

Reserve Depletion and Security Runway in Proof-of-Stake Systems

Paolo Penna ✉ 

IOG, Zurich, Switzerland

Manvir Schneider ✉ 

Cardano Foundation, Zurich, Switzerland

Abstract

Many proof-of-stake protocols finance validator rewards from two sources: transaction fees and a finite reserve of native tokens. This creates a dynamic hand-off problem. Early in the life of the system, fees may be too small to fund the target level of security; later, fees may become sufficient. The central question is whether the reserve provides enough runway for the protocol to remain secure until this fee-only region is reached. We study this problem in a discrete-time stochastic model of validator participation. Token price and transaction demand fluctuate over time, while validators choose participation strategically in each state. We solve the validator entry game in closed form and derive an exact state-dependent reserve threshold: for every token price and demand state, the threshold is the minimal reserve stock necessary and sufficient to sustain a target security level. This threshold separates three regions: infeasibility, reserve-dependent security, and fee-only security. The threshold turns reserve adequacy into a hitting-time problem. Security fails when the reserve first falls below the state-dependent threshold, and a successful hand-off occurs exactly when the fee-only region is reached before that failure time. We derive conservative finite-horizon stress-test guarantees that convert lower confidence bands for token price and demand into reserve requirements, and under lognormal price-demand dynamics we obtain explicit failure-probability and expected hand-off-time bounds. Finally, we extend the model to forward-looking validators and derive the Markov participation condition that captures how current participation affects future reserve-funded rewards. The main implication is that reserve policy should not be evaluated by nominal depletion dates or steady-state reward ratios alone. A protocol can have a large nominal reserve and still be close to security failure after adverse price or demand shocks. Conversely, once demand crosses the fee-only threshold, the reserve becomes redundant for security. This paper provides a tractable equilibrium framework for stress-testing this transition.

2012 ACM Subject Classification Theory of computation → Algorithmic game theory and mechanism design; Security and privacy → Economics of security and privacy; Mathematics of computing → Stochastic processes

Keywords and phrases Blockchain, Proof-of-Stake, Reserve Depletion, Security Runway, Validator Incentives, Tokenomics, Economic Security, Transaction Fees, Nash Equilibrium, Markov-Perfect Equilibrium, Stochastic Games

Digital Object Identifier 10.4230/LIPIcs.AFT.2026.6

Related Version *Full Version:* <https://arxiv.org/abs/2606.03587> [22]

Acknowledgements We thank members of the Cardano Foundation and IOG for helpful discussions.

1 Introduction

Proof-of-stake blockchains typically pay validators from two sources: current transaction fees and a reserve of previously minted or retained tokens. The reserve is meant to bridge the early phase of the system, when adoption is still low and fees alone do not support the amount of active stake needed for security. Cardano is a canonical example of such a design: its reward system combines transaction fees with a controlled drawdown of reserves [4].



© Paolo Penna and Manvir Schneider;
licensed under Creative Commons License CC-BY 4.0
8th Conference on Advances in Financial Technologies (AFT 2026).

Editors: Aggelos Kiayias and Maria Kyropoulou; Article No. 6; pp. 6:1–6:22

Leibniz International Proceedings in Informatics



LIPICs Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

Other examples include Avalanche [3], Algorand [1], Ripple [2], and Hedera [11]. The policy problem is therefore not simply whether fee revenue is positive in the long run. The relevant question is whether the protocol can remain secure *along the entire transition path*.

That transition problem has two distinct sources of fragility. First, fee income is uncertain because transaction demand is uncertain. Second, even if the reserve is large in token units, its purchasing power is stochastic because validator costs are incurred in an external numeraire while the reserve is held in the native token. A reserve that looks ample at one token price can become inadequate after a price drawdown.

This creates a governance problem that is not visible from the reserve balance alone. Suppose two protocols have the same number of reserve tokens. The first has strong current demand and a high token price; the second faces weak demand and a price drawdown. Their nominal reserves are identical, but their ability to fund validator participation is not. Conversely, a protocol with a declining reserve may be safe if fee demand has already grown enough to support the target security level without subsidies. The relevant object is therefore a state-contingent reserve requirement.

The security-runway perspective is designed to capture exactly this issue. A protocol has runway at time t if the reserve stock at that date is large enough, given the current state of demand and token price, to keep equilibrium security above the required minimum. Runway ends not when the reserve literally reaches zero, but when it first becomes too small to finance the security shortfall.

This paper studies the following fundamental question:

Can a blockchain protocol guarantee secure operation throughout the transition from reserve-funded to fee-only validator rewards, and under what conditions does this transition fail?

We answer this question by deriving a reserve threshold that can be interpreted as a state-contingent runway criterion: given price, demand, and a target security level, it tells whether the current reserve is sufficient to sustain equilibrium security. This allows reserve policy to be stress-tested state by state.

Our Contribution

We study the reserve-to-fee transition in a stochastic state model with two exogenous state variables: token price and user demand for blockspace. Validators choose participation strategically, so equilibrium security and fees respond to the current reserve, price, and demand state. Our contributions are as follows.

- (i) We solve the validator participation game in closed form and prove existence and uniqueness of a symmetric Nash equilibrium (Theorem 3). The equilibrium shows that the security value of a given reserve stock depends on the current token price.
- (ii) We derive an exact state-dependent reserve threshold (Theorem 10). For each price-demand state, the threshold is the minimal reserve stock necessary and sufficient to sustain a target security level.
- (iii) We use this threshold to formulate the dynamic hand-off problem (Section 4). Security persists while the reserve remains above the threshold, and hand-off succeeds when the fee-only region is reached before failure (Proposition 13, Theorem 15, Corollary 17).
- (iv) We extend the baseline model to a forward-looking Markov environment (Section 5). We prove finite-horizon Markov perfect equilibrium existence and characterize how continuation values alter current validator participation (Theorem 22, Propositions 23, 25, and 27).

- (v) We specialize the runway analysis to geometric Brownian token price and discrete-time lognormal demand (Section 6). This yields explicit finite-horizon failure-probability guarantees and bounds on hand-off timing (Theorem 30, Proposition 31).

Related Literature

Our analysis is related to three strands of literature. First, the Ouroboros line of work studies the protocol and security foundations of proof-of-stake systems and provides the incentive-theoretic background for stake-based consensus [18, 9]. Second, economic analyzes of proof-of-stake study equilibrium incentives, reward design, and the distributional effects of staking rewards [23, 10, 15]. Third, the blockchain fee-market literature studies how congestion pricing and transaction fees support decentralized infrastructