the detriment of the bank. Given the increasing experimentation, and in some cases adoption, of tokenized deposits, in this paper we use an existing blockchain liquidity transformation protocol to investigate whether tokenized deposits may be more or less stable than traditional deposits.

We analyze, theoretically and empirically, a liquid staking protocol on the Ethereum blockchain – the Lido protocol. This protocol is analogous to traditional banks because it pools relatively liquid deposits (in the form of ETH) and invests them in an illiquid asset. The illiquid asset is the consensus mechanism on the Ethereum blockchain. (We describe the mechanics of this process in detail in an appendix.) Briefly, to participate, validators must post collateral typically provided by third parties. If the validators perform successfully, they obtain ETH rewards. Validators' rewards, analogous to returns on banks loans, are then repaid to the collateral providers. The staked collateral is illiquid because stakers must wait when exiting the network to make sure that they can be held accountable for any misbehavior.[1] Lido protocol, acting as an intermediary, pools small deposits that are below the threshold to provide validator collateral, selects validators to provide with collateral, and provides returns to its depositors as well as a buffer to handle withdrawals.

When the protocol first launched, depositors who had liquidity needs could join a Lido withdrawal queue. More recently, a market has emerged that facilitates direct trade between tokens representing staked ETH and ETH. In short, instead of withdrawing the tokens that they have staked in Lido, an impatient depositor can simply sell their staked tokens. These tokens are freely tradable and thus allow users to quickly enter and exit the staking market. With over $61 billion invested, liquid staking is the largest application today in a system of smart contracts on the blockchain, also known as Decentralized Finance (DeFi).

Liquid staking transforms an illiquid instrument into something that is liquid and tradable. To exit their position, liquid stakers can either directly sell their tokens on a secondary market, potentially with price impact, or they can redeem their tokens with the liquid staking protocol after incurring private waiting costs. Redemption requests are served by the liquid staking protocol from a treasury, or liquidity reserve. Currently, traditional bank deposits can only be redeemed with one counter-party - the issuing bank and only at face value. However, if deposits are tokenized, they become exchangeable outside the repository bank. In this way, tokenized deposits are more liquid than demand deposits as they can be used without first being withdrawn.

We emphasize that deposit tokenization is still experimental. Citi Token Services offers blockchain based token payments to institutional investors. In 2023 the Swiss Banking Association ran a tokenized deposit project with a dozen banks, while in Hong Kong, VISA and two banks com-

[1] Stakers on Ethereum need to give a withdrawal notice after which they go through an exit queue that allows about 2025 validators to exit per day. Subsequently, departed validators must wait for 256 Epochs (about 27 Hours) before they can withdraw their funds.

pleted a tokenized deposit pilot program. As part of project Agorá at the Bank of International settlements, financial firms explored tokenized deposits for cross border payments. Tokenized deposits that are freely exchangeable resemble a liquid market for deposits.

We present a stylized run model in the spirit of Goldstein and Pauzner (2005). Investors use an intermediary to access an illiquid asset that pays a random return. Our technical innovation is to include a market for claims that generates an informative price that informs withdrawal decisions. Additionally, we endogenize both non-strategic and strategic withdrawals by relating the former to general yield conditions in other decentralized protocols and the latter to the information revealed in the market signal. We compare the stability of the liquid protocol to a traditional intermediary for which there is no market for claims.

An important role in our model falls to secondary markets. Stakers who are hit by a liquidity shock and do not want to wait in the redemption queue can sell the token immediately on the secondary market. Token demand on this market is also subject to random shocks which result in price fluctuations. Stakers combine information from token prices with their own signal to decide if they should run or not. Because market prices influence the stakers' run decisions, a secondary market can increase financial fragility of the staking protocol through the coordination channel; however, as we show both theoretically and numerically, the net effect of a sufficiently deep secondary market is stabilising.

Our analysis yields several novel insights into the impact of secondary markets on the stability of staking protocols. The secondary market plays multiple roles. First, it serves as a signal to external institutional investors, who are the source of emergency funding to the protocol in the event of a run and are crucial for stemming liquidity runs. Second, the secondary market acts as a mechanical shock absorber by diverting redemption pressure. Third, it creates informational fragility by coordinating staker panic. We find that a deep secondary market and a more collateralizable market price lower the probability of panic runs and protocol failure. A numerical exercise calibrated to the Lido protocol quantifies the magnitude: liquid staking reduces the liquidity-driven run probability significantly relative to traditional staking when the secondary market is moderately deep, but this advantage erodes, and can reverse, when the market is thin.

Interestingly, we find that institutional markets, namely those in which the protocol itself could raise additional funds, can be stabilizing. However, we find that other stabilization tools, such as a large retained treasury, can be destabilizing. This arises because while more money on hand can satisfy redemption demand, it also reduces the per-token payoff for remaining in the protocol and thus encourages redemptions. Similarly, we demonstrate that increasing staking rewards stabilizes the protocol up to a certain point (by reducing liquidity run incentives) but severely destabilizes it thereafter (by increasing insolvency risk, which feeds back into liquidity run risk).

To ground our theoretical framework, we present novel stylized facts using data from the Lido protocol. We document that the stETH discount Granger-causes redemption activity, while the redemption queue is itself more influenced by market sentiment than by past discounts, reflecting that withdrawal decisions are also driven by outside opportunities. These observations confirm our model's core premise: stakers actively choose between the immediate liquidity of the secondary market and the delayed redemption queue based on prevailing market frictions.

We also map the model's comparative statics to the data and find that staking rewards and the protocol treasury are negatively related to the stETH discount (i.e., associated with deeper discounts), while the collateralizability of stETH is positively related to it (i.e., associated with narrower discounts).

The Lido protocol can be viewed as an experiment of the Jacklin Critique. Jacklin (1987) famously argued that for effective liquidity risk sharing, a deposit contract should function like dividend-paying equity paired with a secondary market. Furthermore, the critique suggests that frictionless arbitrage between this secondary market and the intermediary's redemption window would inevitably unravel the bank's deposit franchise. Our model formalizes this argument by analyzing tradable and redeemable deposit contracts, but we extend it by introducing complex informational feedback loops. We demonstrate that the secondary market plays dual roles. Mechanically, it acts as a liquidity outlet that absorbs early withdrawals, thereby reducing direct redemption pressure on the protocol. Informationally, the market price serves as a public signal: it acts as proof of collateral that allows the protocol to secure external liquidity, while simultaneously allowing stakers to update their beliefs and coordinate their run decisions. Ultimately, our model establishes a key limit to the Jacklin critique: the franchise-destroying arbitrage relies on a frictionless secondary market. Once information frictions and market noise are introduced, the arbitrage is impeded, allowing the liquid staking protocol to coexist with the secondary market.

The seminal work of Diamond and Dybvig (1983) provide a stylized framework to understand liquidity risk in banks. Observing that demand deposits facilitate liquidity risk sharing, they illustrate how runs can occur. More recently, Goldstein and Pauzner (2005), show how signals about the underlying state can coordinate run decisions among depositors and make the run equilibrium unique. This important paper links the coordination problem of Diamond and Dybvig (1983) to the solvency problem underlying bank failure. Correia, Luck, and Verner (2024) gather a comprehensive data set to investigate bank failures. Attempting to disentangle liquidity and solvency motives for bank failure. They document that bank failures with runs are strongly related to weak fundamentals. Jacklin (1987) argues that dividend paying equity will also facilitate risk sharing, if there is a market for ex dividend shares. However, the recent work by Goldstein, Ozdenoren, and Yuan (2013) shows that the secondary market price might affect coordination decisions of market participants in the form of trading frenzy and hence affect the real decisions. How changes in the contract between the suppliers of capital and the intermediary affect optimality are explored in Green and Lin (2003) and Peck and Shell (2003). We contribute to this literature by studying tradable and redeemable deposit contracts, which is the first in literature.

2 Liquid Staking and Lido

In proof of stake systems, tokens are staked, i.e. pledged as collateral, to support the consensus mechanism. Value accrues to the staked tokens because they are remunerated by validator rewards. These rewards are not riskless as validators can be penalized through "slashing" if they validate a block that is inappropriate or behave inappropriately in other ways. For example, slashing can occur if a validator signs conflicting blocks, or if they experience significant down-

time. Slashing is effectively a fine. To prevent large scale coordinated attacks some blockchains impose correlated slashing penalties if a large fraction of the validators misbehave. In Ethereum the usual slashing penalty of 1 ETH out of 32 ETH staked increases to 25 ETH when 25% of the validators get slashed. We emphasize that slashing can occur even if the validator is reputable. Concretely, slashing could be a systematic event as professional validators share common architecture. For example, consider the infrastructure concentration as shown in Figure 1.[2] About 22% of all Ethereum nodes run on Amazon (AWS). An unexpected outage could trigger a massive slashing event.

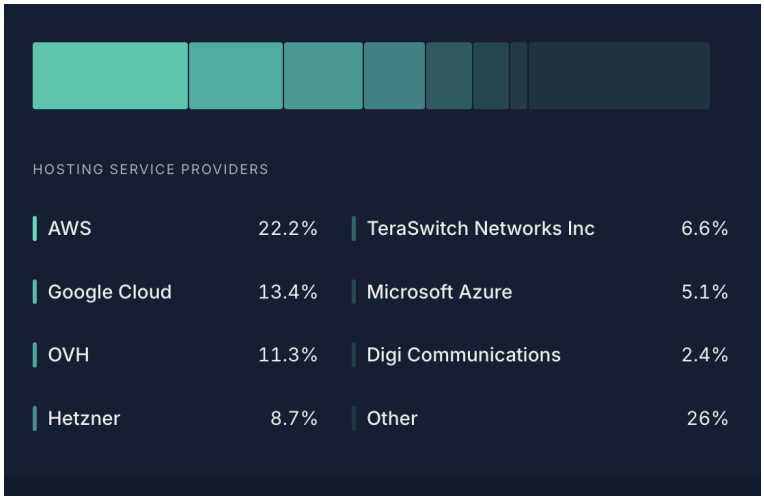

Figure 1. The distribution of Cloud Service Providers among the professional validators.

Directly staking on the blockchain requires chunks of 32 ETH (about $85,000) and is technically non-trivial. The nodes have to be properly configured, need to be connected to maximal extractable value (MEV) relays to maximize revenue, and always have to be online to avoid slashing. To allow users with less capital and technical sophistication to participate in the staking market a class of intermediaries has arisen that pools smaller amounts of cryptocurrencies, and deposits them with different validators to reap the rewards which they pay out. It is important to note that ETH staked with validators is illiquid. That is, it cannot be withdrawn at will. The DeFi solution to this has been to introduce "staking derivatives." These tokens represent claims on the underlying staked ETH. Figure 2 illustrates how a stylized staking protocol operates. The primary focus of our study (and empirical work) is the Lido protocol, the largest liquid staking protocol.

When users deposit their ETH in the protocol, they essentially receive receipt tokens that are a claim on the underlying staked assets. The holder of these receipt tokens will eventually be able to redeem them (if the underlying have not been slashed). However, because the underlying staked tokens are not liquid, redemption is slow. Specifically, there is a queue (which is observable) and waiting costs. Redemption requests are handled sequentially and the tokens remained locked for days to weeks. This lock-up mechanism prevents large numbers of tokens leaving the system. For the stakers who are impatient, there is a liquid market in which the

[2]Current Data are available from `https://explorer.rated.network/network`

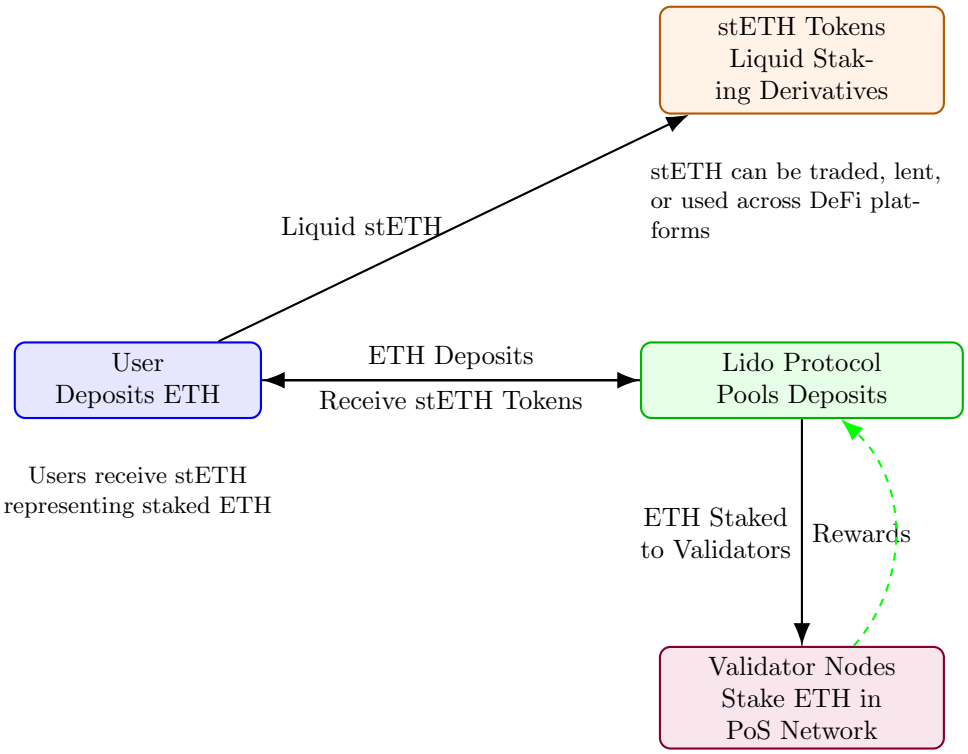

Figure 2. Schematic Showing how Liquid Staking Operates

staked tokens can be exchanged for the underlying. Because of slashing risk it is not only a market for liquidity, as there is a small risk associated with the receipt tokens. Specifically in a large correlated slashing event would put the fundamental value of the remaining funds well below par. Since it is not known where specific servers of a staking protocol reside, e.g. how many LIDO nodes are on AWS, there is ambiguity how much a specific liquid staking protocol is exposed to the slashing. The historical price relationship between ETH and staked ETH shows that in our sample, no correlated slashing event occurred, and so stETH trades close to par with ETH.

3 Model

Our interest in liquid staking is motivated by banking stability; and first we describe the nomenclature and institutional environment that differs from the more familiar banking setting. In doing so, we try to point out how the decentralized finance environment relates to standard intermediation.

Illiquid Asset: Common with canonical banking problems, there is an illiquid asset which generates a risky return. In our framework, returns arise from the consensus mechanism of a proof of stake (PoS) blockchain, which we take to be Ethereum. The PoS consensus mechanism is maintained by staked ETH tokens which in turn earn rewards for attesting blocks, and

additionally may earn what are known as Execution Layer (EL) rewards such as "maximum extractible value (MEV)" and transaction fees. (See the Appendix for details.) In this economy, ETH is the numeraire.

Validators: Charged with adding valid blocks to the blockchain, validators lock up tokens as collateral. Importantly, solo-staking requires a minimum investment, currently set at 32 ETH. If the block is valid, newly minted tokens accrue to the validator, whereas if a block is invalid, the tokens are "slashed." Thus, the return on staked tokens is risky, and depositing collateral with a validator is akin to investing in an illiquid, risky, asset.

Staking Service: A staking service provider pools deposited tokens to reach the required minimum staking size, and stakes the tokens on the depositors' behalf. The staking service provider passes through the stake reward (90% in Ethereum) to the depositors. However, if the service provider acts maliciously or unreliably, the token depositors are exposed to operational risk. Additionally, as noted above tokens deposited with validators are illiquid and bear risk. We characterize the fundamentals of the staking service as an unobservable random variable $\tilde{\theta}$, which represents the quality of the staking service provider, including the reward, the security, possible slashing penalty, and the balance sheet strength. Economically, the staking service performs deposit aggregation and investment in illiquid asset as do banks.

Withdrawals: Unlike small demand deposits in the banking system which, as a means of payment, can be withdrawn at any time (up to a limit), tokens deposited with a staking service can be liquid or illiquid. In a "traditional" staking service, staked tokens are locked, cannot be transacted or used as collateral to earn yield, rendering them illiquid. To withdraw their staked tokens when facing liquidity shocks, depositors in traditional staking face a lock-up period. By contrast, a liquid staking service provider solves this liquidity problem by providing the depositors with a receipt in the form of new tokens, denoted as $stETH$. These tokens can either be redeemed, sold or used directly as collateral in other applications. The focus of our paper is on the effect of this liquid staking market on the stability of the staking service.

Agents: Our model comprises a blockchain with a proof-of-stake (PoS) consensus mechanism that generates value, a liquid or traditional staking service provider, and a continuum of token holders. By assumption, all small token holders do not have the minimal required size to stake their tokens and thus have to deposit them into a staking service provider. There are two types of staking service: traditional and liquid. In the latter case, the receipt tokens can be traded.

Sequence of Events

The sequence of events for this two-period, three-date investment game, is illustrated in Figure 3. A mass S of small token holders, each holding a unit of ETH has access to a staking protocol that deposits all stakes with validators. The protocol promises a staking reward of $1 + r$ per token at date $t = 2$, conditional on its survival. The protocol's ability to meet this obligation depends on an unobservable fundamental $\tilde{\theta}$, which captures the overall quality of the staking service, including validator performance, slashing risk, and balance sheet strength. Of the total funds raised, T is held in a treasury as a liquidity reserve and $(I - T)$ is invested in staking, generating a random gross return of $\left(\tilde{\theta} + 1\right)$ per unit invested. The realized value of the staked

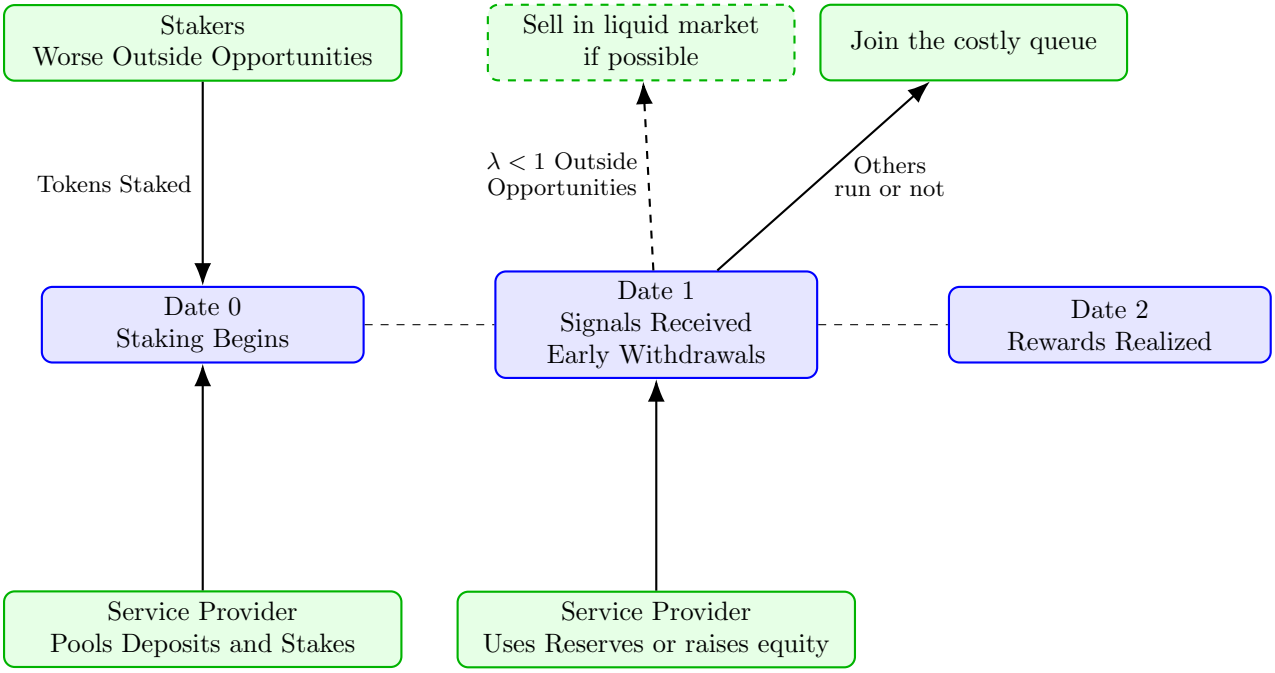

Figure 3. The Sequence of Events in the model

assets, $(I-T)\left(\tilde{\theta}+1\right)$, determines how much external liquidity the protocol can raise in the event of a run (see below).

We assume that the random variable $\tilde{\theta} = \bar{\theta} + \sigma_\theta \tilde{\epsilon}_\theta$ where $\tilde{\epsilon}_\theta$ is normally distributed with a mean of 0 and a standard deviation of 1, i.e., a standard normal variable.

Before depositing in the staking service, each small token holder receives an independent signal $\widetilde{c}$ drawn from a distribution $F_0(\cdot)$. Here, $\widetilde{c}$ represents the payoff from a decentralized application, generating yield, that is available to them. The draws of this outside opportunity are independent because there is no centralized source of information about cryptocurrency investment possibilities, and different investors have differential access to different blockchains. Small token holders invest in the staking service if $r \geq c$, and thus a mass of $F_0(r)S$ will deposit funds in the staking protocol. For notational compactness, we let $I = F_0(r)S$ denote the mass of small token holders who deposit their tokens in the staking protocol. This is thus, the size of the protocol.

At date 1, the small token holders who joined the staking service receive a new independent signal about another possible return $\widetilde{c}$, drawn from $F_1(\cdot)$. Once again, small token holders compare r to c, and if $\widetilde{c} \geq r$, they withdraw from the staking service. Let $\lambda = [1 - F_1(r)]F_0(r)S$ denote the mass of stakers who withdraw for these outside investment reasons. Subsequently, each remaining staker observes a private signal $s_i = \tilde{\theta} + \sigma_s \tilde{\epsilon}_i$ about the staking service's fundamental. Here $\tilde{\epsilon}_i$ is an independently drawn standard normal random variable. We denote the p.d.f. and the c.d.f. of the standard normal distribution by ϕ and Φ, respectively, and the precision of the noise $\tau_s = 1/\sigma_s^2$. Based on this signal, the remaining stakers decide if they should withdraw their tokens for strategic reasons or wait for the final payoff at date 2. For notational compactness,

we denote $\kappa = I - \lambda$. They can withdraw via one of the two alternatives: a fast route by selling the staked token in the market or a slow one by redeeming it via redemption queue. Joining the redemption queue, albeit slow and relatively more costly since the staking protocol takes much longer time to meet redemption requests, it comes with the benefit of observing the stETH price traded in the market.

We denote the action of a strategic staker, by $a_i \in \{0, 1\}$, where $a_i = 0$ represents "not withdraw" and $a_i = 1$ represents "withdrawal." Let $A \equiv \int_0^1 a_i di \in [0, 1]$ denote the mass of withdrawals from strategic token depositors with measure 1. For expositional ease, we assume the total withdrawal from strategic token depositors is $\kappa\Phi^{-1}(A) > 0$, and denote the total withdrawal from the protocol as $\mathcal{W} = \kappa\Phi^{-1}(A) + \lambda$.

Strategic withdrawal at $t = 1$ depends on stakers' beliefs about the viability of the staking service. Withdrawing early is costly, and stakers take this cost as given, which we denote by η. Strategic withdrawal can occur either by the selling claims to the staked tokens (if a market is operating), at a price p or by joining a redemption queue. In practice, mirroring the Ethereum blockchain itself, staked ETH can be withdrawn but only with a lag, by joining a redemption queue with a slow service rate. We model these market microstructure in details in the next two sections.

4 Traditional Staking

In the traditional staking case, the staking service now has various means to fulfill withdrawals: It can maintain a treasury, T, which is unstaked ETH that it can pay out, or it can raise external financing based on their investment and the fundamentals of the protocol $(I - T)\tilde{\theta}$. Formally, the protocol fails if

$$\text{ETH Redemption Queue} \geq \underbrace{T + l}_{\text{Internal and External ETH Liquidity}}, \tag{1}$$

where l is the external liquidity that the protocol can raise. We assume that the external liquidity provider does not observe the protocol's fundamental $\tilde{\theta}$ directly, but instead receives a noisy signal $\tilde{\theta} + \sigma_l\tilde{\epsilon}_l$, where $\tilde{\epsilon}_l$ is an independently drawn standard normal random variable and $\sigma_l > 0$ measures the noise in the external provider's assessment. The provider extends liquidity based on this signal and the invested amount $(I - T)$:

$$l = \rho\,(I - T)\,(\theta + \sigma_l\tilde{\epsilon}_l + 1), \tag{2}$$

where $0 < \rho < 1$ reflects the cost of illiquidity of external finance. Larger ρ, the more liquid the external financing market. When $\sigma_l = 0$, the external provider observes the fundamental perfectly, and we recover the benchmark case.

The protocol faces two distinct failure channels. First, a *liquidity failure* at date 1: the protocol fails if aggregate withdrawal demand exceeds the treasury plus external liquidity $(T+l)$. Second, a *solvency failure* at date 2: even if the protocol survives the withdrawal episode, it must honor

the promised return $1 + r$ per token to the remaining stakers. The protocol's total resources at date 2 are the treasury T plus the return on staked assets $(I - T)(\theta + 1)$. Solvency requires

$$(I - T)(\theta + 1) + T \geq I(1 + r), \tag{3}$$

which yields a solvency threshold

$$\theta^{sol} = \frac{Ir}{I - T}. \tag{4}$$

When $\theta < \theta^{sol}$, the protocol cannot honour its promised return regardless of withdrawal behaviour. We denote liquidity failure by $\delta_{liq} = f$ (withdrawal demand exceeds $T + l$ at date 1) and solvency failure by $\delta_{sol} = f$ ($\theta < \theta^{sol}$ at date 2). Conversely, $\delta_{liq} = s$ and $\delta_{sol} = s$ indicate survival of the respective channel. The two failures are distinct events: the protocol can survive the run but still be insolvent ($\delta_{liq} = s$, $\delta_{sol} = f$), or it can fail the run despite sound fundamentals ($\delta_{liq} = f$, $\delta_{sol} = s$).

Strategic stakers face four possible states. The payoff to strategic stakers from early withdrawal (denominated in ETH) from redeeming in the protocol redemption queue is

$$u_i(a_i = 1) = \begin{cases} 1 - \eta & \text{if } \delta_{liq} = s,\ \delta_{sol} = s \\ 1 - \eta & \text{if } \delta_{liq} = s,\ \delta_{sol} = f \\ \frac{T}{I} & \text{if } \delta_{liq} = f,\ \delta_{sol} = s \\ \frac{T}{I} & \text{if } \delta_{liq} = f,\ \delta_{sol} = f \end{cases} \tag{5}$$

Early withdrawers redeem before the solvency outcome is realized, so their payoff depends only on whether the protocol has sufficient liquidity at date 1. If liquidity survives ($\delta_{liq} = s$), the redemption queue yields the face value $1-\eta$; if liquidity fails ($\delta_{liq} = f$), the protocol is disbanded and pays out the treasury pro-rata.

The payoff from late withdrawal via redemption queue is:

$$u_i(a_i = 0) = \begin{cases} 1 + r & \text{if } \delta_{liq} = s,\ \delta_{sol} = s \\ 0 & \text{if } \delta_{liq} = s,\ \delta_{sol} = f \\ \frac{T}{I} - \psi & \text{if } \delta_{liq} = f,\ \delta_{sol} = s \\ \frac{T}{I} - \psi & \text{if } \delta_{liq} = f,\ \delta_{sol} = f \end{cases} \tag{6}$$

Late stakers are exposed to both channels. If both survive ($\delta_{liq} = s, \delta_{sol} = s$), they receive the full promised return $1 + r$. If the protocol survives the run but is insolvent ($\delta_{liq} = s, \delta_{sol} = f$), the protocol cannot honor its obligations and late stakers receive 0. If the protocol fails the run ($\delta_{liq} = f$), it is disbanded and late stakers receive $T/I - \psi$ regardless of solvency, where ψ reflects disbanding costs.

Combining the payoffs across the four states, the difference in payoffs between late and early withdrawal by the strategic depositors is:

$$\begin{aligned}\pi &= u_i(a_i = 0) - u_i(a_i = 1) \\ &= \begin{cases} r + \eta & \text{if } \delta_{liq} = s,\ \delta_{sol} = s \\ -(1-\eta) & \text{if } \delta_{liq} = s,\ \delta_{sol} = f \\ -\psi & \text{if } \delta_{liq} = f \end{cases} \end{aligned} \tag{7}$$

Note that in the state $(\delta_{liq} = s, \delta_{sol} = f)$, the protocol survives the run but cannot pay at maturity. In this state, late stakers lose their full investment $(1-\eta)$ relative to early withdrawers who exited at face value. This creates a new motive for strategic withdrawal that is absent when solvency is guaranteed.

Following the global game literature, we focus on monotone equilibria defined as follows.

Definition 1 *A monotone equilibrium is an equilibrium where a strategic token depositor withdraws if and only if his private signal is below a cutoff g. Or, $a(\tilde{s}_i) = 1$ if $\tilde{s}_i \leq g$ for constants g, and $a(\tilde{s}_i) = 0$ otherwise.*

Therefore, we study monotone equilibria where a strategic token depositor withdraws if and only if his private signal $s_i \leq g$. Because the protocol faces both liquidity and solvency risk, the marginal staker's indifference condition must account for all four states $(\delta_{liq}, \delta_{sol})$.

Conditional on signal s_i, the staker assesses the probability of each state. Liquidity failure occurs when $\tilde{v} \equiv \tilde{\theta} + c_{tr}\tilde{\epsilon}_l < \theta^*_{tr}$, where θ^*_{tr} is the (liquidity) failure threshold of traditional staking and $c_{tr} = \rho(I-T)\sigma_l\sigma_s/(\kappa + \sigma_s\rho(I-T))$ translates external noise into threshold uncertainty. Solvency failure occurs when $\tilde{\theta} < \theta^{sol}$. Since $\tilde{v}$ and $\tilde{\theta}$ are jointly normal conditional on s_i with correlation $\rho_c = 1/(\sigma_v\sqrt{\tau_\theta + \tau_s})$, the state probabilities involve the bivariate normal CDF Φ_2. Here $\sigma_v = \sqrt{1/(\tau_\theta + \tau_s) + c_{tr}^2}$ is the effective posterior standard deviation.

Setting the expected payoff difference to zero at $s_i = g$ yields the indifference condition:

$$(1+r)\,\Phi_2(-a, -b; \rho_c) = \psi + (1 - \eta - \psi)(1 - \Phi(b)), \tag{8}$$

where

$$a = \sqrt{\tau_\theta + \tau_s}\left(\theta^{sol} - \frac{\tau_s g}{\tau_\theta + \tau_s}\right), \qquad b = \frac{\theta^*_{tr} - \frac{\tau_s g}{\tau_\theta + \tau_s}}{\sigma_v}, \tag{9}$$

with $\bar{\theta} = 0$ without loss of generality.

This leads to the following Proposition.

Proposition 1 *Suppose that $\rho(I-T) > \kappa\sigma_s\tau_\theta$. Then there exists a unique equilibrium in monotone strategy characterized by a signal threshold g and a failure threshold θ^*_{tr} that jointly satisfy:*

(i) The liquidity condition *(relating θ^*_{tr} to g):*

$$\theta^*_{tr} = \frac{g + \frac{\sigma_s}{\kappa}\left[\lambda - T - \rho(I-T)\right]}{1 + \frac{\sigma_s\rho(I-T)}{\kappa}}. \tag{10}$$

(ii) The four-state indifference condition *(8), which involves the bivariate normal CDF Φ_2 because liquidity failure depends on both $\tilde{\theta}$ and $\tilde{\epsilon}_l$ while solvency failure depends only on $\tilde{\theta}$.*

The protocol fails due to illiquidity when $\tilde{\theta} + c_{tr}\tilde{\epsilon}_l < \theta^*_{tr}$, and fails due to insolvency when $\tilde{\theta} < \theta^{sol} = Ir/(I-T)$. The two failure channels are distinct and can occur independently.

Equation (10) is one equation in two unknowns (g, θ^*_{tr}); it traces out the set of (g, θ^*_{tr}) pairs that are mechanically consistent with the protocol's liquidity balance. The indifference condition (8) provides the second equation, which pins down the marginal staker's behavior given the joint distribution of $(\tilde{\theta}, \tilde{v})$. The equilibrium is the unique (g, θ^*_{tr}) that satisfies both. In general, the indifference condition involves the bivariate normal CDF Φ_2 because liquidity failure ($\tilde{v} < \theta^*_{tr}$) and solvency failure ($\tilde{\theta} < \theta^{sol}$) are correlated but distinct events, so no closed-form solution for θ^*_{tr} exists.

Closed form when solvency is non-binding. When the solvency threshold is non-binding ($\theta^{sol} \ll \theta^*_{tr}$), the parameter $a \to -\infty$ and $\Phi_2(-a, -b; \rho_c) \to \Phi(-b)$. The indifference condition (8) then reduces to $\Phi(b) = (r+\eta)/(r+\eta+\psi)$, recovering the standard global games formula. Substituting this back into (10) eliminates g and yields the closed-form solution

$$\theta^*_{tr}\big|_{\theta^{sol}\ \text{n.b.}} = \frac{(I-T)(1-\rho) - \kappa\left[1 + \sigma_s(\tau_\theta + \tau_s)\sigma_v \Phi^{-1}\left(\frac{r+\eta}{r+\eta+\psi}\right)\right]}{\rho(I-T) - \kappa\sigma_s\tau_\theta}, \tag{11}$$

where $\sigma_v = \sqrt{1/(\tau_\theta + \tau_s) + c_{tr}^2}$ is the effective posterior standard deviation of the composite variable $\tilde{v} = \tilde{\theta} + c_{tr}\tilde{\epsilon}_l$ defined above, with $c_{tr} = \rho(I-T)\sigma_l\sigma_s/(\kappa + \sigma_s\rho(I-T))$. Equation (11) is not an alternative to (10) but a special-case solution that holds only when the solvency channel is inactive.

When the solvency threshold is binding ($\theta^{sol} > \theta^*_{tr}$), the state ($\delta_{liq} = s, \delta_{sol} = f$) has positive probability, and the risk of insolvency despite liquidity survival raises the incentive for strategic withdrawal. This pushes g upward and, through the liquidity condition, raises θ^*_{tr} as well. The equilibrium failure region expands on both dimensions: the solvency threshold directly adds a failure zone for $\theta \in [\theta^*_{tr}, \theta^{sol})$, and the anticipation of insolvency amplifies the liquidity run.

The condition $\rho(I-T) > \kappa\sigma_s\tau_\theta$ ensures an interior liquidity threshold. The left side, $\rho(I-T)$, measures how much additional external liquidity the protocol can raise per unit improvement in θ: better fundamentals increase the value of the staked assets $(I-T)(\theta+1)$, and the external provider lends a fraction ρ against that value. The right side, $\kappa\sigma_s\tau_\theta$, measures how much additional aggregate withdrawal pressure a unit improvement in θ generates through the common prior. When the prior is precise (high τ_θ), stakers weight it heavily relative to their private signal, so an increase in θ shifts all posteriors in the same direction, creating correlated withdrawal responses. The condition requires that the protocol's marginal borrowing capacity exceeds this collective prior-driven withdrawal sensitivity. When the condition fails, an improvement in fundamentals paradoxically worsens the protocol's net position since withdrawals grow faster than resources, and the failure threshold is not well-defined. The condition is naturally satisfied when private signals are sharp (σ_s small, so stakers anchor on their own information rather than

the prior), external finance is liquid (ρ large), or the prior is diffuse (τ_θ small). This resonates with the central insight of the global games literature: sufficiently precise private information is needed to break common-knowledge-driven coordination and sustain a unique equilibrium.

Figure 4 illustrates how the equilibrium is determined by the intersection of the liquidity-condition function $g_{\text{liq}}(\theta^*_{tr})$ and the indifference function $g_{\text{indiff}}(\theta^*_{tr})$.

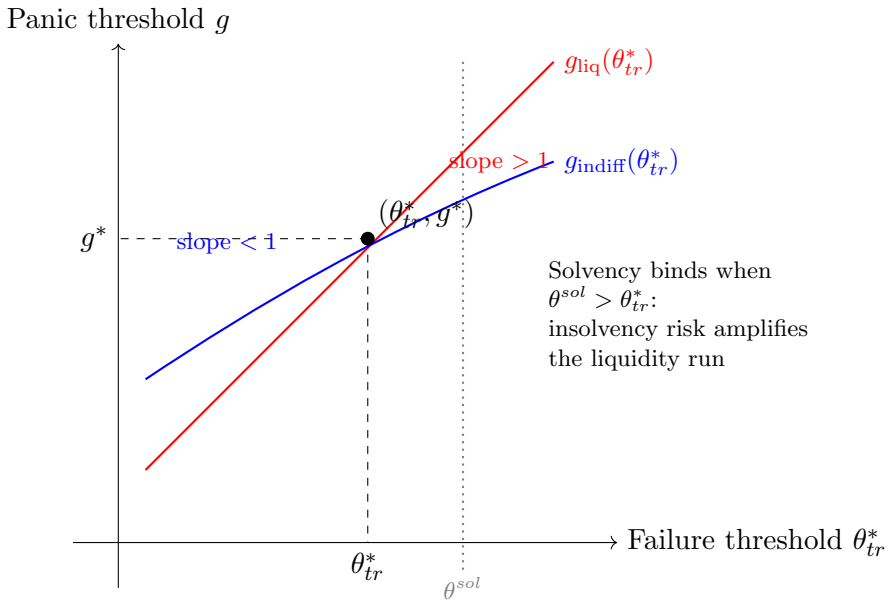

Figure 4. Equilibrium determination in traditional staking. The linear function $g_{\text{liq}}(\theta^*_{tr})$ (red) inverts the liquidity-balance condition: a higher θ^*_{tr} requires a higher signal cutoff, with slope $1+\sigma_s\rho(I-T)/\kappa > 1$. The function $g_{\text{indiff}}(\theta^*_{tr})$ (blue) is the signal threshold at which the marginal staker is indifferent, accounting for all four payoff states; it is increasing (a more fragile protocol requires a better signal to stay) but with slope less than one under the interiority condition $\rho(I-T) > \kappa\sigma_s\tau_\theta$. Their unique intersection determines the equilibrium (θ^*_{tr}, g^*). The vertical dotted line marks the solvency threshold θ^{sol}; when it lies to the right of θ^*_{tr}, the insolvency state $(\delta_{liq} = s, \delta_{sol} = f)$ has positive probability, shifting g_{indiff} upward and amplifying the run.

Next, we study the equilibrium properties of the traditional staking environment. First, we find that a larger treasury might not lower the probability of protocol failure. Its impact depends on the liquidity of the external market.

Corollary 1 *Define*

$$\bar{\rho} = \frac{\sigma_s\tau_\theta}{\Psi + \sigma_s\tau_\theta}. \tag{12}$$

When $\rho < \bar{\rho}$, $\frac{\partial\theta^*_{tr}}{\partial T} > 0$. *When* $\rho > \bar{\rho}$, $\frac{\partial\theta^*_{tr}}{\partial T} < 0$.

When the external funding market is illiquid and imposes large haircut (that is, $\rho < \bar{\rho}$), increasing the treasury destabilizes the protocol by increasing the failure threshold (that is, $\frac{\partial\theta^*_{tr}}{\partial T} > 0$). This is because by diverting cash from investing in the fundamentals, due to illiquidity of the external

market, little external financing can be raised to prevent the run. The loss of investment returns outweighs the safety benefit of cash in treasury.

Alternatively, when the external funding market is highly liquid and imposes lower haircut (that is, $\rho > \bar{\rho}$), increasing the treasury stabilizes the protocol by decreases the failure threshold (that is, $\frac{\partial \theta^*_{tr}}{\partial T} < 0$). The safety of holding cash (which suffers no liquidation haircut) outweighs the opportunity cost of lost investment returns.

Second, we find that the staking reward has an ambiguous effect on the failure threshold due to the interaction of the solvency and liquidity channels.

Corollary 2 *The effect of the staking reward r on the equilibrium failure threshold θ^*_{tr} is non-monotonic. The liquidity threshold θ^*_{tr} is decreasing in r, while the solvency threshold $\theta^{sol} = Ir/(I-T)$ is increasing in r. Define $r^\dagger$ as the reward level at which $\theta^*_{tr}(r^\dagger) = \theta^{sol}(r^\dagger)$. Then:*

- *When $r < r^\dagger$, the liquidity channel dominates ($\theta^*_{tr} > \theta^{sol}$), and $\frac{d\theta^*_{tr}}{dr} < 0$: a higher reward stabilizes the protocol.*
- *When $r > r^\dagger$, the solvency channel dominates ($\theta^{sol} > \theta^*_{tr}$), and $\frac{d\theta^*_{tr}}{dr} = \frac{I}{I-T} > 0$: a higher reward destabilizes the protocol.*

For low staking rewards, the benefit of a higher r is that it makes holding the token more attractive relative to withdrawing, reducing run pressure. However, once the reward exceeds the critical level $r^\dagger$, further increases in r raise the bar for solvency without additional liquidity benefits. A protocol promising returns it cannot deliver is more fragile, not less – rational stakers anticipate the solvency shortfall and incorporate it into their withdrawal decisions.

We next examine how the noise in the external liquidity provider's signal affects the stability of the traditional staking protocol.

Corollary 3 *Suppose the liquidity channel is binding ($\theta^*_{tr} > \theta^{sol}$) and $r + \eta > \psi$. Then $\frac{d\theta^*_{tr}}{d\sigma_l} < 0$: greater noise in the external liquidity provider's signal lowers the failure threshold and stabilizes the protocol. When the solvency channel binds ($\theta^{sol} > \theta^*_{tr}$), σ_l has no effect on θ^*_{tr}.*

The noise σ_l enters the equilibrium through the effective variance $\sigma_v^2 = 1/(\tau_\theta + \tau_s) + c_{tr}^2$. A higher σ_l raises c_{tr} and hence σ_v, making the failure event less predictable from the staker's perspective. In the global games framework, this reduces the precision with which stakers can coordinate their withdrawal decisions. With less certainty about whether the protocol will actually fail, since a good or bad draw of $\tilde{\epsilon}_l$ can respectively save or doom the protocol, stakers are less responsive to marginal changes in their private signal, weakening the strategic complementarity that drives runs. The signal threshold g falls, withdrawal demand decreases, and the failure threshold θ^*_{tr} declines.

This result echoes the broader finding in Proposition 1 that noise dampens coordination: just as sufficiently noisy private signals are required for equilibrium uniqueness, noisier external

liquidity assessments reduce the effectiveness of run coordination. The condition $r + \eta > \psi$ ensures that the benefit of staying exceeds the cost of late failure, so that stakers are willing to tolerate greater uncertainty about the failure event. When the solvency threshold binds, σ_l has no effect because solvency depends on the true fundamental θ, not on the liquidity provider's noisy assessment.

We next examine how the depth of the external funding market, parameterized by ρ, affects the failure threshold.

Corollary 4 *Suppose the liquidity channel is binding ($\theta^*_{tr} > \theta^{sol}$) and the equilibrium failure threshold satisfies $\theta^*_{tr} > -1$. Then $\frac{d\theta^*_{tr}}{d\rho} < 0$: a more liquid external funding market lowers the failure threshold and stabilizes the protocol.*

A higher ρ raises the marginal external liquidity per unit improvement in fundamentals, $\rho(I-T)$, which appears on both sides of the protocol's balance sheet. When the external provider observes the fundamental perfectly ($\sigma_l = 0$, so that σ_v does not depend on ρ), differentiating the closed form (11) gives

$$\frac{\partial \theta^*_{tr}}{\partial \rho} = -\frac{(I-T)(1+\theta^*_{tr})}{\rho(I-T) - \kappa\sigma_s\tau_\theta} < 0,$$

since the denominator is positive under the interiority condition of Proposition 1 and $1+\theta^*_{tr} > 0$. Intuitively, a deeper external funding market makes the protocol's invested base more loanable: each unit improvement in θ is more leveraged into liquidity, both shrinking the queue's reach for given fundamentals and raising the slope at which the protocol's resources respond. Combined with Corollary 1, this clarifies the role of ρ: when external markets are very illiquid ($\rho < \bar{\rho}$), the protocol must rely on idle treasury and increasing T is stabilizing; when external markets are deep ($\rho > \bar{\rho}$), productive assets dominate and the protocol benefits from a smaller treasury. In both regions, deeper external markets (higher ρ) work to lower the failure threshold.

Finally, we consider the impact of the mass of non-strategic, liquidity-shocked stakers λ that join the redemption queue mechanically.

Corollary 5 *The effect of λ on the equilibrium failure threshold θ^*_{tr} is governed by the same condition as the treasury effect. The sign of $\frac{d\theta^*_{tr}}{d\lambda}$ depends on ρ relative to $\bar{\rho}$:*

- *When $\rho < \bar{\rho}$, $\frac{d\theta^*_{tr}}{d\lambda} < 0$: more forced sellers stabilize the protocol.*
- *When $\rho > \bar{\rho}$, $\frac{d\theta^*_{tr}}{d\lambda} > 0$: more forced sellers destabilize the protocol.*

Since $I = \kappa + \lambda$ is fixed, an increase in λ mechanically reduces the strategic mass: $d\kappa/d\lambda = -1$. The total derivative of the closed-form threshold (11) with respect to λ is therefore $d\theta^*_{tr}/d\lambda = -\partial\theta^*_{tr}/\partial\kappa$, which gives

$$\frac{d\theta^*_{tr}}{d\lambda} = \frac{(I-T)[\rho\Psi - (1-\rho)\sigma_s\tau_\theta]}{[\rho(I-T) - \kappa\sigma_s\tau_\theta]^2}.$$

The numerator is exactly the negative of the numerator $\mathcal{N}$ that determines the sign of $\partial\theta^*_{tr}/\partial T$ in Corollary 1. The economic intuition mirrors the treasury result: when external markets are illiquid ($\rho < \bar{\rho}$), a larger λ shrinks the strategic mass κ, weakening the coordination engine of the run enough to overpower the mechanical increase in queue demand. When external markets are liquid ($\rho > \bar{\rho}$), the direct queue-congestion effect dominates. This result has a sharp counterpart in liquid staking, where forced sellers are routed to the secondary market and barely affect the queue.

5 The Liquid Staking

We assume the existence of a secondary market for staked tokens. To obtain liquidity, stakers now face a choice: they can sell their tokens in the market for immediate execution, or they can submit a redemption request to the protocol and wait for a response.

Importantly, the market for staked tokens reveals information about the protocol's fundamentals. This information affects the amount of external liquidity available to the protocol and helps coordinate strategic stakers' redemption decisions.

To model this informational impact on the protocol's external liquidity, we assume that external liquidity providers base their funding decisions on their own noisy signals regarding the fundamentals, as well as on the traded token price, P. This is specified as:

$$l = (I - T)\left(\rho \underbrace{(\theta + \sigma_l \tilde{\epsilon}_l + 1)}_{\text{Liquidity provider's signal}} + \beta \underbrace{(y(P) + 1)}_{\text{Market price signal}}\right), \tag{13}$$

where $y(P)$ is the sufficient statistic for θ extracted from the staked token price, $0 < \beta < 1$ is the weight placed on the market price signal per unit of invested capital, and $\tilde{\epsilon}_l$ is the same external liquidity noise as in the traditional staking case. The impact of both signals on external liquidity is proportional to the protocol's invested base, $(I - T)$.

To model how the market price serves as an additional signal for strategic token holders coordinating a run, we assume that token holders requiring urgent liquidity, for informational or non-informational reasons, prefer to utilize the secondary market. This is motivated by two empirical observations. First, the redemption queue imposes a significant time delay, whereas the secondary market offers immediate settlement for sell orders.[3] Second, the staking service has an alternative means to fulfill withdrawals beyond running down its treasury: it can sell the staked tokens from the redemption queue in the market for unstaked ETH at the market price. Thus, any redemption demand can be indirectly routed to the market and executed with a delay. Rather than waiting for the protocol to do this, early withdrawers can execute the trade themselves immediately.

Consequently, two types of stakers naturally utilize the market exclusively for withdrawals: non-strategic, liquidity-shocked token holders (λ), and strategic stakers who receive exceptionally weak signals about the protocol's fundamentals and wish to exit their positions immediately.

[3] See empirical evidence.

All strategic token holders optimize their exit strategies based on their private signals. Those who receive signals above the "exceptionally weak" threshold opt to wait and observe the market price and use it as an additional signal before making a final decision on whether to exit their staked position. This latter group must exit via the redemption queue; because market makers face a downward-sloping demand curve, the market impact of additional selling would drive the price too low for these late sellers.

In this section, we study how strategic stakers make this set of decisions optimally. We first delineate their payoff from this set of decisions. As in the traditional staking case, stakers face four possible states $(\delta_{liq}, \delta_{sol})$. The payoff to strategic stakers from early withdrawal (denominated in ETH) by selling in the market is p, regardless of the protocol's state and the seller exits immediately. The payoff to strategic stakers from early withdrawal (denominated in ETH) by redeeming in the protocol redemption queue is

$$u_i(a_i = 1) = \begin{cases} 1 - \eta & \text{if } \delta_{liq} = s,\ \delta_{sol} = s \\ 1 - \eta & \text{if } \delta_{liq} = s,\ \delta_{sol} = f \\ \frac{T}{I} & \text{if } \delta_{liq} = f,\ \delta_{sol} = s \\ \frac{T}{I} & \text{if } \delta_{liq} = f,\ \delta_{sol} = f \end{cases} \tag{14}$$

As in the traditional case, queue redeemers exit before the solvency outcome is realized, so their payoff depends only on δ_{liq}.

The payoff from late withdrawal via the redemption queue is:

$$u_i(a_i = 0) = \begin{cases} 1 + r & \text{if } \delta_{liq} = s,\ \delta_{sol} = s \\ 0 & \text{if } \delta_{liq} = s,\ \delta_{sol} = f \\ \frac{T}{I} - \psi & \text{if } \delta_{liq} = f,\ \delta_{sol} = s \\ \frac{T}{I} - \psi & \text{if } \delta_{liq} = f,\ \delta_{sol} = f \end{cases} \tag{15}$$

Late stakers are exposed to both channels. If both survive, they receive the full promised return $1 + r$. If the protocol survives the run but is insolvent $(\delta_{liq} = s, \delta_{sol} = f)$, late stakers receive 0. If the protocol fails the run $(\delta_{liq} = f)$, it is disbanded and late stakers receive $T/I - \psi$ regardless of solvency.

The solvency threshold $\theta^{sol} = Ir/(I - T)$ from (4) applies equally here. For the protocol to fully survive, both conditions must hold: redemption demand at date 1 must be less than $T + l$, and $\theta \geq \theta^{sol}$ at date 2.

5.1 The Market for Staked Tokens

We assume a group of liquidity traders acts as market makers for the staked tokens. Each provider receives a private signal about the protocol's fundamental, $\tilde{s}_m = \tilde{\theta} + \sigma_m \tilde{\epsilon}_m$, where $\tilde{\epsilon}_m$ is an independent standard normal random variable and $\sigma_m > 0$ scales the signal noise. Informed by this signal, they submit a downward-sloping demand curve to accommodate selling requests:

$$D(P) = b\tilde{s}_m - \Phi^{-1}\left(\frac{\gamma P}{1 + r}\right). \tag{16}$$

Here, $D(P)$ is the token-denominated demand and $b > 0$ measures its sensitivity to the private signal. The parameter γ represents price elasticity, determined by the market makers' risk appetite and alternative uses for the staked tokens (e.g., as collateral). A higher γ decreases their willingness to absorb selling pressure, making it observationally equivalent to a higher risk premium. Finally, setting $b = 0$ recovers the benchmark of uninformed market makers with a constant demand intercept.[4]

To determine the price in the liquid staking market, we characterize the supply of tokens that are sold there as opposed to redeemed by joining the queue. The market clearing condition is hence,

$$D(P) + \sigma_\xi \tilde{\xi} = \text{Aggregate Sales} = \underbrace{(\kappa \text{Sales from Strategic Stakers} + \lambda)}_{\text{Aggregate Sales}}, \tag{17}$$

where the right side of the equation is the aggregate selling from the staked token holders. The price will not be fully revealing because of two sources of noise: the demand shock $\sigma_\xi \tilde{\xi}$, which captures market trade independent of the staking service, and the noise $\sigma_m \tilde{\epsilon}_m$ in the market makers' signal. Both $\tilde{\xi}$ and $\tilde{\epsilon}_m$ are independent standard normal random variables. As before, we denote $\tau_\xi = 1/\sigma_\xi^2$ and $\tau_m = 1/\sigma_m^2$.

5.2 Learning from the Market Price and Redemption Queue

At date 1, strategic stakers need to decide, based on their private signal, to withdraw or wait a bit more to observe an additional market price signal. If they choose withdraw immediately, they sell their token in the market. They can delay making this decision by waiting a bit more to observe the market price. By then, they can choose a slow withdraw route by submitting a redemption request via the redemption queue.

As we have noted earlier, the payoffs to their decision induce a cost to delay in the withdrawal decision. Strategic stakers thus rationally trade off the lower cost to selling in the market against the option value of waiting in the queue and decide by observing the market price. We look for Bayes Nash equilibria for this game.

Strategic stakers who wait in the queue observe two noisy signals, s_i and p. Let $z(s_i, p)$ denote a sufficient statistic of θ for (s_i, p). We look for equilibria in which agents choose monotone threshold strategies, such that they choose to sell their staked token in the market if $s_i < g$; otherwise, they withdraw via the redemption queue if $x(s_i, p) < x^*$ and stay with the protocol otherwise.

Before proceeding to analyze such threshold equilibria, we first characterize the information held by stakers who choose to observe the price, given the assumed strategies for the early sellers in the market.

[4]We utilize this reduced-form specification for the demand curve to keep the inference problem extracted from the market price tractable.

Since those strategic token holders who choose to withdraw their tokens by selling in the market whenever their private signal about the fundamental value is sufficiently low, or $\tilde{s}_i \leq g$, or equivalently, $\sigma_s \tilde{\epsilon}_i \leq g - \tilde{\theta}$. Thus, strategic aggregate selling is:

$$\kappa \Phi^{-1}\left(A(s_i)\right) = \kappa \frac{g - \tilde{\theta}}{\sigma_s}. \tag{18}$$

We note that the Φ^{-1} formulation is a reduced-form approximation that keeps the price inference framework Gaussian and linear; the true bounded mass κA is replaced by the unbounded $\kappa \Phi^{-1}(A)$, which can in principle be negative or exceed κ. The approximation is standard in the global games literature and accurate when private signals are sufficiently precise.[5]

The total withdrawal selling amount comprises the strategic investors in Equation (18) plus the stakers who had access to a better investment opportunity, which we denoted by λ. The queue is thus:

$$\tilde{Q} = \left(\kappa \frac{g - \tilde{\theta}}{\sigma_s} + \lambda\right).$$

Combining the market selling and the market buying (17), we obtain the market price via the market clearing condition as follows:

$$\underbrace{b\tilde{s}_m - \Phi^{-1}\left(\frac{\gamma P}{1+r}\right) + \sigma_\xi \tilde{\xi}}_{\text{Noisy Token Demand}} = \underbrace{\left(\kappa \frac{g - \tilde{\theta}}{\sigma_s} + \lambda\right)}_{\text{Redemption Selling}},$$

$$P = \frac{1+r}{\gamma} \Phi\left(\left(b + \frac{\kappa}{\sigma_s}\right)\tilde{\theta} + b\sigma_m \tilde{\epsilon}_m + \sigma_\xi \tilde{\xi} - \frac{\kappa g}{\sigma_s} - \lambda\right). \tag{19}$$

Note that the token price is between $(0, 1+r]$ since $\gamma \geq 1$.

Given that other strategic stakers can condition on the price in the token market and draw inferences about the fundamental value from that price, we can reorganize (19) and collect the terms that are in this group of strategic stakers' information set to the left side of the equation and the rest to the right side. Recalling that $I = \kappa + \lambda$, we obtain:

$$y(P) = \underbrace{\frac{1}{b + \kappa/\sigma_s}\left(\Phi^{-1}\left(\frac{\gamma P}{1+r}\right) + \frac{\kappa g}{\sigma_s} + \lambda\right)}_{\text{Price Information}} = \tilde{\theta} + \frac{b\sigma_m \tilde{\epsilon}_m + \sigma_\xi \tilde{\xi}}{b + \kappa/\sigma_s}. \tag{20}$$

Intuitively, $y(P)$ is distributed normally and is a sufficient statistic for the information in the market price, P. The price does not fully reveal the fundamental $\tilde{\theta}$ because of two sources of noise: the market makers' signal noise $\tilde{\epsilon}_m$ and the demand shock $\tilde{\xi}$. We characterize the variance of $y(P)$ given $\tilde{\theta}$ as

$$\sigma_p^2 = \frac{b^2 \sigma_m^2 + \sigma_\xi^2}{\left(b + \kappa/\sigma_s\right)^2}.$$

[5] See Goldstein and Pauzner (2005) for the analogous approximation in their bank-run model.

We denote the precision of $y(P)$ as a signal for $\tilde{\theta}$ as:

$$\tau_p = 1/\sigma_p^2 = \frac{(b + \kappa/\sigma_s)^2}{b^2\sigma_m^2 + \sigma_\xi^2}. \tag{21}$$

When $b = 0$ (uninformed market makers), the precision reduces to $\tau_p = \kappa^2\tau_s\tau_\xi$, recovering the benchmark. A higher b increases the informativeness of the price (raises τ_p) provided that the market makers' signal is not too noisy relative to demand noise (σ_m not too large). Since the rest of strategic token holders observe $y(P)$, they will use it to update beliefs about $\tilde{\theta}$. Let $\nu(\tilde{\theta}|s_i, P)$ denote their updated belief on the value of the staking intermediary given the market price.

Upon observing $s_i, y\,(P)$, the rest of strategic token holders believe that $\tilde{\theta}$ is distributed normally with mean $\left(\tau_\theta\bar{\theta} + \tau_s s_i + \tau_p y\,(P)\right) / \left(\tau_\theta + \tau_s + \tau_p\right)$ and variance $1/\left(\tau_\theta + \tau_s + \tau_p\right)$. Therefore, define:

$$z = \frac{\tau_\theta\bar{\theta} + \tau_s s_i + \tau_p y\,(P)}{\tau_\theta + \tau_s + \tau_p}. \tag{22}$$

Hence, a sufficient statistic for information about θ held by this group of strategic stakers is:

$$\theta|s_i, p \equiv \theta|z \sim N\left(\frac{\tau_\theta\bar{\theta} + \tau_s s_i + \tau_p y\,(P)}{\tau_\theta + \tau_s + \tau_p}, \frac{1}{\tau_\theta + \tau_s + \tau_p}\right). \tag{23}$$

For notational simplicity, we use $\hat{\tau} = \tau_\theta + \tau_s + \tau_p$.

5.3 Redemption Queue and the Survival of Liquidity Staking

When strategic stakers follow monotone strategies outlined above, at any θ, the redemption length is determined by $\kappa Pr\,(s > g, x < x^*)$. Conditional on $s > g$, the threshold signal x for joining the queue involves the expected payoff $(1 + \theta)$ scaled by liquidity factors:

$$x = (I - T)\left[\rho(z + 1) + \beta(y(P) + 1)\right].$$

Expanding this using the components of z and $y(P)$:

$$x = (I - T)\left[\rho\left(\frac{\tau_\theta\bar{\theta} + \tau_s(\tilde{\theta} + \sigma_s\epsilon_i) + \tau_p\left(\tilde{\theta} + \frac{b\sigma_m\tilde{\epsilon}_m + \sigma_\xi\tilde{\xi}}{b + \kappa/\sigma_s}\right)}{\hat{\tau}} + 1\right) + \beta\left(\tilde{\theta} + \frac{b\sigma_m\tilde{\epsilon}_m + \sigma_\xi\tilde{\xi}}{b + \kappa/\sigma_s} + 1\right)\right] \tag{24}$$

We denote the random component as $\tilde{\omega}$:

$$\tilde{\omega} = (I - T)\left[\frac{\rho\tau_s\sigma_s}{\hat{\tau}}\tilde{\epsilon}_i + \left(\frac{\rho\tau_p}{\hat{\tau}} + \beta\right)\frac{b\sigma_m\tilde{\epsilon}_m + \sigma_\xi\tilde{\xi}}{b + \kappa/\sigma_s}\right], \tag{25}$$

and re-organize (24) to obtain the condition for joining the redemption queue, $\omega < M(\theta)$:

$$\omega < x^* - (I-T)\left[\frac{\rho\tau_\theta\bar{\theta}}{\hat{\tau}} + \left(\frac{\rho(\tau_s+\tau_p)}{\hat{\tau}} + \beta\right)\tilde{\theta} + (\rho+\beta)\right]. \tag{26}$$

The protocol survives the liquidity test ($\delta_{liq} = s$) if:

$$\kappa Pr\left(s > g, x < x^*\right) < T + (I-T)\left[\rho(1+\theta+\sigma_l\tilde{\epsilon}_l) + \beta(y(P)+1)\right]. \tag{27}$$

Because stakers do not observe $\tilde{\epsilon}_l$, their decision thresholds (g, x^*) are unchanged in form. However, the actual liquidity survival depends on both θ and $\tilde{\epsilon}_l$. Define the net stress function $F(\theta, \epsilon_l) = \kappa \Pr(s > g, x < x^*|\theta) - T - (I-T)[\rho(1+\theta+\sigma_l\epsilon_l) + \beta y(P)]$. The protocol fails the liquidity test when $F(\theta, \epsilon_l) > 0$. Since F is linear in ϵ_l, the liquidity failure event given θ can be written as:

$$\delta_{liq} = f \quad \Longleftrightarrow \quad \epsilon_l < \frac{F(\theta, 0)}{(I-T)\rho\sigma_l}, \tag{28}$$

where $F(\theta, 0)$ is the net stress evaluated at the expected external liquidity (i.e., $\epsilon_l = 0$).

The liquidity failure threshold θ^*_{liq} is defined implicitly by $F(\theta^*_{liq}, 0) = 0$:

$$\kappa Pr\left(s > g, x < x^*|\theta^*_{liq}\right) = T + (I-T)\left[\rho(1+\theta^*_{liq}) + \beta(y(P)+1)\right]. \tag{29}$$

The solvency channel is independent: $\delta_{sol} = f$ iff $\theta < \theta^{sol} = Ir/(I-T)$. As in the traditional case, the two failure channels are distinct and can occur independently. From (28), the conditional probability of liquidity failure given θ is:

$$\Pr(\delta_{liq} = f \mid \theta) = \Phi\left(\frac{F(\theta, 0)}{(I-T)\rho\sigma_l}\right). \tag{30}$$

When $\sigma_l = 0$, this reduces to a step function at θ^*_{liq}, recovering the benchmark.

Let $F(\theta)$ denote the "Net Stress" on the protocol:

$$F(\theta) = \kappa Pr\left(s > g, x < x^*\right) - \left(T + (I-T)\left[\rho(1+\theta) + \beta(y(P)+1)\right]\right). \tag{31}$$

Note that $F(\theta)$ retains its original form in expectation over the aggregate noise sources ($\tilde{\epsilon}_l$, $\tilde{\epsilon}_m$, $\tilde{\xi}$), since all have zero mean. The liquidity failure threshold θ^* is therefore defined by evaluating the net stress at the expected values of the aggregate market noise ($\epsilon_l = \epsilon_m = \xi = 0$), which yields a single deterministic threshold rather than a failure boundary that shifts randomly with the price realization. We show that there exists a critical θ above which the protocol survives and below which it fails when σ_s is sufficiently large.

Lemma 1 *When σ_s is sufficiently large, the function $F(\theta)$ is strictly decreasing in θ and crosses zero exactly once. Consequently, there exists a unique fundamental threshold θ^* such that the protocol survives if and only if $\theta > \theta^*$.*

The following Lemma derives the properties of function $\theta^*(g)$.

Lemma 2 *The fundamental threshold $\theta^*(g)$ is a smooth function of the signal threshold g. Furthermore, for sufficiently large noise σ_s, the function is strictly decreasing: $\frac{d\theta^*}{dg} < 0$.*

The primary economic insight from Lemma 2 is that the secondary market performs as a liquidity outlet for the liquidity staking protocol. Unlike traditional bank runs where every panicked depositor must withdraw directly from the bank, liquid staking offers an alternative exit. When stakers panic (represented by a higher signal threshold g), they prefer to sell their tokens on the open market rather than wait in the slow redemption queue. Crucially, this transfers ownership to risk-tolerant market makers without forcing the protocol to pay out immediate reserves. As more nervous agents exit via the market, the protocol's redemption queue actually shrinks. With fewer immediate liabilities to service, the fundamental threshold required for the protocol's survival (θ^*) lowers, making the system paradoxically more robust during moments of high anxiety.

Additionally, secondary market price provides information to external liquidity providers (ie., via the $y(P)$ term), whose resource stabilizes the protocol. While increased panic naturally depresses the secondary market price, it also clarifies the signal for external liquidity providers. Sophisticated external lenders understand that a price drop during high panic is driven by liquidity needs rather than necessarily poor fundamentals. Our model implies that lenders counteracts this panic component; they attribute the low price to the "market frenzy" rather than underlying asset insolvency. This allows the protocol to maintain access to external credit lines even when token prices are falling, preventing the credit freeze that typically accelerates a collapse.

Together, these mechanisms distinguish liquid staking from traditional banking models. In a standard Diamond-Dybvig bank run, panic is self-reinforcing and destructive: fear leads to withdrawals, which drains reserves, ensuring failure. In contrast, Lemma 2 demonstrates that the secondary market acts as a shock absorber. By allowing impatient agents to trade with risk-tolerant buyers, the market effectively "quarantines" the panic outside of the protocol's balance sheet. Consequently, an increase in the propensity to run unintentionally protects the protocol by diverting outflows away from the treasury and into the secondary market, strictly decreasing the likelihood of insolvency.

5.4 Characterizing Strategic Staker's Monotone Threshold Strategy

As outlined above, strategic stakers follow a monotone strategy (g, x^*) or equivalently, (g, z^*).[6]

There are two indifference equations for these stakers, both of which must account for the four states $(\delta_{liq}, \delta_{sol})$.

The "late" indifference equation governs the strategic stakers who have observed the price and decide between joining the redemption queue or holding the token till maturity. The payoff

[6]Because x is a strictly monotonic function of z, a strategy defined by a threshold on one variable is mathematically identical to a strategy defined on the other.

difference (hold minus queue) across states is: $r+\eta$ if $(\delta_{liq}=s,\delta_{sol}=s)$; $-(1-\eta)$ if $(\delta_{liq}=s,\delta_{sol}=f)$; and $-\psi$ if $\delta_{liq}=f$. Setting the expected payoff difference to zero at the queue threshold x^*:

$$\begin{aligned}&(r+\eta)\ \Pr(\delta_{liq}=s,\ \delta_{sol}=s\mid x^*)\\&-(1-\eta)\ \Pr(\delta_{liq}=s,\ \delta_{sol}=f\mid x^*)-\psi\ \Pr(\delta_{liq}=f\mid x^*)=0.\end{aligned}\tag{32}$$

The staker conditioning on x^* (equivalently z^*) has posterior $\tilde{\theta}\mid z^*\sim N(z^*,1/\hat{\tau})$. Since $\tilde{\epsilon}_l$ is independent, the conditional probability of liquidity failure given θ is $\Pr(\delta_{liq}=f\mid\theta)=\Phi(F(\theta,0)/[(I-T)\rho\sigma_l])$ from (30). The probabilities in (32) therefore involve integrating over the posterior on θ:

$$\begin{aligned}\Pr(\delta_{liq}=s,\ \delta_{sol}=s\mid z^*)&=\int_{\theta^{sol}}^{\infty}\left[1-\Phi\Big(\frac{F(\theta,0)}{(I-T)\rho\sigma_l}\Big)\right]\hat{\tau}^{1/2}\phi\Big(\hat{\tau}^{1/2}(\theta-z^*)\Big)\,d\theta,\\\Pr(\delta_{liq}=s,\ \delta_{sol}=f\mid z^*)&=\int_{-\infty}^{\theta^{sol}}\left[1-\Phi\Big(\frac{F(\theta,0)}{(I-T)\rho\sigma_l}\Big)\right]\hat{\tau}^{1/2}\phi\Big(\hat{\tau}^{1/2}(\theta-z^*)\Big)\,d\theta.\end{aligned}\tag{33}$$

When $\sigma_l=0$, $\Pr(\delta_{liq}=f\mid\theta)$ becomes the indicator $\mathbf{1}_{\theta<\theta^*_{liq}}$, and if additionally $\theta^{sol}\ll\theta^*_{liq}$ (solvency non-binding), the condition reduces to $\Pr(\theta>\theta^*_{liq}\mid x^*)(r+\eta)-\Pr(\theta<\theta^*_{liq}\mid x^*)\psi=0$.

There is also the "early" indifference condition for the strategic stakers who are deciding whether to sell to the market early. These stakers trade off the expected benefit of selling early against the retention value, which is the expected benefit of waiting and then acting optimally given the price observation.

With four payoff states, the retention value $V(g)$ for the marginal investor with signal g accounts for both the queue/hold decision and the solvency outcome. Those who wait and then hold $(x>x^*)$ receive: $1+r$ if $(\delta_{liq}=s,\delta_{sol}=s)$; 0 if $(\delta_{liq}=s,\delta_{sol}=f)$; and $T/I-\psi$ if $\delta_{liq}=f$. Those who wait and then queue $(x<x^*)$ receive: $1-\eta$ if $\delta_{liq}=s$; and T/I if $\delta_{liq}=f$. Thus, the marginal investor must satisfy:

$$\begin{aligned}E\left[p|g\right]\quad&=(1+r)\ \Pr(\delta_{liq}=s,\ \delta_{sol}=s,\ x>x^*\mid g)\\&+0\cdot\Pr(\delta_{liq}=s,\ \delta_{sol}=f,\ x>x^*\mid g)+\left(\frac{T}{I}-\psi\right)\Pr(\delta_{liq}=f,\ x>x^*\mid g)\\&+(1-\eta)\ \Pr(\delta_{liq}=s,\ x<x^*\mid g)+\frac{T}{I}\ \Pr(\delta_{liq}=f,\ x<x^*\mid g).\end{aligned}\tag{34}$$

The probabilities in (34) involve the joint distribution of $\tilde{\theta}$, $\tilde{v}$, and the queue threshold x, since the liquidity-failure event, the solvency event, and the queue/hold decision are all correlated through the common price noise. In the general case, these probabilities require bivariate (or trivariate) normal CDFs, exactly paralleling the use of Φ_2 in the traditional staking indifference (8). When $\sigma_l=0$ and $\theta^{sol}\ll\theta^*_{liq}$, the solvency state collapses ($\delta_{sol}=s$ always), the liquidity failure becomes a deterministic threshold on θ, and the early indifference reduces to a simpler form.

We show in the following proposition that for a given θ^*, the cutoff g exists and is unique. Moreover, the symmetric equilibrium g^* (in which all stakers use the same threshold) is unique:

the aggregate indifference function $\mathcal{G}(g) \equiv E[p|s = g; g] - V(g; g)$ is strictly decreasing in g because the expected selling proceeds $E[p|g]$ fall as aggregate panic increases (more selling depresses the price) while the retention value $V(g)$ rises (a higher g makes the posterior on θ more optimistic for a staker at the margin).

Proposition 2 *Assume the interest rate satisfies* $1 + r > \frac{\gamma T}{I\Phi(b\bar{\theta}-I)}$ *and market friction satisfies* $\gamma > \Phi(b\bar{\theta} - \lambda)$. *Then, for any fixed belief* θ^*, *there exists a unique signal threshold* g *such that the marginal staker is indifferent between selling and retaining the token.*

The first condition, $1 + r > \frac{\gamma T}{I\Phi(b\bar{\theta}-I)}$, represents a constraint that preserves the secondary market's function during stress. It ensures that the fundamental value of the token remains sufficiently high relative to the protocol's liquidation value. Economically, this condition guarantees that even in a worst-case panic scenario, the market price of the token offers a better payoff than the pro-rata liquidation value T/I a staker would receive from a failed protocol's treasury. If this condition were violated, the market price would fall below the liquidation value of the treasury, destroying the incentive to trade. Rational agents would abandon the secondary market entirely and flood the redemption queue to claim their share of the reserves, thereby negating the liquidity outlet mechanism served by the secondary market that distinguishes liquid staking from traditional banking.

The second condition, $\gamma > \Phi(b\bar{\theta} - \lambda)$, establishes a necessary liquidity premium to prevent the market from unravelling during normal times. This inequality requires that market friction (γ) be sufficiently large relative to market depth. If the market were perfectly liquid and frictionless (a very small γ), the current market price would arbitrage perfectly with the future reward ($p \approx 1+r$). In such a scenario, risk-averse agents would have no incentive to hold the risky asset to maturity; they would prefer to sell immediately to lock in the profit risk-free. By enforcing a minimum level of illiquidity, our model ensures that the market price trades at a discount to the future reward. This discount creates an opportunity cost for selling, incentivizing agents with positive private signals to hold the token rather than exit, thus sustaining the equilibrium where agents sort themselves based on information.

5.5 Equilibrium: Existence and Uniqueness

We now combine the analysis of the staker's decision (Proposition 2) and the protocol's survival (Lemmas 1, and 2) to prove the existence and uniqueness of the joint equilibrium pair (θ^*, x^*, g^*).

Note that in the proof of the existence and uniqueness of the equilibrium, the variable x^* (and by extension z^*) is substituted out because it functions as a dependent variable. The "late indifference" condition (i.e., eq. (32)) ensures that for any hypothesized failure threshold θ^*, there is exactly one belief x^* that renders a staker indifferent between queuing and holding. Consequently, x^* is strictly determined by, θ^*. This relationship allows the proof to reduce the dimensionality of the system from three variables (g, x^*, θ^*) to just two interacting forces: the Staker's reaction $g(\theta^*)$ and the Protocol's reaction $\theta^*(g)$. By substituting $x^*(\theta^*)$ out, the analysis simplifies to finding the intersection of these two reaction functions in the (θ^*, g) plane,

making the explicit simultaneous solution for x^* redundant for establishing uniqueness. We state the existence and the uniqueness of the monotone threshold equilibrium in the following theorem and the detailed proof is the Appendix.

Theorem 1 *For sufficiently large private signal noise σ_s, there exists a unique equilibrium characterized by the pair of thresholds (θ^*, g^*).*

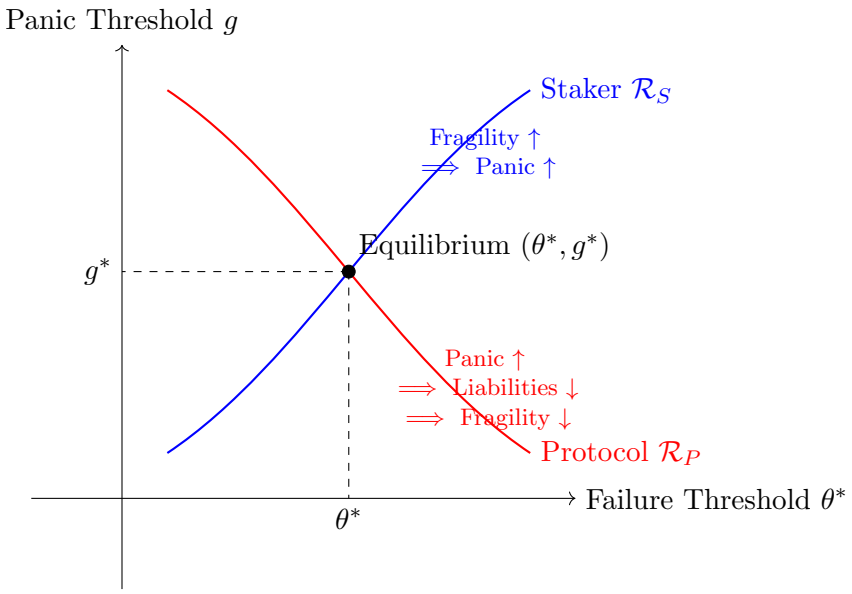

Figure 5. Uniqueness of Equilibrium. The Staker's reaction function (blue) is upward sloping: as the protocol becomes riskier (higher θ^*), agents require a higher signal g to stay. The Protocol's reaction function (red) is downward sloping: as agents panic more (higher g), they exit via the secondary market, reducing the protocol's redemption queue liability and lowering the failure threshold θ^*.

Figure 5 illustrates the determination of a unique equilibrium through the interaction of two opposing economic forces: the stakers' strategic panic and the protocol's mechanical survival constraint. The upward-sloping *Staker Reaction Curve* represents the standard "fear" mechanism found in run models: if stakers perceive the protocol as more fragile (a higher failure threshold θ^*), they demand a higher signal of quality (g) to remain invested, leading to increased aggregate panic. Counteracting this is the downward-sloping *Protocol Reaction Curve*, which captures the liquidity outlet effect served by the secondary market unique to liquid staking. As the panic threshold (g) rises and more agents rush to exit, they primarily sell tokens on the secondary market rather than crowding the redemption queue; this diversion of outflows reduces the protocol's immediate liability, thereby lowering the fundamental threshold (θ^*) required for solvency. The unique intersection of these curves defines the equilibrium, demonstrating that the secondary market stabilizes the system by ensuring that increased panic forces a reduction in the protocol's structural fragility.

In standard Global Game frameworks, such as Morris and Shin (2003), the uniqueness of equilibrium relies on private signals being sufficiently precise (small σ). The fundamental friction

in these bank run models is coordination failure: if information is public or signals are too noisy, agents prioritize coordinating with one another over evaluating fundamentals, leading to self-fulfilling multiple equilibria (i.e., "I run because I believe you will run"). By introducing precise private information, the model compels agents to anchor their decisions on the fundamental state rather than on the actions of peers. This breaks the common knowledge required for a run, collapsing the multiplicity into a single unique equilibrium. Thus, in the canonical literature, signal precision is a prerequisite for uniqueness.

In sharp contrast, our liquid staking model requires private signals to be sufficiently noisy (large σ) to guarantee uniqueness. This reversal arises because the stabilizing mechanism is structural rather than coordinative. In our model, increased staker panic (g) paradoxically stabilizes the protocol by diverting liabilities to the secondary market. However, this mechanism relies on external lenders maintaining credit lines despite falling token prices. If private signals were too precise (small σ), the secondary market price would efficiently aggregate private information, crashing violently during a panic. External lenders, observing this sharp price drop, would infer insolvency with certainty and immediately retract funding, overwhelming the stabilizing effect of the secondary market. Therefore, the model requires noisy signals to dampen the price impact, ensuring that lenders attribute price drops to temporary market noise rather than fundamental insolvency. This result aligns with Angeletos and Werning (2006) and Ozdenoren and Yuan (2008), who demonstrate that when agents coordinate via an endogenous public signal (such as price), private signals must be sufficiently noisy to sustain a unique equilibrium.

6 Equilibrium Properties: Traditional and Liquid Staking Comparison

The liquid staking equilibrium pair (θ^*, g^*) is jointly determined by the Protocol Reaction function $\mathcal{R}_P$ and the Staker Reaction function $\mathcal{R}_S$ established in Theorem 1. Each parameter $\mu \in \{T, \gamma, \beta, r, \sigma_\xi\}$ generates a *direct effect* by shifting $\mathcal{R}_P$ for given g^*, and an *indirect effect* by first shifting $\mathcal{R}_S$ (changing g^*), which then moves θ^* along $\mathcal{R}_P$. Formally, let $F(\theta^*, g^*; \mu) = 0$ denote the Net Stress condition and $G(\theta^*, g^*; \mu) = 0$ denote the staker indifference condition. Applying the Implicit Function Theorem to this two-equation system yields the decomposition

$$\frac{d\theta^*}{d\mu} = \frac{1}{\Lambda}\left(\underbrace{\left.\frac{\partial \mathcal{R}_P}{\partial \mu}\right|_g}_{\text{Direct (Protocol)}} + \underbrace{\frac{\partial \mathcal{R}_P}{\partial g} \cdot \left.\frac{\partial g}{\partial \mu}\right|_{\mathcal{R}_S}}_{\text{Indirect (Staker)}} \right), \tag{35}$$

where $\Lambda = 1 - \frac{\partial \mathcal{R}_P}{\partial g}\frac{\partial \mathcal{R}_S}{\partial \theta^*} > 1 > 0$, because $\frac{\partial \mathcal{R}_P}{\partial g} < 0$ (Lemma 2) and $\frac{\partial \mathcal{R}_S}{\partial \theta^*} > 0$ (proof of Theorem 1).

Corollary 6 (Comparative Statics) *In the liquid staking equilibrium:*

(i) ***Staking reward (r):*** *The sign of $\frac{d\theta^*}{dr}$ is ambiguous. A higher return directly stabilizes the liquidity threshold by shrinking redemption queue demand, but indirectly destabilizes*

it by reducing the secondary market liquidity outlet. Moreover, the solvency threshold $\theta^{sol} = Ir/(I-T)$ is strictly increasing in r, introducing a third destabilizing channel that dominates for sufficiently large r.

(ii) ***Treasury size (T):*** *The sign of $\frac{d\theta^*}{dT}$ is ambiguous and is subject to the same condition on external market liquidity as in Corollary 1, with an additional destabilizing indirect channel operating through the Staker Reaction that is absent in traditional staking.*

(iii) ***Market price weight (β):*** *$\frac{d\theta^*}{d\beta} < 0$. Greater reliance on the secondary market price for external liquidity robustly stabilizes the protocol. The direct effect through the Protocol Reaction dominates a small, destabilizing indirect effect through the queue threshold x^*.*

(iv) ***Market friction (γ):*** *$\frac{d\theta^*}{d\gamma} > 0$. A less liquid secondary market (higher γ) is unambiguously destabilizing.*

(v) ***Market noise (σ_ξ):*** *The sign of $\frac{d\theta^*}{d\sigma_\xi}$ is ambiguous. Greater noise weakens the informational channel for external liquidity (τ_p falls) but simultaneously dampens the coordination role of the price signal.*

(vi) ***External liquidity noise (σ_l):*** *The sign of $\frac{d\theta^*}{d\sigma_l}$ is ambiguous. The direct effect through the Protocol Reaction is stabilizing: noisier external liquidity reduces the precision of the failure assessment, lowering the queue threshold x^* and shrinking redemption demand. The indirect effect through the Staker Reaction is destabilizing: the lower panic threshold g^* weakens the secondary market liquidity outlet.*

(vii) ***External liquidity efficiency (ρ):*** *$\frac{d\theta^*}{d\rho} < 0$. A more liquid external funding market is unambiguously stabilizing. The direct effect through the Protocol Reaction is the same as in Corollary 4 and dominates a small, same-signed indirect effect through the Staker Reaction.*

(viii) ***Mass of forced sellers (λ):*** *The sign of $\frac{d\theta^*}{d\lambda}$ is ambiguous. The direct channel is destabilizing: λ absorbs market depth, depresses the expected secondary market price, and weakens the liquidity outlet through the Staker Reaction. However, since $\kappa = I - \lambda$, a larger λ also shrinks the strategic mass, weakening the coordination engine of the run (stabilizing). The net effect depends on the equilibrium regime, paralleling the ambiguity in Corollary 5 for traditional staking. In liquid staking, forced sellers are routed to the secondary market rather than the queue, so the direct channel operates through market congestion rather than queue congestion.*

The proof is given in the Appendix.

The comparative statics reveal a rich interplay between the secondary market and the protocol's stability that has no counterpart in traditional staking. The key structural insight, embedded in the decomposition (35), is that every parameter operates through at most two channels: a direct shift of the Protocol Reaction $\mathcal{R}_P$ (holding g^* fixed) and an indirect shift via the Staker Reaction $\mathcal{R}_S$ (changing g^*, which then moves θ^* along $\mathcal{R}_P$). Because $\frac{\partial \mathcal{R}_P}{\partial g} < 0$ (Lemma 2), the indirect channel always reverses the sign of the staker shift: a parameter that raises g^* lowers θ^* through the market outlet, and vice versa. We discuss each parameter in turn.

Staking reward (r) and the solvency-liquidity tradeoff. The staking reward r operates through three channels. First, a *liquidity channel*: a higher r increases the benefit of staying relative to queueing ($r + \eta$ in the $(\delta_{liq} = s, \delta_{sol} = s)$ state), so stakers tolerate a higher failure risk before joining the queue. This lowers x^*, reduces redemption demand, and the liquidity threshold θ^*_{liq} declines. Second, a *market-outlet channel* unique to liquid staking: a higher r raises both the market price ($p = \frac{1+r}{\gamma}\Phi(\cdot)$) and the holding payoff $(1 + r)$, but the holding payoff increases by proportionally more since the market price is discounted by $\gamma > 1$. Stakers lower their panic threshold g^*, weakening the market liquidity outlet and raising θ^*_{liq} through the indirect channel. Third, a *solvency channel*: the solvency threshold $\theta^{sol} = Ir/(I - T)$ is strictly increasing in r. A protocol that promises returns it cannot deliver becomes more fragile regardless of withdrawal behavior.

In both traditional and liquid staking, the solvency channel operates through the four-state indifference condition: stakers anticipate the probability of the state $(\delta_{liq} = s, \delta_{sol} = f)$ in which the protocol survives the run but cannot pay $1 + r$, and this prospect amplifies their incentive to withdraw early. In traditional staking, the first and third channels oppose each other, producing the non-monotonic effect characterized in Corollary 2: for small r the liquidity channel dominates (stabilizing), while for large r the solvency channel dominates (destabilizing). In liquid staking, all three channels interact. For small r, the liquidity channel dominates and higher r is stabilizing. For intermediate r, the market-outlet channel may reverse the sign even before solvency binds. For large r, the solvency channel dominates and higher r is unambiguously destabilizing.

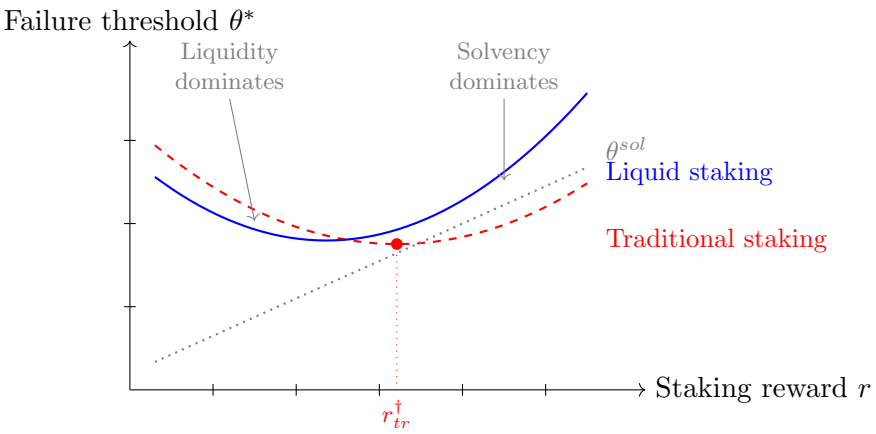

Figure 6. Effect of staking reward r on the failure threshold θ^*. Both traditional staking (dashed red) and liquid staking (solid blue) exhibit non-monotonic patterns. For small r, the liquidity channel dominates: a higher reward reduces run incentives and lowers θ^*. For large r, the solvency channel dominates: the rising obligation threshold θ^{sol} (dotted) makes the protocol more fragile. Liquid staking can exhibit an earlier reversal due to the additional market-outlet channel. (Schematic illustration.)

Figure 6 illustrates this. Traditional staking (dashed) now displays a U-shaped pattern reflecting the solvency-liquidity tradeoff from Corollary 2. Liquid staking (solid) can display a more complex non-monotonic pattern due to the additional market-outlet channel.

Treasury size (T) and the amplification of traditional-staking instability. The treasury comparative static connects directly to Corollary 1 for traditional staking, but with an important amplification. In liquid staking, the direct effect is $\partial F/\partial T = (\rho + \beta)(1 + \theta^*) - 1$, which is destabilising when $(\rho+\beta)(1+\theta^*) > 1$. Compared with traditional staking (where the condition is $\rho(1 + \theta^*) > 1$), the additional $\beta\theta^*$ term means that increasing the treasury is destabilising over a wider parameter range: a larger treasury shrinks the invested base $(I - T)$, which now reduces both the fundamental-based and market-price-based channels of external liquidity.

Liquid staking also adds a strictly destabilising indirect channel through the Staker Reaction. When T increases, the failure payoff T/I rises, which raises the retention value V in the early indifference condition. Stakers at the margin therefore require a worse private signal before they prefer to sell in the market; the panic threshold g^* falls. A lower g^* weakens the market liquidity outlet (Lemma 2), raising θ^*. This indirect channel is strictly positive (destabilising) regardless of ρ, so the region in which the treasury is net stabilising is strictly smaller under liquid staking than under traditional staking.

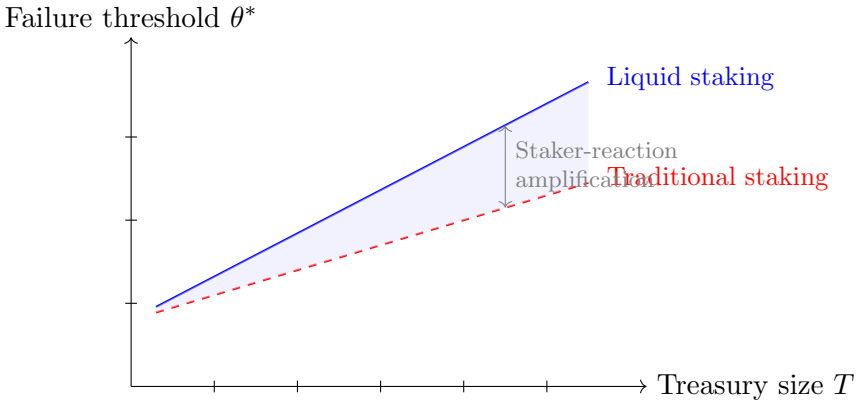

Figure 7. Effect of treasury size T on the failure threshold θ^*. In the destabilising regime ($(\rho+\beta)(1+\theta^*) > 1$), both the direct channel (forgone investment reducing both liquidity channels) and the indirect channel (lower panic threshold g^*) raise θ^* in liquid staking (solid blue), producing a steeper upward slope than in traditional staking (dashed red). The shaded gap represents the additional fragility introduced by the Staker Reaction channel that is absent in traditional staking. (Schematic illustration.)

Figure 7 illustrates this contrast for the destabilising regime, in which both direct and indirect channels are destabilising for liquid staking. For the stabilising regime, the same figure applies with a less steep upward slope for traditional staking and a comparison that reveals the liquid staking curve always lies above (more fragile) for the same T.

Market price weight (β). The parameter β enters the model through the resources available to the protocol: $T + (I - T)[\rho(1 + \theta^*) + \beta(y(P) + 1)]$. A higher β amplifies the contribution of the market price signal $y(P)$ to external liquidity, scaled by the invested base $(I - T)$. When fundamentals are sound ($\theta^* > 0$), the signal is on average positive, so external lenders extend more credit, reducing the failure threshold. Crucially, β has no counterpart in traditional staking because there is no secondary market from which external lenders can extract a signal. There is also a small indirect channel: β enters the queue threshold x through the $\beta(y(P) + 1)$ term, so

a higher β shifts the late indifference threshold x^*, raises the retention value $V(g)$, and slightly lowers the panic threshold g^*; via $\partial \mathcal{R}_P / \partial g < 0$, this is weakly destabilizing. However, this indirect effect is second-order, so the result $\frac{d\theta^*}{d\beta} < 0$ is robust: the market price acts as a form of collateral that expands the protocol's borrowing capacity, and institutional arrangements permitting market prices to serve as credit signals are stability-enhancing in liquid staking.

Market friction (γ) and the liquidity-outlet mechanism. The result $\frac{d\theta^*}{d\gamma} > 0$ is perhaps the most distinctive finding of the comparative statics, because it operates *entirely* through the indirect (Staker Reaction) channel. Since the price signal $y(P)$ is constructed to be invariant to γ (equation (20)), the parameter does not enter the failure condition or the late indifference that pins x^*, so the direct effect is zero. The indirect effect arises because γ scales the market price: $p = \frac{1+r}{\gamma}\Phi(\cdot)$. A higher γ lowers the expected proceeds from selling in the secondary market, making the liquidity outlet less attractive to stakers. Stakers therefore require a *worse* private signal before they are willing to sell early, pushing the panic threshold g^* down. But a lower g^* means fewer agents use the market as an escape valve; redemption pressure accumulates in the queue instead, raising the failure threshold θ^*.

Figure 8 illustrates this mechanism. In traditional staking (dashed line), γ has no effect because there is no secondary market: the failure threshold is flat. In liquid staking (solid line), the failure threshold rises with γ, reflecting the diminishing role of the market outlet. The gap between the two curves widens as γ grows, highlighting how a liquid secondary market is most valuable precisely when it is deep.

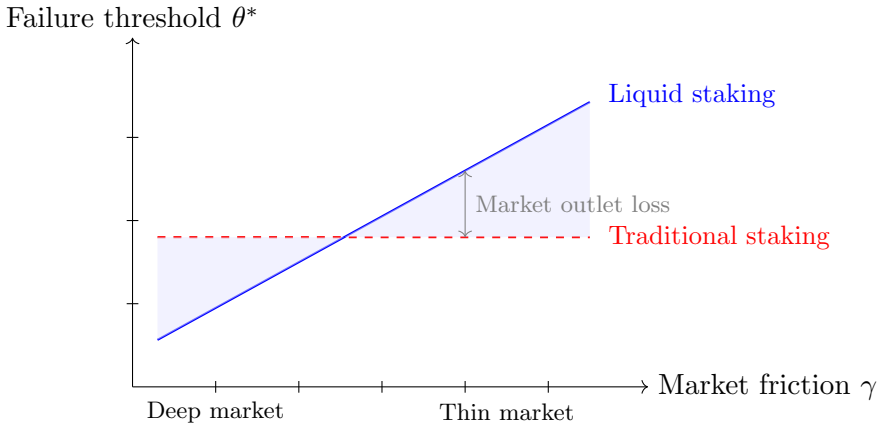

Figure 8. Effect of market friction γ on the failure threshold θ^*. In liquid staking (solid blue), a thinner secondary market raises θ^* because fewer stakers use the market as a liquidity outlet, concentrating redemption pressure in the queue. In traditional staking (dashed red), γ has no effect because there is no secondary market. The shaded region represents the stability gain from market depth that is unique to liquid staking. (Schematic illustration.)

Market noise (σ_ξ). The noise parameter σ_ξ enters the model through two distinct channels with opposing signs. First, it reduces the precision of the price signal: $\tau_p = (b + \kappa/\sigma_s)^2/(b^2\sigma_m^2 + \sigma_\xi^2)$, so a noisier market makes the signal $y(P)$ less informative about the fundamental θ. External lenders who observe $y(P)$ face greater uncertainty about whether a low price reflects poor

fundamentals or market noise. If lenders respond by tightening credit in times of uncertainty, the $(I-T)\beta y(P)$ resource channel weakens, raising θ^* (destabilising direct effect). Second, noisier prices reduce the coordination role of the secondary market: stakers who condition on $y(P)$ form less precise posteriors, which dampens the extent to which prices synchronise withdrawal decisions. This is the flip side of the uniqueness result in Theorem 1, which required σ_s to be sufficiently large precisely to prevent the price from perfectly coordinating a run. The opposing effects make the net impact of σ_ξ ambiguous in general, and numerical analysis is needed to determine which channel dominates for specific parameter configurations. The result is consistent with the broader finding in the literature that market transparency has dual effects: it improves resource allocation but can also amplify coordination failures (Goldstein, Ozdenoren, and Yuan, 2013).

External liquidity noise (σ_l) and the coordination-dampening paradox. In traditional staking, Corollary 3 establishes that noisier external liquidity is unambiguously stabilizing: it reduces the precision of stakers' failure assessments, dampening their ability to coordinate on a run. Liquid staking preserves this stabilizing direct channel but introduces a countervailing indirect channel through the secondary market.

The direct effect operates through the queue threshold x^*. A higher σ_l increases the variance of the failure event, making stakers who observe the price less certain about protocol survival. Under the condition $r + \eta > \psi$, these stakers tolerate greater ambiguity and lower their queue threshold x^*, reducing redemption demand. Holding the panic threshold g^* fixed, this lowers θ^* (stabilizing).

The indirect effect operates through the panic threshold g^*. The same increase in failure uncertainty also makes selling in the secondary market relatively less urgent: stakers who observe only their private signal s_i are less sure that the protocol will fail, so they require a worse signal before choosing to sell. The panic threshold g^* falls, weakening the secondary market as a liquidity outlet. Since fewer agents sell their tokens to risk-tolerant market makers, more redemption pressure accumulates in the queue, raising θ^* (destabilizing).

Which channel dominates depends on the depth of the secondary market and the relative mass of early sellers versus queue joiners. When the market is deep (γ small) and the liquidity outlet is important, the indirect destabilizing channel is stronger, and noisier external liquidity can be destabilizing in contrast to the traditional staking result. This highlights that the secondary market, while serving as a shock absorber, also creates a channel through which uncertainty about external support can paradoxically increase fragility.

External liquidity efficiency (ρ). The parameter ρ acts in liquid staking through the same direct channel as in traditional staking. Higher ρ raises external liquidity by $(I-T)(1+\theta^*)$ per unit, shifting the Protocol Reaction down: $\partial F/\partial\rho = -(I-T)(1+\theta^*) < 0$ when $\theta^* > -1$, so $\partial\mathcal{R}_P/\partial\rho < 0$. The indirect effect through the Staker Reaction is more subtle. A higher ρ also enters the queue threshold x^* (which scales with $\rho(z+1)$), making queueing more attractive because the protocol is more likely to honour it. This raises the retention value $V(g)$ in the early indifference, lowers the panic threshold g^*, and, via $\partial\mathcal{R}_P/\partial g < 0$, raises θ^* slightly. The direct channel dominates because it scales with the protocol's full invested base, while the

indirect channel operates only through a marginal substitution between market and queue. The net effect is unambiguously stabilizing, mirroring the traditional-staking result of Corollary 4: external funding-market depth is one of the few parameters whose stability benefit translates without offset into liquid staking. Notably, in contrast to the parameters that are unique to liquid staking (β, γ), the stabilizing role of ρ is shared with traditional staking and operates through the same balance-sheet logic.

Mass of forced sellers (λ) and the market congestion channel. The mass of non-strategic, liquidity-shocked sellers λ has a markedly different role under liquid staking than under traditional staking. In the traditional case (Corollary 5), λ mechanically expands the redemption queue and is destabilizing on the queue side. In liquid staking, these forced sellers are routed to the secondary market instead, so they do not enter the queue directly. Rather, they absorb depth in the market: from the price equation (19), $\partial P/\partial\lambda < 0$, and the early-indifference condition gives

$$\frac{\partial G}{\partial \lambda} = \frac{\partial E[p|g]}{\partial \lambda} = -\frac{1+r}{\gamma\sqrt{1+\sigma_Z^2}}\,\phi\left(\frac{\mu_Z(g)}{\sqrt{1+\sigma_Z^2}}\right) < 0,$$

because $\mu_Z(g)$ contains $-\lambda$ as an additive term. Lower expected proceeds from selling early make the market outlet less attractive, so the panic threshold g^* falls. Through $\partial\mathcal{R}_P/\partial g < 0$, this raises θ^*: as fewer strategic stakers use the secondary market as an escape valve, more redemption pressure crowds the queue, and the protocol becomes more fragile. However, this direct market-congestion channel is offset by the indirect κ channel: since $\kappa = I - \lambda$, a larger λ shrinks the strategic mass, weakening the coordination engine of the run and lowering θ^*. The net effect is ambiguous and depends on the equilibrium regime, just as for traditional staking. In our calibration (Table 6, Panel B), the two channels approximately cancel for liquid staking: the run probability barely moves as λ varies from 0.06% to 6.5%.

Taken together, Corollary 6 identifies two unambiguously signed results unique to liquid staking: market depth is stabilising (γ up is bad) and market price informativeness in the credit market is robustly stabilising (β up is good). The external funding depth (ρ) is also stabilising. The remaining five parameters (r, λ, σ_ξ, T, σ_l) have ambiguous effects because the direct and indirect channels conflict. These ambiguous cases share a common structure: a parameter that is stabilising through the Protocol Reaction can simultaneously weaken the secondary market liquidity outlet through the Staker Reaction, reversing or attenuating the net effect. This feedback between market participation and protocol resilience is the central mechanism distinguishing liquid from traditional staking.

7 Mapping the Model to Data: Quantifying Run Risks for Traditional and Liquid Staking

7.1 Data Sources

We construct a daily dataset spanning May 2023 to April 2026, beginning when the Lido withdrawal queue was activated (the Shapella upgrade, April 12, 2023). The empirical analysis combines on-chain data, market price data, third-party governance data, and a market sentiment index.

Lido withdrawal queue data (on-chain). We fetch all withdrawal-related events directly from the Ethereum blockchain by querying the Lido `WithdrawalQueueERC721` contract[7] via a public Ethereum RPC node (`https://ethereum-rpc.publicnode.com`). We retrieve two event types: `WithdrawalRequested` (each individual stETH withdrawal request, including the requestor, owner, amount in stETH, amount in shares, and block number) and `WithdrawalsFinalized` (each oracle-reported finalization batch). From these events, we construct (i) a request-level panel of 119,673 individual withdrawal requests with their wait times to finalization and (ii) a daily panel of queue activity (number of requests, total stETH requested, queue size, finalized amounts).

Lido protocol balance sheet (on-chain). We query the Lido stETH contract [8] at daily block intervals (every 7,200 blocks, approximately 24 hours) for three protocol-level state variables: (i) `getBufferedEther()` (the unallocated ETH buffer), (ii) `getTotalPooledEther()` (the total ETH staked through the protocol), and (iii) `getTotalShares()` (the total stETH shares outstanding). We compute the share rate as the ratio of total pooled ether to total shares; the staking APR is then derived as the annualized daily change in the share rate, smoothed over a 30-day rolling window.

stETH and ETH price data. Hourly price, trading volume, and market capitalization series for both stETH and ETH are obtained from CoinMarketCap's API for the full sample. From these series we compute (i) the stETH/ETH discount, $(p_{\text{stETH}} - p_{\text{ETH}})/p_{\text{ETH}}$, at hourly and daily frequencies (closing, mean, minimum), (ii) realized volatility from hourly log returns over a rolling 7-day window, annualized by $\sqrt{\text{hours per year}}$, and (iii) stETH outstanding, inferred as market capitalization divided by price (since CoinMarketCap reports market cap directly). The hourly data are aggregated to daily frequency by taking the daily minimum, maximum, mean, and closing values of the discount.

wstETH collateral parameters on Aave V3. Loan-to-value (LTV) ratios, liquidation thresholds, and liquidation bonuses for wrapped stETH on Aave V3 are collected from Aave governance proposals (Aave Improvement Proposals, AIPs). We record four discrete parameter changes during the sample period: AIP-121 (Aave V3 launch with wstETH at 68.5% LTV in January 2023), AIP-252 (LTV increase to 71% in June 2023), AIP-390 (parameter update to

[7] `0x889edC2eDab5f40e902b864aD4d7AdE8E412F9B1`

[8] `0xae7ab96520DE3A18E5e111B5EaAb095312D7fE84`

75% LTV in March 2024), and the most recent governance proposal raising LTV to 79% in October 2024.

Crypto Fear and Greed Index. Daily values of the Crypto Fear and Greed Index (FnG) are obtained from the public API at `https://api.alternative.me/fng/`. The index ranges from 0 (extreme fear) to 100 (extreme greed), with five categorical labels (Extreme Fear, Fear, Neutral, Greed, Extreme Greed). The full available history runs from February 2018 to the present; we restrict the sample to dates within the regression window (May 2023–April 2026).

Sample construction. The four data sources are merged on calendar date. After the merge, the regression sample contains 1,047 – 1,058 daily observations (the exact count depends on the variables included, since the on-chain queries occasionally fail and the sentiment index has occasional gaps). All replication code and intermediate data files are publicly available.

7.2 Empirical Moments of T, r, ρ/β, and η

Table 1 and Figures 9–11 reveal four facts about the Lido protocol's balance sheet that are central to understanding its stability and closely related to the key parameters in our model.

First, the buffer is remarkably small. The median daily buffer is 4,155 ETH, which is just 0.05% of the 9.1 million ETH pooled in the protocol (Panel A). By contrast, the median daily redemption request is 9,248 ETH, more than double the buffer. Panel A2 shows that on 76.2% of days, the daily redemption request exceeds the buffer (median Request/Buffer ratio of 2.1). The protocol therefore cannot meet a typical day's withdrawal demand from its reserves alone. Instead, it relies on a continuous flow of validator exits, incoming deposits, and external liquidity sources, including the secondary market and external liquidity providers, to service the queue.

Second, the staking APR declined steadily over the sample, from approximately 5% in mid-2023 to 2.5% by early 2026, with a sample mean of 3.2% (Panel A). This decline, which is also shown over time in Figure 9, reflects the growing number of Ethereum validators competing for a fixed pool of block rewards. In the model, the APR corresponds to the reward r: the return promised to patient stakers conditional on protocol survival. Its maximum is closely tied to the fundamental payoff θ.

Third, the collateral haircut for wstETH (wrapped stETH) on Aave V3 has fallen substantially over the sample (Panel B). Figure 10 shows that through four governance decisions, Aave increased the LTV from 68.5% to 79%, reducing the effective haircut from 31.5% to 21%. This secular improvement in external borrowing terms corresponds to an increase in ρ or β (more liquid external finance) in the model.

Fourth, Figure 11 shows that the wait times for redemption requests to be filled are typically 1–3 days but spiked to 12–13 days in August–September 2025, coinciding with a period of heavy redemption activity and validator exit delays. In the model, wait times correspond to the cost η of early withdrawal via the redemption queue.

Taken together, the declining haircut and the declining APR over the sample period are mutually

consistent: as external liquidity improved (ρ rose) and the reward r fell toward the stable region, the protocol became less fragile.

7.3 Price Discount and Redemption Queue

We now relate the empirical observations on the price discount and the redemption queue to their equilibrium properties derived in our model. Table 2 provides the summary statistics of these variables for the sample period, along with the volatility of ETH and stETH returns and the market capitalization of stETH. The daily closing price discount averages -0.12%, while the minimum intraday discount averages -0.44%. The daily redemption rate averages 0.15% of stETH outstanding.

The risk-return profile of holding ETH over our sample is remarkably unattractive: the full-sample annualized mean return is 5.02%, while the mean annualized volatility (7-day rolling window) is 59.3%. The implied Sharpe ratio is approximately 0.08, an order of magnitude below the 0.5–0.6 range typical of diversified equity indices. Without the staking reward, there is little economic reason for a risk-averse investor to hold ETH rather than cash or Treasury bills. The 3.2% staking APR reported in Table 1 is therefore not a minor supplement but the central compensation for bearing protocol risk: it raises the expected return from 5.0% to roughly 8.2%, increasing the risk premium by more than half. This economic significance is precisely why the Corollary on r matters in the model. Because the staking reward is the marginal rationale for holding stETH in the first place, any indication that the protocol might fail to deliver r, from the solvency channel $\theta^{sol} = Ir/(I-T)$ rising above sustainable levels, can trigger disproportionately large withdrawal responses.

Price Discount, Redemption, and FnG. Table 3 shows how market sentiment, proxied by the Crypto Fear and Greed Index (FnG), affects the price discount and the redemption decision. Panel A reveals an asymmetric pattern in the raw averages. The stETH discount is *narrower* during Extreme Fear (-0.378%) and Fear (-0.363%) than during Neutral (-0.455%) or Greed (-0.479%). Before reading this as a causal effect, however, it is essential to control for the strong autocorrelation in both the discount and the redemption rate, since periods of fear tend to cluster and so do the dependent variables. Panel B therefore reports regressions with one and five lags of the dependent variable.

After controlling for the lagged discount, the continuous FnG coefficient shrinks from -0.0019 to -0.0012 ($t = -2.95$) with one lag and further to -0.0007 ($t = -1.85$) with five lags. The effect remains negative but becomes marginal. The lagged discount itself is highly persistent (Lag1 $= 0.367$, $t = 12.82$), and the longer-lag specification reveals additional persistence at lags 2, 3, and 5 (R^2 rises from 15% to 20%). The column (5) regression with category dummies shows that the Fear dummy remains positive and significant ($+0.056$, $t = 2.01$), meaning that the Fear regime coincides with a narrower-than-average discount even after absorbing persistence. Extreme Fear, Greed, and Extreme Greed are all insignificant once autocorrelation is controlled for.

The redemption result is more robust. With one lag, the continuous FnG coefficient is $+0.00078$

($t = 2.09$); with five lags, it is +0.00059 ($t = 2.66$) and actually *stronger* with more lag controls. This stability is consistent with the predictive regression in Table 4 Panel C (with FnG), which found a positive coefficient of +0.0006 ($t = 2.52$). The category dummy regression in Column (6) shows the effect is concentrated in the Extreme Greed regime (+0.059, $t = 1.88$): during peak bullishness, stakers withdraw to chase higher yields elsewhere, consistent with the outside-opportunity channel as specified in the model.

After controlling for lagged variables, the FnG effect on the *discount* is weak and specification-dependent: the raw univariate relationship appears to reflect autocorrelation in both series more than a true sentiment channel. The FnG effect on *redemptions*, by contrast, is robust to lag specification and consistent with the outside-opportunity interpretation. This asymmetry is consistent with the model: the secondary market price $y(P)$ is the primary channel through which strategic stakers assess protocol risk (Table 4), while market-wide sentiment matters mainly through the outside-opportunity channel for non-strategic stakers (λ).

Lead-Lag between Price Discount and Redemption. We now examine the validity of our model assumption that price informs the redemption decision by studying the lead-lag relationship between the price discount and the redemption rate, while accounting for the effect of market sentiment on both variables. Table 4 uses Granger causality to test the direction of the dynamic relationship between the stETH discount and the redemption rate. The question is whether the secondary market price *leads* withdrawal decisions (as in the model setup), or whether redemptions drive prices, or both. Standard lead-lag correlations cannot answer this question because both variables are autocorrelated. Granger causality resolves this by asking whether one variable has *incremental* predictive power after controlling for the other variable's own history. To ensure that the results are not driven by market-wide sentiment, each test is reported both without and with the Crypto Fear & Greed Index as an additional control.

Panel A of Table 4 finds evidence supporting that yesterday's price discount Granger-causes today's redemption. Without the FnG control, the F-statistic is 12.80 ($p < 0.001$) with 1 lag. With the FnG control, the result remains highly significant: $F = 10.62$ ($p = 0.001$). The result is robust across all lag specifications in both columns. Controlling for market-wide sentiment barely attenuates the effect, confirming that the discount's predictive power reflects protocol-specific information, not general market fear.

The economic magnitude can be read from Panel C of Table 4. Without the FnG control, the coefficient on the lagged discount is −0.0563 ($t = -3.58$); with the FnG control, it is −0.0515 ($t = -3.26$). A 1 percentage point deeper discount yesterday predicts a 0.05–0.06 percentage point higher redemption rate today. Given the mean daily redemption rate of 0.15% (Table 2), this represents roughly one-third of the daily average, an economically meaningful effect.

Panel B of Table 4 does not find evidence that redemption Granger-causes the discount. Without the FnG control, $F = 1.25$ ($p = 0.26$); with the FnG control, the result is even weaker: $F = 0.71$ ($p = 0.40$). Adding sentiment as a control absorbs some residual correlation, further confirming that redemptions have no independent predictive power for the discount. Naive regressions of discount on lagged redemption (without controls) do produce significant coefficients, but the Granger test reveals this as spurious, driven by common factors rather than a causal channel from redemption to price.

Panel C of Table 4 shows the results from key predictive regressions for redemption. The right-hand columns show that the Fear & Greed Index has a small but significant positive coefficient (+0.0006, $t = 2.52$): higher market sentiment (greed) predicts slightly higher redemptions. This likely reflects the outside-opportunity channel ($\widetilde{c}$) in the model: in bullish markets, stakers withdraw to chase higher yields elsewhere. Crucially, including FnG barely changes the discount coefficient (-0.056 to -0.052) and raises R^2 only modestly (from 4.75% to 5.32%), confirming that the discount captures the relevant information channel above and beyond market-wide sentiment.

These results support the model mechanism. The one-directional causality, from price to withdrawal, but not from withdrawal to price, is precisely the mechanism at the heart of our model. The secondary market price $y(P)$ aggregates dispersed information about the protocol's fundamental $\tilde{\theta}$, and strategic stakers use this public signal to update their beliefs and coordinate withdrawal decisions. The empirical finding that the discount Granger-causes redemption but not vice versa confirms that the price acts as an *information channel* rather than a mere reflection of withdrawal pressure. The robustness to controlling for the Fear & Greed Index further confirms that this is a protocol-specific information effect, not a manifestation of market-wide sentiment. This supports the modeling choice to treat the market price as a signal that enters the stakers' indifference conditions, rather than as a passive outcome of their trading.

7.4 Testing Comparative Statics and Simulating Run Risk

Testing Comparative Statics. We empirically examine how the model's comparative statics map to the data. In the model, the equilibrium price discount inferred from equation (19) reflects the panic threshold g^*, which depends on the failure threshold θ^*. The comparative statics of θ^* with respect to the staking reward r, the treasury T, external liquidity ρ, and market sentiment therefore generate predictions about the behavior of the discount. Table 5 reports these findings. Specifically, we test whether the variables corresponding to the model's key parameters explain variation in the stETH discount, after controlling for the persistent component of the discount via its one-day lag. The lag coefficient is large and highly significant in every specification (0.28–0.36, $t > 9$), confirming that the discount is persistent. Since the discount is negative on average (-0.44%), a *positive* coefficient means the variable narrows the discount (stabilizing), and a *negative* coefficient means it deepens the discount (destabilizing).

The APR remains the dominant predictor after the lag control (-0.1211, $t = -8.37$, $R^2 = 19.3\%$) and is essentially unchanged in the multiple regression (-0.1127, $t = -7.48$). Higher staking rewards are associated with *deeper* discounts. The model predicts an ambiguous effect of r: the liquidity channel (higher r reduces run incentives, stabilizing) opposes the solvency channel ($\theta^{sol} = Ir/(I - T)$ rises with r, destabilizing). The negative empirical sign indicates that the solvency channel dominates at observed APR levels, consistent with the model's Corollary on r in the region where r exceeds the critical level.

Higher LTV is associated with a *narrower* discount (+0.0142, $t = 5.82$, Column 4), consistent with the model's prediction that higher ρ/β stabilizes the protocol by expanding external borrowing capacity. LTV is excluded from the multiple regression because it is highly collinear with

APR (corr $= -0.85$).

Market greed is associated with deeper discounts in the univariate-with-lag regression (-0.0013, $t = -3.04$, Column 1), but the effect weakens and becomes insignificant in the multiple regression ($t = -1.27$). APR absorbs most of the sentiment-driven variation, consistent with the interpretation that the secondary market price is the primary information channel and sentiment enters mainly through the outside-opportunity channel (as shown in Table 3 for redemptions).

A larger buffer is associated with a deeper discount (-0.1942, $t = -2.96$, Column 2), consistent with the model's prediction that the treasury can be destabilizing in liquid staking: a larger T diverts capital from the invested base ($I - T$), shrinking external liquidity. The effect weakens in the multiple regression ($t = -1.55$), consistent with the model's ambiguous net prediction for T.

Overall, all four model parameters have the predicted signs even after controlling for discount persistence, and APR alone explains 19.3% of daily discount variation with the lag control. The multiple regression (Column 5, excluding the collinear LTV) achieves $R^2 = 19.6\%$. The lag coefficient of approximately 0.28 in the multiple regression indicates that roughly 72% of daily discount variation is driven by contemporaneous factors rather than inertia.

Simulating Run Risk. We conduct numerical exercises to compare the run probabilities for both traditional and liquid staking protocols. We observe the empirical range of key parameters such as T, r, and ρ from the data, as well as equilibrium quantities such as the price and the price discount. We use observed price discounts and return volatilities to back out the unobservable parameters (such as σ_θ, σ_s, σ_l, σ_b, σ_ξ, and λ), subject to the conditions for equilibrium existence and uniqueness (e.g., $\rho(I-T) > \kappa\sigma_s\tau_\theta$ in Proposition 1, and $1+r > \frac{\gamma T}{I\Phi(b\bar{\theta}-I)}$ and $\gamma > \Phi(b\bar{\theta}-\lambda)$ in Proposition 2). The equilibrium ex ante run probability is stated in the following corollary.

Corollary 7 *The ex ante run probabilities for the liquid and traditional staking protocols are, respectively:*

$$\Pr(\textit{liquid staking run}) = \int_{-\infty}^{\theta^*} d\Phi(\tilde{\theta}), \tag{36}$$

$$\Pr(\textit{traditional staking run}) = \int_{-\infty}^{\theta^*_{tr}} d\Phi(\tilde{\theta}), \tag{37}$$

where, for liquid staking, θ^* is the unique fixed point of the staker's reaction function $\mathcal{R}_S : \theta^* \to g$ given by $F_{staker}(g;\theta^*) = E[p|g] - V(g;\theta^*) = 0$ and the protocol's reaction function $\mathcal{R}_P : g \to \theta^*$ given by $F(\theta^*, g) = 0$, with $F(\theta^*, g)$ defined in (61); for traditional staking, θ^*_{tr} is defined in (11). Because we focus on liquidity-driven runs, the solvency constraint does not bind. Hence $\theta^* < \theta^{sol} = Ir/(I-T)$ and $\theta^*_{tr} < \theta^{sol} = Ir/(I-T)$.

We quantify the run-probability reduction from liquid staking by solving both equilibria numerically at calibrated parameter values. This is a counterfactual comparison. It is not a moment-matching exercise and designed to illustrate how the model's comparative statics translate into run-probability magnitudes under realistic parameter configurations drawn from the Lido data.

The implied unconditional price moments ($E[P] \approx 0.65$, $\text{std}(P) \approx 0.26$) differ from the empirical moments ($E[P] \approx 0.99$, $\text{std}(P) \approx 0.01$) because the model generates more fundamental uncertainty than is realised in the stable post-2022 sample; the exercise should be read as a structural comparison under counterfactual stress, not as a fit to the observed price distribution.

We fix the structural parameters at $T = 0.0008$ (the average Lido buffer share from Panel A of Table 1 is 0.08% of total pooled ETH), $r = 3.22\%$ (sample mean staking APR), $\eta = 0.05$ (a 5% early-withdrawal haircut, reflecting the typical stETH discount during stress periods and the queue-wait cost), and $\lambda = 0.0016$ (the average daily redemption rate is 0.16% of stETH outstanding from Panel B of Table 2). Market parameters are $\beta = 0.30$ (the weight external lenders place on the market price signal, set below the wstETH LTV of 0.75 to reflect that only a fraction of the market price translates into credit supply) and $\rho = 0.10$ (the effective fast-liquidity capacity from the protocol's direct borrowing channel). The market friction parameter γ is varied in the sensitivity analysis; the baseline is $\gamma = 1.30$. We normalise $I = 1$, so $\kappa = I - \lambda$, and set the disbanding cost $\psi = 0.10$. The mean of the protocol fundamental is $\bar{\theta} = 0.045$, comfortably above the solvency threshold $\theta^{sol} = Ir/(I - T) \approx 0.0322$.[9]

We set $\tau_\theta = 10$ ($\sigma_\theta \approx 0.316$) and fix the private signal precision at $\tau_s = 100$ ($\sigma_s = 0.1$), giving a signal-to-prior ratio $\tau_s/\tau_\theta = 10$. The remaining precisions are $\tau_m = \tau_s/2 = 50$ and $\tau_\xi = \tau_s/2 = 50$, with the external lender's signal precision $\tau_l = 0.5\,\tau_\theta = 5$. The market maker's demand sensitivity is set at $b = 0.10$, reflecting the Bayesian posterior weight $b \approx \tau_m/(\tau_m + \tau_\theta + \tau_{\text{other}})$ where τ_{other} captures other signals held by the market maker that are not modelled explicitly.[10]

For traditional staking, we solve the two-equation system from Proposition 1: the liquidity-balance condition (10) and the indifference condition $\theta^*_{tr} - \mu_g = \sigma_v \Phi^{-1}((r+\eta)/(r+\eta+\psi))$. For liquid staking, we solve the full equilibrium from Theorem 1: the protocol reaction $F(\theta^*, g) = 0$ using the bivariate-normal queue probability $\kappa \Pr(s > g^*, x < x^*)$ from Lemma 1 together with resources $T + (I-T)[\rho(1+\theta) + \beta(y(P)+1)]$ as specified in equation (13), and the early indifference $E[p|g] = V(g; \theta^*)$ with the equilibrium price $E[p|g] = (1+r)/\gamma \cdot \Phi(\mu_Z(g)/\sqrt{1+\sigma_Z^2})$.

Because the solvency threshold $\theta^{sol} = 0.0322$ operates identically under both architectures (it does not depend on whether a secondary market exists), we report the *liquidity-only* run probability $\Pr(\tilde{\theta} < \theta^*) = \Phi((\theta^* - \bar{\theta})/\sigma_\theta)$ to isolate the channel that distinguishes liquid from traditional staking.

Table 6 reports the calibrated equilibria and liquidity-only run probabilities, varying γ, λ, and τ_s one at a time. At the baseline ($\gamma = 1.30$, $\lambda = 0.16\%$, $\tau_s = 100$), the traditional staking equilibrium has $\theta^*_{tr} \approx 0.022$ and a liquidity-run probability of 47.1%; liquid staking has $\theta^*_{liq} \approx -0.229$ and a run probability of 19.3%, yielding a gap of 28 percentage points.

At the baseline, the secondary market reduces the liquidity-driven run probability by 28 per-

[9] Validator income on Ethereum has two components: (i) consensus-layer issuance, providing roughly 3% APR; and (ii) execution-layer rewards from priority tips and MEV, which add roughly 1–2% APR. Our $\bar{\theta} = 4.5\%$ corresponds to a baseline validator return that just exceeds the promised payout.

[10] In the model, market makers receive a private signal $\tilde{s}_m$ with precision τ_m. A fully Bayesian market maker would weight this signal at $\tau_m/(\tau_m + \tau_\theta) = 50/60 \approx 0.83$. However, market makers in practice hold additional proprietary information sources (order flow data, cross-exchange signals) whose precision τ_{other} is large; a conservative $\tau_{\text{other}} \approx 450$ gives $b = \tau_m/(\tau_m + \tau_\theta + \tau_{\text{other}}) = 50/510 \approx 0.10$.

centage points: traditional staking has a 47.1% run probability while liquid staking has 19.3%. The gap is driven by the secondary market's dual role: it absorbs panicked stakers who would otherwise congest the redemption queue (reducing the demand side via the bivariate-normal queue probability $\kappa \Pr(s > g^*, x < x^*)$) and provides an additional price signal that external lenders use to extend credit (augmenting the resource side through $\beta(y(P)+1)$ in equation (13)).

We find that market depth is the critical determinant (Panel A). As γ rises from 1.01 to 5.0, the liquid staking threshold θ^*_{liq} rises from -0.351 to $+0.048$, and the gap narrows from 36.6pp to -3.3pp. At $\gamma = 5.0$, liquid staking is actually *less* stable than traditional: the secondary market is so thin that it cannot absorb selling pressure, and its informational role for external lenders is compromised by illiquidity. This confirms and sharpens Corollary 6(iv): a less liquid secondary market (higher γ) is unambiguously destabilising for liquid staking, and at sufficiently high γ the advantage disappears entirely. Importantly, γ does not affect the traditional staking equilibrium at all (there is no secondary market), so the entire effect operates through the liquid staking column.

Numerical exercises show that forced sellers affect both channels but the gap persists (Panel B). As λ rises from 0.06% to 6.5%, both run probabilities decline: traditional from 47% to 39% and liquid from 19% to 17%. The gap narrows from 28pp to 22pp but remains substantial. The decline in the traditional run probability reflects the indirect (Staker Reaction) channel: a larger λ shrinks the strategic mass $\kappa = I - \lambda$, weakening the strategic-coordination engine of the run. For liquid staking, forced sellers exit via the market and barely affect the queue, so the liquid run probability changes little.

Finally, simulations show that signal precision has a modest effect (Panel C). Higher τ_s slightly narrows the gap from 26pp at $\tau_s = 20$ to 27pp at $\tau_s = 500$, but the effect is small and non-monotonic. This reflects two offsetting forces: more precise signals tighten the coordination mechanism (raising traditional fragility) but also make the price signal less incrementally informative for external lenders (reducing liquid staking's advantage) and strategic stakers.

The numerical exercise confirms the model's central prediction: a functioning secondary market reduces the protocol's exposure to liquidity-driven runs by 28 percentage points at the baseline. The most striking finding is the role of market depth (γ): with a deep market ($\gamma = 1.01$), liquid staking's advantage reaches 37pp and the protocol's run probability is just 10.5%; with a thin market ($\gamma = 5.0$), the advantage reverses and liquid staking becomes *more* fragile than traditional. This threshold effect underscores the importance of institutional arrangements supporting deep secondary markets for staked tokens, such as collateral eligibility on Aave, integration with DEX aggregators, and market-maker incentive programmes, as first-order determinants of protocol stability. A protocol that launches a liquid staking derivative without ensuring adequate market depth may find that the secondary market amplifies rather than absorbs coordination risk.

8 Conclusion

The newly emerging decentralized economy provides researchers with a rare opportunity to examine high-value, unregulated institutions in real time, populated by relatively sophisticated participants and observable through granular on-chain data. Liquid staking is a particularly informative case: a protocol pools depositors' funds to stake on Ethereum and issues a tradeable derivative claim, combining bank-like liquidity transformation with two features absent from the canonical Diamond–Dybvig setting: a secondary market in which depositors can exit without burdening the protocol, and an external funding channel that observes the market price. We build a global-game model that incorporates both features, derive its comparative statics, and confront its predictions with three years of on-chain data from Lido.

A key insight is on the roles that the secondary market plays in the stability of liquid staking. First, in liquid staking, panicked depositors prefer to sell their tokens to risk-tolerant market makers rather than queue at the protocol, which transfers their claims out of the protocol's balance sheet. Increased panic therefore mechanically shrinks the redemption queue and lowers the failure threshold; the market acts as a shock absorber rather than a vehicle of contagion. Second, market price is informative. It aggregates dispersed private information and external lenders are able to use it to supply emergency funding to the protocol in times of stress. Third, it serves as a public signal and potentially coordinates runs. Interestingly, this mechanism requires private signals to be sufficiently *noisy*, the opposite of the precision condition in canonical global games. When private signals are too precise, the market price aggregates information efficiently and crashes during a panic; external lenders observing the crash infer insolvency with certainty and retract funding, overwhelming the outlet effect. Noise dampens the price impact and preserves external credit, sustaining a unique equilibrium.

A second contribution is the explicit separation of liquidity failure from solvency failure. Stakers face four payoff states depending on whether the protocol survives the run and whether it can honour its promised return at maturity. The state in which the protocol survives the run but turns out insolvent, which is unique to our framework, generates a powerful amplification mechanism: rational stakers anticipate the insolvency shortfall, which raises their incentive to withdraw early, which in turn tightens the liquidity constraint. This interaction implies that the staking reward r has a non-monotonic effect on stability: stabilising at low levels by raising the value of staying, but destabilising at high levels because the rising obligation tips the protocol into insolvency. Protocols that promise high returns are therefore exposed not only to direct solvency risk but also to amplified liquidity risk through staker anticipation.

A third theme is that several seemingly "safe" design features can have counterintuitive effects. A larger treasury, often viewed as a buffer against runs, can be destabilising because it diverts capital from the productive base $(I - T)$ and shrinks both the fundamental-based and market-price-based channels of external liquidity. In liquid staking the destabilising force is amplified by an additional staker channel: a larger pro-rata failure payoff raises the retention value, lowers the panic threshold, and weakens the market outlet. A larger mass of forced sellers λ has an ambiguous effect on traditional staking: the direct channel congests the queue, but the indirect channel shrinks the strategic mass and weakens coordination; the sign depends on the equilibrium regime. In liquid staking, forced sellers are routed to the secondary market and

barely affect the queue. By contrast, two parameters with no counterpart in traditional staking are unambiguously stabilising: the depth of the secondary market (low friction γ) and the weight that external lenders place on the market price (β). Both highlight the dual role of the market as a liquidity outlet and a credit-signal channel.

Empirically, we test the comparative statics on daily data from the Lido protocol over May 2023–April 2026. All four parameters with measurable proxies, staking APR, collateral haircut on Aave V3, treasury buffer, and an outside-opportunity proxy, move the stETH discount and the redemption rate in the directions the model predicts, even after controlling for autocorrelation and the Crypto Fear and Greed Index. We further document that the discount Granger-causes redemption demand but not vice versa, consistent with the model's view that the secondary market price aggregates information about fundamentals, which then drives staker withdrawal decisions.

A numerical exercise quantifies the model's predictions. At the baseline parameters, liquid staking reduces the liquidity-driven run probability by 28 percentage points relative to traditional staking (19% versus 47%). The advantage depends critically on market depth: with a deep secondary market ($\gamma = 1.01$), the gap reaches 37 percentage points and the liquid staking run probability is just 10.5%; with a thin market ($\gamma = 5.0$), the advantage reverses and liquid staking becomes more fragile than traditional. This finding underscores that a protocol launching a liquid staking derivative without ensuring adequate market depth may find that the secondary market amplifies rather than absorbs coordination risk.

The optimal design of an intermediary engaged in liquidity transformation is intricate, and the structures that have organically emerged in decentralized finance offer a useful laboratory in which to study it. Our results suggest that some of the design features that intuitively look safer, such as a larger reserve buffer, a higher promised return, a perfectly transparent secondary market, can have the opposite effect once the feedback between depositor coordination, external credit, and market prices is taken into account.

References

Angeletos, George-Marios, and Iván Werning, 2006, Crises and Prices: Information Aggregation, Multiplicity, and Volatility, *American Economic Review* 96, 1720–1736.

Correia, Sergio, Stephan Luck, and Emil Verner, 2024, Failing Banks, *Working Paper*.

Diamond, Douglas W., and Philip H. Dybvig, 1983, Bank runs, deposit insurance, and liquidity, *Journal of Political Economy* 91, 401–419.

Goldstein, Itay, Emre Ozdenoren, and Kathy Yuan, 2013, Trading frenzies and their impact on real investment, *Journal of Financial Economics* 109, 566–582.

Goldstein, Itay, and Ady Pauzner, 2005, Demand–deposit contracts and the probability of bank runs, *Journal of Finance* 60, 1293–1327.

Green, Edward J, and Ping Lin, 2003, Implementing efficient allocations in a model of financial intermediation, *Journal of Economic Theory* 109, 1–23.

Jacklin, Charles J., 1987, Demand Deposits, Trading Restrictions, and Risk Sharing, in Edward C. Prescott, and Neil Wallace, eds.: *Contractual Arrangements for Intertemporal Trade* (University of Minnesota Press,).

Morris, Stephen, and Hyun Song Shin, 2003, Global Games: Theory and Applications, in Mathias Dewatripont, Lars Peter Hansen, and Stephen J. Turnovsky, eds.: *Advances in Economics and Econometrics: Theory and Applications, Eighth World Congress* (Cambridge University Press,).

Ozdenoren, Emre, and Kathy Yuan, 2008, Feedback Effects and Asset Prices, *The Journal of Finance* 63, 1939–1975.

Peck, James, and Karl Shell, 2003, Equilibrium Bank Runs, *Journal of Political Economy* 111, 103–123.

Notation

$\widetilde{\theta}$	Payoff from the validators
$s_i = \widetilde{\theta} + \sigma_s \tilde{\epsilon}_i$	private signal to each staker
ϵ_i	standard normal
τ_s	precision of the private signal
τ_θ	inverse of the variance of the fundamental
τ_p	precision of $y(P)$
$y(P)$	information content of price
$\hat{\tau}$	$\tau_\theta + \tau_s + \tau_p$
$\tilde{\xi}$	market demand shock
τ_ξ	precision of the market demand shock
$\tilde{s}_m = \tilde{\theta} + \sigma_m \tilde{\epsilon}_m$	market maker's private signal
σ_m	standard deviation of market maker's signal noise
$\tau_m = 1/\sigma_m^2$	precision of market maker's signal
b	sensitivity of market-maker demand to signal
$\tilde{\epsilon}_l$	noise in external liquidity provider's signal
σ_l	standard deviation of external liquidity noise
c_{tr}	$\rho(I-T)\sigma_l\sigma_s/(\kappa + \sigma_s\rho(I-T))$, noise coefficient (traditional)
σ_v	$\sqrt{1/(\tau_\theta + \tau_s) + c_{tr}^2}$, effective posterior std. dev.

Figures

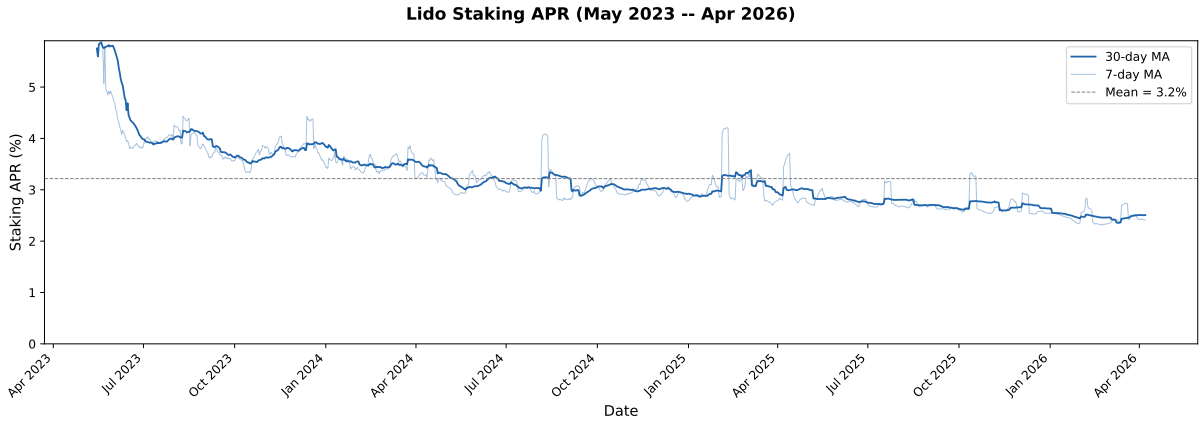

Figure 9. Lido Staking APR. The 30-day moving average APR (dark line) declined from approximately 5% in mid-2023 to around 2.5% by early 2026, with a sample mean of 3.2%. The lighter line shows the noisier 7-day MA. In the model, the staking APR corresponds to the reward parameter r.

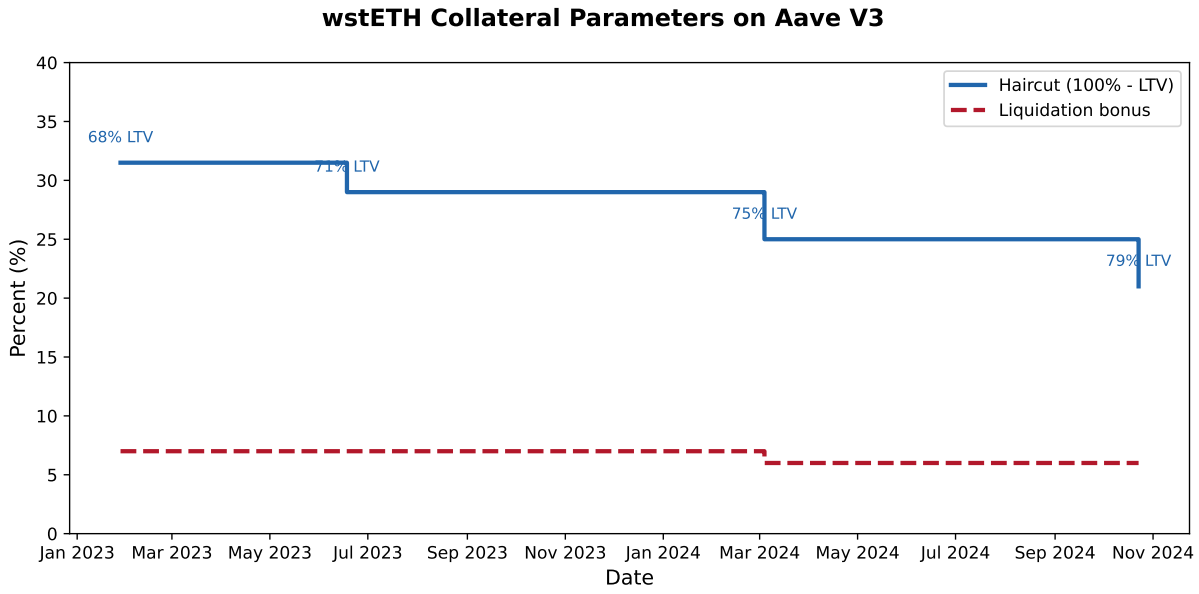

Figure 10. wstETH Collateral Haircut on Aave V3. The solid line shows the collateral haircut (100% − LTV), which has declined from 31.5% to 21% as Aave governance progressively increased the LTV ratio for wstETH. The dashed line shows the liquidation bonus.

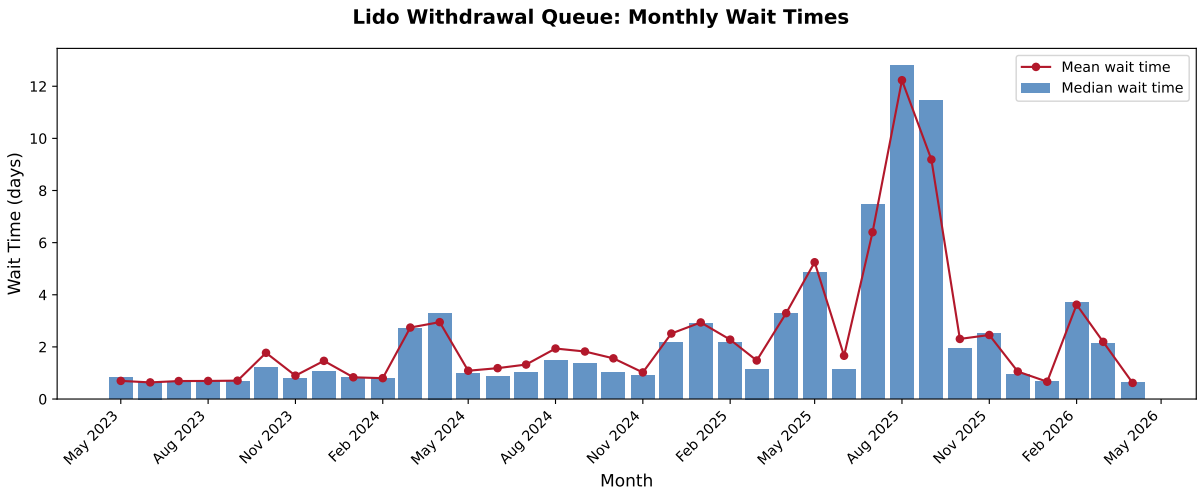

Figure 11. Withdrawal Queue Wait Times by Month. The blue bars show the monthly median wait time (days) for withdrawal requests to be finalized. The red line shows the monthly mean.

Tables

Table 1. Protocol Fundamentals. Panel A reports on-chain data from the Lido contract at daily block intervals. The buffer is unallocated ETH available for withdrawal fulfilment. The staking APR is computed from the daily change in the share rate, annualized and smoothed over 30 days. Panel A2 compares the buffer to daily redemption requests in ETH terms. Panel B reports wstETH collateral parameters on Aave V3. Sample: May 2023–April 2026.

	N	Mean	Std	P25	Median	P75	Max
Panel A: On-Chain Protocol Data							
Buffer (ETH)	1,051	7,028	11,140	1,760	4,155	8,219	168,044
Total pooled (M ETH)	1,051	9.10	0.68	8.79	9.25	9.60	9.90
Buffer (% of pooled)	1,051	0.08	0.14	0.02	0.05	0.09	2.50
Staking APR, 30d (%)	1,051	3.22	0.62	2.79	3.04	3.54	5.87
Panel A2: Buffer vs. Daily Redemption Requests							
Buffer (ETH)	1,049	6,768	10,917	1,760	4,155	8,190	168,044
Daily stETH requested	1,049	13,541	18,808	5,140	9,248	16,838	441,204
Request / Buffer ratio	1,049	6.1	12.7	1.0	2.1	5.2	232.3
Days buffer $\geq$ request: 250/1,049 (23.8%); Days buffer $<$ request: 799/1,049 (76.2%)							

Panel B: wstETH Collateral Parameters on Aave V3					
Date	Event	LTV	Liq. Thresh.	Liq. Bonus	Haircut
Jan 2023	Aave V3 launch	68.5%	79.5%	7.0%	31.5%
Jun 2023	LTV increase	71.0%	80.0%	7.0%	29.0%
Mar 2024	Parameter update	75.0%	82.0%	6.0%	25.0%
Oct 2024	LTV increase	79.0%	83.0%	6.0%	21.0%

Table 2. Summary Statistics of Prices and Redemptions. All panels cover May 2023 to April 2026, beginning when the Lido withdrawal queue was activated. Panel A reports daily stETH/ETH price data. The discount is $(p_{\text{stETH}} - p_{\text{ETH}})/p_{\text{ETH}}$; negative values indicate stETH trades below par. Annualized return is the mean daily log return multiplied by 365, reported as a single number over the full sample. Annualized volatility is computed from hourly log returns over a rolling 7-day window and annualized by $\sqrt{\text{hours per year}}$. Panel B reports withdrawal queue variables and the associated stETH market size (in million). Redemption % is daily stETH requested divided by stETH outstanding (inferred from CoinMarketCap market cap / price). Wait time is at the individual-request level (N = 119,421); all other queue variables are daily.

	N	Mean	Std	P25	Median	P75	Max
Panel A: Daily Price Data							
Discount, close (%)	1,058	−0.12	0.23	−0.21	−0.09	−0.02	2.04
Discount, min (%)	1,058	−0.44	0.30	−0.56	−0.36	−0.24	0.03
stETH annualized vol, 7d (%)	1,058	56.49	22.63	41.59	53.28	66.45	147.32
ETH annualized vol, 7d (%)	1,058	59.34	21.93	44.95	56.69	68.99	143.05
Full-sample annualized mean return: ETH = 5.02%, stETH = 5.11%							
Panel B: Withdrawal Queue and Market Size							
Redemption rate (%)	1,057	0.15	0.25	0.06	0.10	0.18	6.57
stETH requested (daily)	1,057	14,037	19,439	5,139	9,248	16,832	441,204
Daily requests (count)	1,057	113	49	83	103	132	520
Queue size (stETH)	1,057	42,028	50,019	9,694	22,922	55,071	430,545
Wait time (days)	119,421	2.74	3.30	0.76	1.33	3.37	16.93
stETH outstanding (M ETH)	1,057	9.10	0.67	8.79	9.24	9.60	9.88

Table 3. Crypto Fear & Greed Index: Breakdown and Regressions. Panel A reports the Fear & Greed Index (FnG) by sentiment classification, along with the average stETH discount and redemption rate in each regime. Panel B reports regressions of the discount and redemption rate on the FnG index, with lagged dependent variable controls to absorb autocorrelation. Columns (1)–(2) include one lag of the dependent variable. Columns (3)–(4) include five lags. Columns (5)–(6) use FnG category dummies (Neutral omitted) with one lag of the dependent variable. The discount is in %; the redemption rate is in %. Sample: May 2023–April 2026 (daily). ***, **, *: significance at 1%, 5%, 10%.

Panel A: Summary by Sentiment Classification

Classification	FnG range	N	Avg discount (%)	Avg redemption (%)
Extreme Fear	0–25	147	−0.378	0.140
Fear	26–46	200	−0.363	0.133
Neutral	47–54	194	−0.455	0.156
Greed	55–75	419	−0.479	0.153
Extreme Greed	76–94	96	−0.456	0.218

Panel B: Regressions with Lagged Dependent Variable Controls

	1 lag		5 lags		Dummies (1 lag)	
	Discount	Redemption	Discount	Redemption	Discount	Redemption
FnG (0–100)	−0.00122***	0.00078**	−0.00074*	0.00059***		
	(−2.95)	(2.09)	(−1.85)	(2.66)		
Extreme Fear					0.0368	−0.0166
					(1.22)	(−0.61)
Fear					0.0560**	−0.0237
					(2.01)	(−0.94)
Greed					−0.0222	−0.0048
					(−0.93)	(−0.22)
Extreme Greed					−0.0124	0.0586*
					(−0.36)	(1.88)
Lag 1 of dep. var.	0.3674***	0.0497	0.2613***	0.1638***	0.3626***	0.0474
	(12.82)	(1.62)	(8.49)	(5.30)	(12.60)	(1.54)
Lags 2–5 of dep. var.	—	—	Yes	Yes	—	—
Constant	−0.2118***	0.1053***	−0.1243***	0.0756***	−0.2837***	0.1498***
R^2	0.1504	0.0070	0.2030	0.0620	0.1539	0.0101
N	1,055	1,055	1,051	1,051	1,055	1,055

Table 4. Granger Causality: Discount and Redemption Rate. This table tests the direction of predictability between the stETH discount and the redemption rate. Panel A tests whether lagged discount improves the prediction of redemptions after controlling for lagged redemptions. Panel B tests the reverse. Each panel reports results without and with the Fear & Greed Index (FnG) as an additional control, to verify that the results are not driven by market-wide sentiment. Panel C reports the key predictive regressions. Sample: May 2023–April 2026 (daily). ***, **, *: significance at 1%, 5%, 10%.

	Without FnG control		With FnG control	
	F-stat	p-value	F-stat	p-value
Panel A: Does discount Granger-cause redemption?				
H_0: Lagged discount adds no predictive power				
1 lag	12.80	0.0004^{***}	10.62	0.0012^{***}
2 lags	6.42	0.0017^{***}	5.75	0.0033^{***}
3 lags	4.47	0.0040^{***}	3.94	0.0083^{***}
5 lags	2.61	0.0236^{**}	2.25	0.0475^{**}
Panel B: Does redemption Granger-cause discount?				
H_0: Lagged redemption adds no predictive power				
1 lag	1.25	0.2647	0.71	0.3981
2 lags	0.26	0.7678	0.12	0.8909
3 lags	1.59	0.1910	1.85	0.1360
5 lags	1.79	0.1127	1.87	0.0962^{*}

Panel C: Predictive Regressions for Redemption rate$_t$ (%)

	Without FnG		With FnG	
	Coeff	t-stat	Coeff	t-stat
Redemption rate$_{t-1}$ (%)	0.1729	5.67	0.1662	5.44
Min discount$_{t-1}$ (%)	-0.0563^{***}	−3.58	-0.0515^{***}	−3.26
FnG$_{t-1}$ (0–100)			0.0006^{**}	2.52
Constant	0.0979	11.04	0.0713	5.16
R^2	0.0475		0.0532	
N	1,052		1,051	

Table 5. What Drives the stETH Discount? Model Parameters and Price Discount. The dependent variable is the minimum daily stETH discount (%); more negative values indicate deeper discounts. A positive coefficient means the variable is associated with a *narrower* (less negative) discount; a negative coefficient means a *deeper* discount. All specifications include one lag of the discount as a control for autocorrelation, consistent with the predictive regressions in Table 4 and Table 3. Columns (1)–(4) are univariate in the main regressor. Column (5) is a multiple regression excluding LTV, which is omitted from the multiple regression due to high multicollinearity with APR (corr = −0.85). The Fear & Greed Index proxies for risk aversion; Buffer for the treasury T; APR for the staking reward r; LTV for external liquidity ρ (higher LTV = lower haircut = more external liquidity). Sample: May 2023–April 2026 (daily, $N = 1047$). ***, **, *: significance at 1%, 5%, 10%.

	Univariate				Multiple
	(1)	(2)	(3)	(4)	(5)
Discount_{t-1}	0.3608^{***}	0.3642^{***}	0.2838^{***}	0.3289^{***}	0.2808^{***}
	(12.51)	(12.68)	(9.54)	(11.25)	(9.44)
Fear & Greed (0–100)	-0.0013^{***}				-0.0005
	(−3.04)				(−1.27)
Buffer (% of pooled)		-0.1942^{***}			-0.1002
		(−2.96)			(−1.55)
APR, 30d MA (%)			-0.1211^{***}		-0.1127^{***}
			(−8.37)		(−7.48)
LTV (%)				0.0142^{***}	
				(5.82)	
Constant	-0.2121^{***}	-0.2617^{***}	0.0778^{*}	-1.3697^{***}	0.0849^{*}
R^2	0.1463	0.1459	0.1929	0.1658	0.1961
N	1,047	1,047	1,047	1,047	1,047

Table 6. Simulated run probabilities: traditional vs. liquid staking. Each row solves the equilibrium (g^*, θ^*) for both traditional staking (Proposition 1) and liquid staking (Theorem 1 with bivariate-normal queue probability, price-based resources as in equation (13), and early indifference $E[p|g] = V(g;\theta^*)$). Baseline: $T = 0.0008$, $r = 0.0322$, $\eta = 0.05$, $\rho = 0.10$, $\beta = 0.30$, $\psi = 0.10$, $\bar{\theta} = 0.045$, $\tau_\theta = 10$, $\tau_s = 100$, $b = 0.10$, $\gamma = 1.30$. Run probabilities are *liquidity-only*: $\Phi((\theta^* - \bar{\theta})/\sigma_\theta)$. The solvency probability $\Pr(\tilde{\theta} < \theta^{sol})$ is the same for both protocols and is not reported.

	Traditional		Liquid		
	θ^*_{tr}	Pr(liq. run)	θ^*_{liq}	Pr(liq. run)	Gap
Panel A: vary γ (market friction)					
$\gamma = 1.01$	+0.022	47.1%	−0.351	10.5%	36.6pp
$\gamma = 1.10$	+0.022	47.1%	−0.304	13.5%	33.6pp
$\gamma = 1.30$	+0.022	47.1%	−0.229	19.3%	**27.8**pp
$\gamma = 2.00$	+0.022	47.1%	−0.096	32.7%	14.4pp
$\gamma = 3.00$	+0.022	47.1%	−0.017	42.2%	4.9pp
$\gamma = 5.00$	+0.022	47.1%	+0.048	50.4%	−3.3pp
Panel B: vary λ (forced sellers, at $\gamma = 1.30$)					
$\lambda = 0.06\%$	+0.023	47.2%	−0.228	19.4%	27.9pp
$\lambda = 0.16\%$	+0.022	47.1%	−0.229	19.3%	**27.8**pp
$\lambda = 1.0\%$	+0.014	46.1%	−0.232	19.0%	27.0pp
$\lambda = 2.0\%$	+0.003	44.8%	−0.236	18.7%	26.1pp
$\lambda = 4.0\%$	−0.018	42.1%	−0.245	18.0%	24.2pp
$\lambda = 6.5\%$	−0.046	38.7%	−0.256	17.1%	21.6pp
Panel C: vary τ_s (signal precision, at $\gamma = 1.30$)					
$\tau_s = 20$	+0.025	47.5%	−0.201	21.9%	25.6pp
$\tau_s = 50$	+0.025	47.5%	−0.217	20.4%	27.1pp
$\tau_s = 100$	+0.022	47.1%	−0.229	19.3%	**27.8**pp
$\tau_s = 200$	+0.016	46.3%	−0.240	18.4%	27.9pp
$\tau_s = 500$	−0.001	44.3%	−0.251	17.5%	26.8pp

A Proofs

A.1 Proof of Proposition 1

Proof. In a monotone equilibrium, strategic token holders choose to withdraw their tokens whenever their private signal about the fundamental value is sufficiently low or $\tilde{s}_i \leq g$, or equivalently, $\sigma_s \tilde{\epsilon}_i \leq g - \tilde{\theta}$. Thus, strategic aggregate withdrawal/selling is:

$$\kappa \Phi^{-1}\left(A(s_i)\right) = \kappa \frac{g - \tilde{\theta}}{\sigma_s}. \tag{38}$$

Step 1: Liquidity threshold and the noise coefficient c_{tr}. The protocol suffers liquidity failure when withdrawal demand exceeds available resources:

$$\underbrace{\lambda + \kappa \frac{g - \theta}{\sigma_s}}_{\text{Withdrawal demand}} > \underbrace{T + \rho(I - T)(\theta + \sigma_l \epsilon_l + 1)}_{\text{Treasury + external liquidity}}. \tag{39}$$

The left side is decreasing in θ (better fundamentals $\Rightarrow$ fewer strategic withdrawals) with slope $-\kappa/\sigma_s$. The right side is increasing in θ (better fundamentals $\Rightarrow$ more external liquidity) with slope $\rho(I - T)$. The two sides cross at a critical θ that depends on the realization of ϵ_l.

To find this crossing point, we move all terms involving θ and ϵ_l to one side. Moving withdrawal demand to the right and resources to the left:

$$\begin{aligned} T + \rho(I - T)(\theta + \sigma_l \epsilon_l + 1) &> \lambda + \kappa \frac{g - \theta}{\sigma_s} \\ T + \rho(I - T)\theta + \rho(I - T)\sigma_l \epsilon_l + \rho(I - T) &> \lambda + \frac{\kappa g}{\sigma_s} - \frac{\kappa \theta}{\sigma_s}. \end{aligned}$$

Collecting all θ terms on the left and all ϵ_l-free constants on the right:

$$\theta \underbrace{\left[\frac{\kappa}{\sigma_s} + \rho(I - T)\right]}_{\text{total sensitivity to } \theta} + \rho(I - T)\sigma_l \epsilon_l > \frac{\kappa g}{\sigma_s} + \lambda - T - \rho(I - T).$$

The term in brackets, $\kappa/\sigma_s + \rho(I - T)$, is the total sensitivity of the protocol's net position to θ: a unit increase in θ both reduces withdrawals (by κ/σ_s) and increases external liquidity (by $\rho(I - T)$). Dividing both sides by this total sensitivity:

$$\theta + \frac{\rho(I - T)\sigma_l}{\kappa/\sigma_s + \rho(I - T)} \epsilon_l > \frac{g + \frac{\sigma_s}{\kappa}[\lambda - T - \rho(I - T)]}{1 + \frac{\sigma_s}{\kappa}\rho(I - T)}.$$

Multiplying through by σ_s/σ_s in the ϵ_l coefficient to simplify:

$$\theta + \underbrace{\frac{\rho(I-T)\sigma_l\sigma_s}{\kappa + \sigma_s\rho(I-T)}}_{\equiv c_{tr}} \epsilon_l > \underbrace{\frac{g + \frac{\sigma_s}{\kappa}[\lambda - T - \rho(I-T)]}{1 + \frac{\sigma_s}{\kappa}\rho(I-T)}}_{\equiv \theta^*_{tr}}.$$

The protocol *survives* the liquidity test if and only if $\theta + c_{tr}\epsilon_l > \theta^*_{tr}$, or equivalently, the composite variable

$$\tilde{v} \equiv \tilde{\theta} + c_{tr}\tilde{\epsilon}_l$$

exceeds θ^*_{tr}. The composite $\tilde{v}$ is the "effective fundamental" from the protocol's liquidity perspective: it combines the true fundamental $\tilde{\theta}$ with the noise in the external provider's assessment, weighted by c_{tr}.

The coefficient c_{tr} has a transparent interpretation. The numerator $\rho(I-T)\sigma_l\sigma_s$ is the product of three terms: the external borrowing rate $\rho(I-T)$ (how much external liquidity responds to the provider's signal), the provider's noise σ_l, and the private signal noise σ_s. The denominator $\kappa + \sigma_s\rho(I-T)$ is the total sensitivity of the net position to θ (in σ_s-scaled units). Intuitively, c_{tr} is large when external borrowing is important ($\rho(I-T)$ large), the provider's signal is noisy (σ_l large), or private signals are imprecise (σ_s large). When $\sigma_l = 0$ (the provider observes θ perfectly), $c_{tr} = 0$ and $\tilde{v} = \tilde{\theta}$: liquidity failure is determined by fundamentals alone.

Taking the binding constraint in expectation over ϵ_l (since $E[\epsilon_l] = 0$), the failure threshold θ^*_{tr} satisfies:

$$\theta^*_{tr} = \frac{g + \frac{\sigma_s}{\kappa}\left[\lambda - T - \rho(I-T)\right]}{1 + \frac{\sigma_s}{\kappa}\rho(I-T)}, \tag{40}$$

which establishes condition (i) of the Proposition. Note that θ^*_{tr} itself does not depend on σ_l: the noise affects whether the protocol fails for a given θ (through $\tilde{v}$), but the expected threshold is determined by the deterministic balance of withdrawals and resources.

Step 2: Four-state indifference condition. The staker who receives the cutoff signal $s_i = g$ must be indifferent between withdrawing and holding. Conditional on $s_i = g$, the staker faces four possible states $(\delta_{liq}, \delta_{sol})$. Without loss of generality, set $\bar{\theta} = 0$.

The staker's posterior on $\tilde{\theta}$ is $\tilde{\theta}|g \sim N(\mu_g, 1/(\tau_\theta + \tau_s))$ where $\mu_g = \tau_s g/(\tau_\theta + \tau_s)$. Since $\tilde{\epsilon}_l$ is independent of $\tilde{\theta}$ and s_i, the composite variable $\tilde{v} = \tilde{\theta} + c_{tr}\tilde{\epsilon}_l$ has conditional distribution $\tilde{v}|g \sim N(\mu_g, \sigma_v^2)$ where $\sigma_v^2 = 1/(\tau_\theta + \tau_s) + c_{tr}^2$.

Since $(\tilde{\theta}, \tilde{v})$ are jointly normal conditional on g with correlation $\rho_c = 1/(\sigma_v\sqrt{\tau_\theta + \tau_s})$, the expected payoff difference is:

$$\begin{aligned} E[\pi|g] \;=\; & (r+\eta)\ \Pr(\tilde{v} \geq \theta^*_{tr},\ \tilde{\theta} \geq \theta^{sol} \mid g) \\ & -(1-\eta)\ \Pr(\tilde{v} \geq \theta^*_{tr},\ \tilde{\theta} < \theta^{sol} \mid g) \\ & -\psi\ \Pr(\tilde{v} < \theta^*_{tr} \mid g). \end{aligned}$$

Define the standardized thresholds:

$$a = \sqrt{\tau_\theta + \tau_s}\left(\theta^{sol} - \mu_g\right), \qquad b = \frac{\theta^*_{tr} - \mu_g}{\sigma_v}.$$

Using the bivariate normal CDF Φ_2:

$$\begin{aligned}
\Pr(\tilde{v} \geq \theta^*_{tr},\ \tilde{\theta} \geq \theta^{sol} \mid g) &= \Phi_2(-a, -b; \rho_c), \\
\Pr(\tilde{v} \geq \theta^*_{tr},\ \tilde{\theta} < \theta^{sol} \mid g) &= \Phi(-b) - \Phi_2(-a, -b; \rho_c), \\
\Pr(\tilde{v} < \theta^*_{tr} \mid g) &= \Phi(b).
\end{aligned}$$

Substituting and setting $E[\pi|g] = 0$:

$$(r + \eta)\Phi_2(-a, -b; \rho_c) - (1 - \eta)\left[\Phi(-b) - \Phi_2(-a, -b; \rho_c)\right] - \psi\,\Phi(b) = 0.$$

Collecting terms in Φ_2:

$$\begin{aligned}
(1 + r)\Phi_2(-a, -b; \rho_c) &= (1 - \eta)\Phi(-b) + \psi\,\Phi(b) \\
&= \psi + (1 - \eta - \psi)(1 - \Phi(b)),
\end{aligned}$$

which is equation (8).

Step 3: Existence and uniqueness. We establish the equilibrium as a fixed point. For any candidate θ^*_{tr}, condition (i) gives $g = g(\theta^*_{tr})$ linearly. The indifference condition (8) then determines whether this g is consistent.

Define $\Gamma(\theta^*_{tr}) \equiv g_{\text{indiff}}(\theta^*_{tr}) - g_{\text{liq}}(\theta^*_{tr})$, where $g_{\text{liq}}(\theta^*_{tr})$ inverts (40) and $g_{\text{indiff}}(\theta^*_{tr})$ solves (8).

Monotonicity of g_{indiff}: An increase in θ^*_{tr} expands the liquidity failure region, strictly increasing $\Pr(\tilde{v} < \theta^*_{tr}|g)$ and decreasing $\Pr(\tilde{v} \geq \theta^*_{tr}, \tilde{\theta} \geq \theta^{sol}|g)$. This lowers $E[\pi|g]$, so the indifference signal g must rise (the staker requires a better signal to stay). Thus g_{indiff} is strictly increasing in θ^*_{tr}.

Monotonicity of g_{liq}: From (40), $g_{\text{liq}}(\theta^*_{tr}) = \theta^*_{tr}[1 + \sigma_s\rho(I - T)/\kappa] - \sigma_s[\lambda - T - \rho(I - T)]/\kappa$, which is linear and strictly increasing with slope $1 + \sigma_s\rho(I - T)/\kappa > 1$.

Boundary conditions: As $\theta^*_{tr} \to -\infty$, liquidity failure is impossible, so $g_{\text{indiff}} \to -\infty$ while $g_{\text{liq}} \to -\infty$; the slope of g_{liq} exceeds that of g_{indiff} (since the former has slope > 1 while the latter's slope is bounded above by $(\tau_\theta + \tau_s)/\tau_s < 1 + \sigma_s\rho(I - T)/\kappa$ under the interiority condition), so $\Gamma < 0$ for sufficiently negative θ^*_{tr}. As $\theta^*_{tr} \to +\infty$, failure is near-certain and $g_{\text{indiff}} \to +\infty$, but g_{indiff} grows slower than g_{liq}, so $\Gamma > 0$.

By continuity and monotonicity, Γ crosses zero exactly once, establishing existence and uniqueness of the equilibrium pair (g, θ^*_{tr}).

Verification: Since $E[\pi|s_i]$ is strictly increasing in s_i (by first-order stochastic dominance of the posterior), stakers with $s_i < g$ strictly prefer to withdraw and those with $s_i > g$ strictly prefer

to hold. Thus, the monotone strategy is optimal.

Step 4: Recovery of the closed form when solvency is non-binding. When $\theta^{sol} \ll \theta^*_{tr}$, we have $a \to -\infty$ and $\Phi_2(-a, -b; \rho_c) \to \Phi(-b)$. The indifference condition reduces to $(1+r)\Phi(-b) = \psi + (1-\eta-\psi)\Phi(-b)$, which gives $\Phi(b) = (r+\eta)/(r+\eta+\psi)$. Inverting: $g = \frac{\tau_\theta+\tau_s}{\tau_s}\theta^*_{tr} - \frac{(\tau_\theta+\tau_s)\sigma_v}{\tau_s}\Phi^{-1}\Big(\frac{r+\eta}{r+\eta+\psi}\Big)$. Substituting into (40) and simplifying (using the algebra in the non-binding case), we obtain the closed form (11). ■

A.2 Proof of Corollary 1

Proof. The equilibrium (g, θ^*_{tr}) is determined by two equations: the liquidity condition (40) and the indifference condition (8). Applying the Implicit Function Theorem to this system with respect to T:

When solvency is non-binding ($\theta^{sol} \ll \theta^*_{tr}$), the indifference reduces to the univariate form, and the derivative of the failure threshold with respect to T is:

$$\frac{\partial \theta^*_{tr}}{\partial T} = \frac{\kappa\left[\sigma_s\tau_\theta(1-\rho) - \rho\left(1 + \sigma_s(\tau_\theta+\tau_s)\sigma_v\Phi^{-1}\left(\frac{r+\eta}{r+\eta+\psi}\right)\right)\right]}{\left[\rho(I-T) - \kappa\sigma_s\tau_\theta\right]^2}. \tag{41}$$

The sign is determined by the numerator $\mathcal{N} = \sigma_s\tau_\theta(1-\rho) - \rho\Psi$ where $\Psi = 1 + \sigma_s(\tau_\theta + \tau_s)\sigma_v\Phi^{-1}\left(\frac{r+\eta}{r+\eta+\psi}\right) > 0$. The threshold $\bar{\rho}$ follows from setting $\mathcal{N} = 0$.

When solvency is binding (θ^{sol} near θ^*_{tr}), there is an additional destabilizing effect: a larger T raises $\theta^{sol} = Ir/(I-T)$, expanding the insolvency region. This increases $\Pr(\delta_{liq} = s, \delta_{sol} = f|g)$, which raises g through the indifference condition, further increasing θ^*_{tr}. The treasury is therefore destabilizing over a strictly wider parameter range than in the solvency-non-binding case. ■

A.3 Proof of Lemma 1

Proof. The function $F(\theta)$ is given by:

$$F(\theta) = \underbrace{\kappa\int_{N(\theta)}^{\infty}\int_{-\infty}^{M(\theta)} g(\epsilon,\omega)\,d\omega\,d\epsilon}_{\text{Redemption Demand}} - \underbrace{\left[T + (I-T)(\rho+\beta)(1+\theta) + (I-T)\rho\sigma_l\tilde{\epsilon}_l + (I-T)\beta\frac{b\sigma_m\tilde{\epsilon}_m + \sigma_\xi\tilde{\xi}}{b+\kappa/\sigma_s}\right]}_{\text{Available Funding}}, \tag{42}$$

where the integration limits depend on θ:

$$N(\theta) = \frac{g-\theta}{\sigma_s}, \tag{43}$$

$$M(\theta) = x^* - (I-T)\left[\frac{\rho\tau_\theta\bar{\theta}}{\hat{\tau}} + \left(\frac{\rho(\tau_s+\tau_p)}{\hat{\tau}} + \beta\right)\theta + (\rho+\beta)\right]. \tag{44}$$

The protocol is stable if $F(\theta) < 0$. We observe that when $\theta \to \infty$, redemption demand vanishes and funding grows, so $F(\theta) < 0$. Conversely, when $\theta \to -\infty$, $F(\theta) > 0$. By continuity, a solution exists. To ensure uniqueness and stability, we require $F'(\theta) < 0$.

We differentiate $F(\theta)$ using the Leibniz Integral Rule for variable limits:

$$\frac{\partial F}{\partial \theta} = \kappa \frac{\partial}{\partial \theta} \left(\int_{N(\theta)}^{\infty} \int_{-\infty}^{M(\theta)} g(\epsilon, \omega)\, d\omega\, d\epsilon \right) - K, \tag{45}$$

where $K = (I - T)(\rho + \beta)$ represents the marginal increase in funding.

Applying Leibniz's rule:

$$\frac{\partial F}{\partial \theta} = \kappa \left[-N'(\theta) \int_{-\infty}^{M(\theta)} g(N(\theta), \omega)\, d\omega + M'(\theta) \int_{N(\theta)}^{\infty} g(\epsilon, M(\theta))\, d\epsilon \right] - K. \tag{46}$$

Substituting the slopes $N'(\theta) = -1/\sigma_s$ and $M'(\theta) = -(I - T)\left(\frac{\rho(\tau_s + \tau_p)}{\hat{\tau}} + \beta\right)$:

$$\frac{\partial F}{\partial \theta} = \kappa \left[\frac{1}{\sigma_s} \int_{-\infty}^{M(\theta)} g(N(\theta), \omega)\, d\omega - (I - T) \left(\frac{\rho(\tau_s + \tau_p)}{\hat{\tau}} + \beta \right) \int_{N(\theta)}^{\infty} g(\epsilon, M(\theta))\, d\epsilon \right] - K. \tag{47}$$

To define the marginal densities correctly, we define the conditional means of the random variables. Since ω is a linear combination of ϵ and ξ, they are jointly normal. Let $\mu_{\omega|N} = E[\omega|\epsilon = N(\theta)]$ and $\mu_{\epsilon|M} = E[\epsilon|\omega = M(\theta)]$. The integral terms can be expressed using the conditional cumulative distribution functions:

$$\begin{aligned} \int_{-\infty}^{M(\theta)} g(N(\theta), \omega)\, d\omega &= \phi(N(\theta)) \Phi\left(\frac{M(\theta) - \mu_{\omega|N}}{\sigma_{\omega|N}} \right), \\ \int_{N(\theta)}^{\infty} g(\epsilon, M(\theta))\, d\epsilon &= \hat{\phi}(M(\theta)) \Phi\left(\frac{\mu_{\epsilon|M} - N(\theta)}{\sigma_{\epsilon|M}} \right). \end{aligned}$$

Thus, the condition for stability ($F'(\theta) < 0$) is:

$$\underbrace{\frac{\kappa}{\sigma_s} \phi(N(\theta)) \Phi\left(\frac{M(\theta) - \mu_{\omega|N}}{\sigma_{\omega|N}} \right)}_{\text{Destabilizing Force}} < \underbrace{\kappa (I - T) \left(\frac{\rho(\tau_s + \tau_p)}{\hat{\tau}} + \beta \right) \hat{\phi}(M(\theta)) \Phi\left(\frac{\mu_{\epsilon|M} - N(\theta)}{\sigma_{\epsilon|M}} \right) + K}_{\text{Stabilizing Forces}}. \tag{48}$$

This inequality holds when the noise in private signals, σ_s, is sufficiently large. Specifically, as $\sigma_s \to \infty$, the term $\frac{1}{\sigma_s} \to 0$, causing the destabilizing force (Left-Hand Side) to vanish, while the stabilizing funding term K remains positive.

It is immediate that as $\theta \to -\infty$, $F(\theta)$ approaches $+\infty$. As $\theta \to +\infty$, $F(\theta)$ approaches $-\infty$. When σ_s is sufficiently large, the destabilizing term on the LHS vanishes (since it is proportional

to $1/\sigma_s$) and the stabilizing terms on the RHS remain strictly positive. Thus, the inequality holds, and $F(\theta)$ is strictly decreasing over its entire domain.

A strictly decreasing function can cross zero at most once. Therefore, there exists a unique fundamental threshold θ^* such that $F(\theta) > 0$ (failure) for all $\theta < \theta^*$ and $F(\theta) < 0$ (survival) for all $\theta > \theta^*$. ■

A.4 Proof of Lemma 2

Proof. From Lemma 1, the threshold θ^* is defined implicitly by the condition that the net stress is zero, $F(\theta^*, g) = 0$. Since F is composed of smooth Gaussian functions, $\theta^*(g)$ is differentiable. By the Implicit Function Theorem:

$$\frac{d\theta^*}{dg} = -\frac{\partial F/\partial g}{\partial F/\partial \theta^*}. \tag{49}$$

As for the denominator ($\partial F/\partial \theta^*$), Lemma 1 establishes that for sufficiently large σ_s, the Net Stress is strictly decreasing in fundamentals: $\frac{\partial F}{\partial \theta^*} < 0$.

For the numerator ($\partial F/\partial g$), we differentiate the Net Stress function with respect to the panic threshold g:

$$F(\theta, g) = \kappa \underbrace{\int_{N(\theta,g)}^{\infty} (\dots)}_{\text{Queue Demand}} - \underbrace{[\cdots + \beta y(P(g))]}_{\text{Resources}}. \tag{50}$$

An increase in g affects F through the Queue term. A higher g raises the integration limit $N(\theta, g) = \frac{g-\theta}{\sigma_s}$. This reduces the probability mass of stakers entering the queue (as more stakers exit via the market). Note that the extracted price signal $y(P) = \tilde{\theta} + (b\sigma_m\tilde{\epsilon}_m + \sigma_\xi\tilde{\xi})/(b + \kappa/\sigma_s)$ is independent of g: the g and λ terms cancel exactly when substituting the equilibrium price into the extraction function (20). Therefore $\frac{\partial \text{Resources}}{\partial g} = 0$, and $\frac{\partial F}{\partial g} < 0$ holds entirely because the queue demand strictly decreases with g.

The slope is the ratio of two negative numbers, with a leading negative sign:

$$\frac{d\theta^*}{dg} < 0. \tag{51}$$

Thus, $\theta^*(g)$ is strictly decreasing. ■

A.5 Proof of Proposition 2

Proof. Let $F(s) = E[p|s] - V(s)$ denote the net incentive to sell in the secondary market. We characterize the components $E[p|s]$ and $V(s)$ explicitly to establish the existence and uniqueness of the threshold g.

First we derive expected price $E[p|s]$. The market price is given by $p = \frac{1+r}{\gamma}\Phi(Z)$, where

$Z = \left(b + \frac{\kappa}{\sigma_s}\right) y(P) - \frac{\kappa g}{\sigma_s} - \lambda$. Conditional on the private signal s, the variable Z is normally distributed:

$$Z \mid s \sim \mathcal{N}(\mu_Z(s), \sigma_Z^2), \tag{52}$$

where the conditional mean and variance are derived from the signal structure:

$$\mu_Z(s) = \left(b + \frac{\kappa}{\sigma_s}\right) \frac{\tau_s s}{\tau_\theta + \tau_s} - \frac{\kappa g}{\sigma_s} - \lambda, \tag{53}$$

$$\sigma_Z^2 = \left(b + \frac{\kappa}{\sigma_s}\right)^2 \frac{1}{\tau_\theta + \tau_s} + b^2 \sigma_m^2 + \sigma_\xi^2. \tag{54}$$

Using the identity $E[\Phi(Z)] = \Phi\left(\frac{\mu_Z}{\sqrt{1+\sigma_Z^2}}\right)$, the expected price is:

$$E[p|s] = \frac{1+r}{\gamma} \Phi\left(\frac{\mu_Z(s)}{\sqrt{1+\sigma_Z^2}}\right). \tag{55}$$

Since $\mu_Z(s)$ is linear and increasing in s, $E[p|s]$ is strictly increasing in s.

Second, we derive the retention value $V(s)$. The retention value $V(s)$ is the expected payoff from either queuing or holding, depending on the realization of fundamentals θ and the noise ν. We define two integration boundaries to simplify the expression.

Let $A(s)$ be the standardized threshold for the fundamental θ^* (Survival/Failure boundary):

$$A(s) = \frac{\theta^* - E[\theta|s]}{\sqrt{Var[\theta|s]}} = \sqrt{\tau_\theta + \tau_s} \left(\theta^* - \frac{\tau_s s}{\tau_\theta + \tau_s}\right). \tag{56}$$

Let $D(s)$ be the threshold for the queueing decision variable ν (Queue/Hold boundary):

$$D(s) = \frac{\hat{\tau}}{\tau_p} \left(\theta^* - \frac{\tau_s s}{\hat{\tau}}\right) - \frac{\sqrt{\hat{\tau}}}{\tau_p} \Phi^{-1}\left(\frac{r+\eta}{r+\eta+\psi}\right). \tag{57}$$

Using these limits, the expected retention value $V(s)$ can be written compactly as an integral over the posterior densities:

$$V(s) = \underbrace{\int_{-\infty}^{A(s)} \left[\Phi(D(s)) \frac{T}{I} + (1 - \Phi(D(s))) \left(\frac{T}{I} - \psi\right)\right] d\Phi(\tilde{\theta}_s)}_{\text{Failure Region } (\theta < \theta^*)}$$
$$+ \underbrace{\int_{A(s)}^{\infty} \left[\Phi(D(s))(1-\eta) + (1 - \Phi(D(s)))(1+r)\right] d\Phi(\tilde{\theta}_s)}_{\text{Survival Region } (\theta > \theta^*)}, \tag{58}$$

where $\Phi(\tilde{\theta}_s)$ denotes the standard normal CDF with respect to the standardized fundamental.

To show the uniqueness of g, we examine the slope of $F(s) = E[p|s] - V(s)$. As s increases, the posterior distribution of θ shifts to the right (First Order Stochastic Dominance). There are two effects. First, the probability of the "Survival Region" (payoff $\approx 1 + r$) increases relative to the "Failure Region" (payoff $\approx T/I$). Second, the threshold $D(s)$ decreases ($D'(s) < 0$), meaning stakers are less likely to panic-queue and more likely to hold to maturity. Both effects imply that $V(s)$ is strictly increasing in s. While $E[p|s]$ is also increasing in s, for sufficiently large noise σ_s, the marginal gain in retention value (driven by the binary survival outcome) dominates the marginal gain in Expected Price (driven by linear aggregation). Thus $F'(s) < 0$. Given the boundary conditions established previously ($F(-\infty) > 0$ and $F(+\infty) < 0$), there exists a unique root g. ■

A.6 Proof of Theorem 1

Proof. We establish uniqueness by characterizing the equilibrium as the intersection of two reaction functions in the (θ^*, g) plane.

First, we characterize the staker's reaction function $\mathcal{R}_S : \theta^* \to g$. From Proposition 2, for any fixed failure threshold θ^*, the marginal staker condition is given by $F_{staker}(g; \theta^*) = E[p|g] - V(g; \theta^*) = 0$. By the Implicit Function Theorem, the slope is:

$$\frac{dg}{d\theta^*} = -\frac{\partial F_{staker}/\partial \theta^*}{\partial F_{staker}/\partial g}. \tag{59}$$

As shown in Proposition 2, $\frac{\partial F_{staker}}{\partial g} < 0$. We now sign $\frac{\partial F_{staker}}{\partial \theta^*}$. An increase in the failure threshold θ^* expands the failure region ($\theta < \theta^*$), strictly decreasing the retention value $V(s)$. Thus, the net incentive to sell increases: $\frac{\partial F_{staker}}{\partial \theta^*} > 0$. Consequently, the staker's best response is strictly increasing:

$$\frac{dg}{d\theta^*} > 0. \tag{60}$$

Intuitively, if the protocol is more fragile (higher θ^*), stakers require a higher private signal to retain their tokens.

Second, we characterize the protocol's reaction function $\mathcal{R}_P : g \to \theta^*$. From Lemma 2, the fundamental threshold $\theta^*(g)$ is defined by the condition that the Net Stress on the protocol is zero: $F(\theta^*, g) = 0$.

$$F(\theta^*, g) = \underbrace{\kappa \Pr(s > g, z < z^* \mid \theta^*)}_{\text{Queue Demand}} - \underbrace{[T + (I - T)\left[\rho(1 + \theta^*) + \beta(y(P(g)) + 1)\right]]}_{\text{Resources + Liquidity}}. \tag{61}$$

The slope is determined by the Implicit Function Theorem:

$$\frac{d\theta^*}{dg} = -\frac{\partial F/\partial g}{\partial F/\partial \theta^*} < 0 \tag{62}$$

This confirms that in the presence of a liquid market, higher panic selling ($g \uparrow$) paradoxically stabilizes the protocol by reducing the immediate liability of the redemption queue.

To show existence and uniqueness, we define the equilibrium as a fixed point in θ^*. Let $g(\theta^*)$ be the Staker's reaction function and $\Theta(g)$ be the Protocol's reaction function . Substitute the staker's response into the protocol's response to define the composite function $\Gamma(\theta^*)$:

$$\Gamma(\theta^*) = \Theta(g(\theta^*)) - \theta^*. \tag{63}$$

An equilibrium exists if and only if $\Gamma(\theta^*) = 0$. we first show $\Gamma(\theta^*)$ is monotone. We differentiate Γ with respect to θ^*:

$$\frac{d\Gamma}{d\theta^*} = \underbrace{\frac{d\Theta}{dg}}_{(-)} \cdot \underbrace{\frac{dg}{d\theta^*}}_{(+)} - 1. \tag{64}$$

From Step 1, $dg/d\theta^* > 0$. From Step 2, $d\Theta/dg < 0$. Thus, the product is negative. Consequently, $\frac{d\Gamma}{d\theta^*} < -1$, implying $\Gamma(\theta^*)$ is strictly decreasing.

We then check the boundary conditions. As $\theta^* \to -\infty$ (protocol is ultra-safe), stakers rarely panic ($g \to -\infty$). With no panic, the protocol's actual failure threshold $\Theta(g)$ is high (dominated by queue stress). Thus $\Theta > \theta^*$, implying $\Gamma(-\infty) > 0$. As $\theta^* \to \infty$ (protocol is ultra-fragile), stakers always panic ($g \to \infty$). With high panic, the protocol's actual failure threshold $\Theta(g)$ is low (stress relieved by market). Thus $\Theta < \theta^*$, implying $\Gamma(\infty) < 0$.

Since $\Gamma(\theta^*)$ is continuous, strictly decreasing, and spans positive to negative values, by the Intermediate Value Theorem, there exists a unique θ^* such that $\Gamma(\theta^*) = 0$. The corresponding unique panic threshold is $g^* = g(\theta^*)$. ■

A.7 Proof of Corollary 6

Proof. We apply equation (35) case by case, using $\Lambda > 0$, $\partial\mathcal{R}_P/\partial g < 0$ (Lemma 2), and $\partial G/\partial g < 0$ (Proposition 2).

Part (i): β. The parameter β enters the resources in F scaled by $(I - T)$:

$$\frac{\partial F}{\partial \beta} = -(I - T)y(P).$$

Evaluating at the equilibrium threshold $\theta^* > 0$, $E[y(P)|\theta^*] = \theta^* > 0$, so $\partial F/\partial\beta < 0$ in expectation and the direct effect is stabilizing: $\partial\mathcal{R}_P/\partial\beta < 0$. There is also a small indirect channel: β enters the queue threshold $x = (I - T)[\rho(z + 1) + \beta(y(P) + 1)]$, so an increase in β shifts the late indifference threshold x^*, which raises the retention value $V(g)$ and lowers the panic threshold g^*. Since $\partial\mathcal{R}_P/\partial g < 0$, this indirect channel is weakly destabilizing. However, the indirect effect is second-order (it operates through the late indifference and back through the early indifference), so the direct stabilizing channel dominates. Hence $\frac{d\theta^*}{d\beta} < 0$ robustly, though not exactly as a single-channel result.

Part (ii): γ. The failure condition and the late indifference that pins x^* do not contain γ (since $y(P)$ is defined to remove γ from the price signal, as seen from equation (20)), so $\partial\mathcal{R}_P/\partial\gamma\big|_g = 0$.

However, γ enters the expected market price in the early indifference:

$$\frac{\partial G}{\partial \gamma} = \frac{\partial E[p|g]}{\partial \gamma} = -\frac{1+r}{\gamma^2}\,\Phi\left(\frac{\mu_Z(g)}{\sqrt{1+\sigma_Z^2}}\right) < 0.$$

A higher γ lowers the expected proceeds from selling early, making retention more attractive. The staker threshold therefore shifts:

$$\left.\frac{\partial g}{\partial \gamma}\right|_{\mathcal{R}_S} = -\frac{\partial G/\partial \gamma}{\partial G/\partial g} = -\frac{(-)}{(-)} < 0.$$

Substituting into (35):

$$\frac{d\theta^*}{d\gamma} = \frac{1}{\Lambda}\cdot \underbrace{\frac{\partial \mathcal{R}_P}{\partial g}}_{<0}\cdot \underbrace{\frac{\partial g}{\partial \gamma}\big|_{\mathcal{R}_S}}_{<0} \;>\; 0.$$

The mechanism is: illiquid market $\Rightarrow$ lower selling proceeds $\Rightarrow$ lower panic threshold g^* $\Rightarrow$ weaker liquidity outlet (Lemma 2) $\Rightarrow$ higher θ^*.

Part (iii): r. The late indifference fixes the failure probability at the queue threshold as $\Pr(\theta < \theta^*|x^*) = \frac{r+\eta}{r+\eta+\psi}$, which is increasing in r. Thus x^* falls, shrinking the set of stakers who queue: $\partial F/\partial r < 0$, so the direct effect is stabilizing, $\partial \mathcal{R}_P/\partial r < 0$.

For the indirect effect, r raises both the expected market price ($E[p|g] = \frac{1+r}{\gamma}\Phi(\cdot)$) and the retention value V via the holding payoff $1+r$. Since $\gamma > 1$ (Proposition 2), the market price trades at a discount to the staking reward, so the full-unit increase in the holding payoff typically dominates:

$$\frac{\partial G}{\partial r} = \frac{\partial E[p|g]}{\partial r} - \frac{\partial V}{\partial r} = \frac{1}{\gamma}\Phi(\cdot) - \Pr(\theta > \theta^*, x > x^*|g) < 0,$$

when $\frac{1}{\gamma}\Phi(\cdot) < \Pr(\theta > \theta^*, x > x^*|g)$. In this case $\partial g/\partial r\big|_{\mathcal{R}_S} < 0$, and the indirect term $\frac{\partial \mathcal{R}_P}{\partial g}\cdot \frac{\partial g}{\partial r}\big|_{\mathcal{R}_S} = (-)(-) > 0$ is destabilizing. The two effects work in opposite directions, and the net sign is ambiguous.

Part (iv): σ_ξ. A higher σ_ξ reduces the precision of the price signal, $\tau_p = \tau_s\kappa^2/\sigma_\xi^2$, and lowers the posterior precision $\hat{\tau} = \tau_\theta + \tau_s + \tau_p$. The noise term in $y(P) = \tilde{\theta} - \frac{\sigma_s\sigma_\xi}{\kappa}\tilde{\xi}$ has increased variance, but its conditional mean $E[y(P)|\theta] = \theta$ is unchanged, leaving expected resources unaffected while raising their variance. A less precise posterior shifts the joint distribution underlying both indifference conditions, with conflicting effects on g^*. The two channels have opposing signs; the net effect is ambiguous.

Part (v): T. The direct effect on the Protocol Reaction now reflects the fact that both liquidity channels scale with $(I-T)$. Since expected funding is $T + (I-T)(\rho+\beta)(1+\theta^*)$, we have:

$$\frac{\partial F}{\partial T} = (\rho+\beta)(1+\theta^*) - 1,$$

which is positive (destabilizing) when $(\rho+\beta)(1+\theta^*) > 1$. Compared with traditional staking

(where $\partial F/\partial T = \rho(1+\theta^*) - 1$), the additional $\beta(1+\theta^*)$ term means that the treasury is destabilizing over a wider parameter range: a larger treasury shrinks the invested base $(I-T)$, which now reduces both the fundamental-based and market-price-based channels of external liquidity.

In liquid staking there is also an indirect channel through the Staker Reaction. The failure payoff T/I in the early indifference raises the retention value by $\frac{1}{I}\Pr(\theta < \theta^*|g) > 0$:

$$\frac{\partial G}{\partial T} = -\frac{1}{I}\Pr(\theta < \theta^*|g) < 0,$$

so $\partial g/\partial T\big|_{\mathcal{R}_S} < 0$. Via $\partial \mathcal{R}_P/\partial g < 0$, this indirect term is strictly positive (destabilizing). Hence, the stabilizing effect of a larger treasury is strictly weaker under liquid staking than under traditional staking, and the new formulation amplifies this asymmetry because the market-price channel $\beta y(P)$ also scales with $(I-T)$. ■

B Description of ETH staking and withdrawals

C Detailed Description of the LIDO protocol

B.1 Steps to a withdrawal as a user.

To withdraw a stETH position using the Lido directly a user must interact with the WithdrawalQueueERC721 contract. This contract is responsible for queuing withdrawal requests, and finalizing withdrawal requests following daily oracle reports. This contract is also responsible for minting ERC-721 unstETH NFTs that correspond to withdrawal requests, and burning these NFTs upon the finalization of the withdrawal process. WithdrawalQueueERC721 contract extends WithdrawalQueueBase.sol and WithdrawalQueue.sol. These abstract classes handle the main withdrawal queuing, fulfillment logic,support for larger withdrawals, and conversion of wstETH respectively.

The process from user finalization request to finalization begins by a user initiated withdrawal, which is confirmed by the approval of stETH to the WithdrawalQueueERC721 contract or signing a ERC-2612 Permit. The contract will strictly reserve an amount of eth that corresponds to the requested amount of stETH to be withdrawn, this is because Lido withdrawals can only have nominal or discounted finalization rates. In other words queued stETH cannot accrue rewards and positively increase eth exchange rate, but it still remains liable to slashings that may be socialized during the queueing process. The request is now placed in the queue, and the corresponding ustETH NFT is minted.

Lido relies on off-chain oracles run by each oracle node operator to do the computationally expensive process of updating protocol state and delivering a finalized decision on withdrawals. In this context, an oracle refers to a system that acts as a bridge between the blockchain and external data sources or computational processes. The oracles used by Lido are operated by trusted Node Operators who are responsible for gathering data from the Ethereum network, processing it off-chain and submitting the results back to Lido in the form of a report. These off-chain oracles send a report to Lido every 225 Ethereum Epochs (24 hrs) which then has to pass the consensus committee for further processing. Following a consensus on the report, the report is handled by the Accounting Oracle which triggers a core state update and the fulfillment of processed withdrawals. The Accounting Oracle is a specialized on-chain contract responsible for executing the daily core state updates that keep Lido protocol synchronized with the most recent data.

The claim is now made available for the user, which when triggered will send the finalization amount of eth to the user and burn the unstETH NFT that corresponds with withdrawal, unstETH shares will have been burned at this point as well.

B.2 Fulfilment Execution.

In order to calculate the finalization batches, which determine the finalization decision that will be sent to Lido at the end of the current frame, the off-chain oracle must consider the following three factors :

1. **Report fulfillment cutoff:** When the protocol is operating in standard mode, referred to as turbo mode, there is a single cutoff point for withdrawals known as "new requests border". This cutoff is constant and is determined by the following:

$$\text{ref\_epoch} - \text{NEW\_REQUESTS\_BORDER}.$$

2. **How much ETH is available to finalize requests:** The off-chain oracles calculate how much ether is available for processing withdrawals in the current report in the following way:

$$\text{eth\_amount} = \text{withdrawal\_vault\_balance} + \text{el\_rewards\_vault\_balance} + \text{reserved\_buffer}.$$

The *reserved_buffer* is determined by taking the minimum amount of ether in the buffer to fulfill the unfinalized requests that pertain to the ref_slot. This is how Lido optimizes buffered ether, which is also a major part in determining the buffer size daily.

3. **The finalization rate for each of the batches that will be finalized:** The off-chain oracles determine this rate based on a simulated post-rebase state of Lido, which occurs after the core state update following the report. This ensures the that the rate is calculated using the most up to date information:

$$\text{finalization\_share\_rate} = \frac{\text{simulation.post\_total\_pooled\_ether}}{\text{simulation.post\_total\_shares}}.$$

Lido will only provide users with a nominal rate (1:1, as it was when the withdrawal request was made) or a discounted rate that applies to queued withdrawals who's initial value exceeds the finalization share rate. This up-to-date *share_rate* is derived from the simulation to ensure the protocol can fulfill as many withdrawals within current frame and socialize a possible negative rebase to the queued stETH. If protocol level losses occur, a discount may be applied to withdrawals above the rate, ensuring a finalized batch rate is equal to or lower than the nominal rate in standard cases.

Now the off-chain oracle is ready to start creating the finalization batches by iterating through all the withdrawal requests. The batches are decided based on whether requests were made at the same time, or have a similar withdrawal rate. **Nominal withdrawals** ($\leq$ max_share_rate) are batched and processed at the initial withdrawal rate. **Discounted withdrawals** ($>$ max_share_rate) are adjusted to a rate that accounts for losses in the following way:

$$\text{ethToFinalize} = \text{shares} \times \text{maxShareRate}.$$

Batches are either nominal or discounted. Once the *eth_amount* that's available has been ex-

hausted by the batching, the process completes. Unfinalized withdrawals will carry over to the next report when there is more liquidity in the protocol to fulfill them.

B.3 Lido Buffer Size Information

The Lido buffer is where unallocated Ethereum in the protocol is held. The buffer is made up of recently deposited but unallocated ETH, the contents of the EL Rewards Vault, and the Withdrawals Vault (CL validator withdrawals). Since the ETH in the buffer is unproductive, Lido tries to keep its balance as low as possible, as would be expected. The size of the buffer is not mediated by an external entity with special privilege but is instead determined by user demand. This demand is reflected in off-chain oracle reports and buffer resources allocated by Lido's suite of contracts. The two processes on Lido that determine the buffer size are as follows:

- **Withdrawal Fulfillment**: Lido's first priority when it comes to allocating buffer resources is withdrawal fulfillment. Lido reserves as much ETH as needed to fulfill unfinalized withdrawals.
- **Restaking Mechanism**: If there is still ETH left in the buffer following withdrawal processing, it is restaked to the beacon chain by the Staking Router under certain conditions.

The depositor bot, which is responsible for handling the restaking/deposit mechanism, first checks to see if the amount in the buffer is greater than the unfinalized stETH. If not, then the ETH in the buffer is held for the next report to help fulfill those withdrawals. However, if there is enough ETH in the buffer to justify the gas price of reallocating it to the beacon chain, the depositor bot performs a batch deposit and allocates ETH in the following ways:

- The curated set of validators is called the **NodeOperatorsRegistry**.
- The **Staking Router** will receive some of this ETH to be restaked in order to allocate stake to the **Community** and **DVT Staking Modules**.

B.4 Briefly, conditions for Validator exit

The mechanics of exiting Lido validators, much like processing user withdrawals, relies on reports from off-chain oracles. These reports are decoupled from the withdrawal life cycle, as they are sent to Lido every 8 hours, and the report is processed by another special on chain oracle called ValidatorExitBusOracle.sol. .These off-chain oracles are responsible for determining both when pre-signed validators are to exit and if they are to exit. The decision to exit validators is made if unfinalized stETH exceeds ETH liquidity in Lido. The algorithm that decides which validators to exit can be found in the links below. At the time of writing, it takes validators approximately 9.3 days + and additional 256 epochs (28hrs) to exit the Consensus Layer (CL) and have their funds deposited into the Withdrawals Vault, which is when funds are available for withdrawal fulfilment.

On the user side, wait times vary depending on the withdrawal demand and the amount of liquidity in the Lido buffer to fulfill that demand. As the number of validators on the Ethereum

beacon chain increases, so does the sweep delay. To account for this, Lido's use of a semi-automatic validator exit process helps ensure that Node Operators meet pre-signed validator exit quotas to keep up with the demand. This ensures that users can withdraw much earlier than the usual CL wait times due to there being constant flow of liquidity for every report due to these semi-automatic withdrawals. Additionally, as mentioned in the withdrawals section, users cannot accrue rewards while in the queue, to prevent potential exploitation by sophisticated actors, but they remain liable to slashing for similar reasons.

B.5 Profit Sharing

Lido has two primary ways of generating income: staking on the Beacon Chain and earning rewards for block attestations on the Proof of Stake (POS) network, and earning Execution Layer (EL) rewards such as MEV and transaction fees following the merge. In each oracle report, funds in these vaults are included in the rewards calculation as follows:

$$\begin{aligned}\text{rewards} = &(\text{beacon balance new} - \text{beacon balance old}) \\ &- (\text{validators} \times 32\ \text{ETH}) \\ &+ \text{withdrawn EL vault} + \text{withdrawn WL vault}.\end{aligned}$$

These rewards are distributed to stETH shareholders (90%) and to the Lido treasury and Node Operators (10%) in the form of a commission. This distribution occurs at the protocol level every 24 hours when a core protocol state update is triggered by the oracle report. For the 10% commission to be paid out, the rewards report needs to show a positive result, excluding the EL rewards. If the report is positive, Lido mints shares proportional to 10% of the reward, with half going to the Lido treasury and the other half to the Node Operators. However, if the report is negative, no commission is distributed, which aligns incentives for the Node Operators.

For users, the reward is rebased, whether positive or negative, based on the value of their shares in the pool, which is a function of how much ETH is held in the protocol.

B.6 Max and min for rewards payout function of the oracle. What happens in case of slashing if loss > min?

The rebasing functionality on Lido is triggered on each oracle report, initiated when the Accounting Oracle processes a daily protocol state report. Lido has a dedicated Sanity Checker contract to ensure the reported data is within a reasonable range for validator Consensus Layer (CL) balances (TVL). For instance, LIP-23 limits the potential daily drop in Lido's TVL from 5% to 0.19%, reducing the risk of malicious oracle committee activity that could cause up to 20% loss in TVL before Lido's governance can respond.

Lido handles slashing in two modes: turbo mode and bunker mode. In turbo mode, which is the protocol's nominal state, Lido doesn't wait for validators to exit in order to socialize slashings. Instead, in each report, oracles simulate how much Lido will be in the protocol, following a

report induced rebase, at the end of the frame and adjusts a max share rate accordingly.

In the rare case of a mass slashing event, Lido enters bunker mode. In this state, Lido waits for validators to be penalized and exited before fulfilling any queued withdrawals. Under bunker mode, no withdrawals are processed. However, triggering this mode is exceptionally difficult due to strict conditions, and since only around 500 validators have been slashed since the beginning of Ethereums PoS, the probability of this occurring is low.

Non-Academic Relevant Sources

1. Lido Docs: `https://docs.lido.fi/contracts/withdrawal-queue-erc721`

2. Lido deployed contracts: `https://github.com/lidofinance/lido-dao/blob/master/contracts/0.8.9/WithdrawalQueueBase.sol`

3. Lido on Ethereum. Withdrawals landscape: `https://hackmd.io/@lido/SyaJQsZoj`

4. Accounting Oracle: `https://docs.lido.fi/contracts/accounting-oracle`

5. Off-chain oracle source code: `https://github.com/lidofinance/lido-oracle/tree/develop/src`

6. Accounting Oracle info post by Lido: `https://hackmd.io/@lido/r1wF0OLhj`

7. Staking Router: `https://hackmd.io/@lido/rJ0dCOVdo`

8. Depositor bot source code: `https://github.com/lidofinance/depositor-bot/blob/main/src/bots/depositor.py`

9. Info on withdrawals: `https://hackmd.io/@lido/HknYRrCws?type=view`

10. Detailed Bunker Mode: `https://docs.google.com/document/d/1NoJ3rbVZ1OJfByjibHPA91Ghqk487tT0djA`

11. LIP-23 research proposal passed: `https://hackmd.io/@lido/negative-rebase-lip`

12. Associated Slashings Research: `https://hackmd.io/@lido/r1Qkkiv3j`

13. Oracle Sanity checker source code: `https://github.com/lidofinance/lido-dao/blob/master/contracts/0.8.9/sanity_checks/OracleReportSanityChecker.sol`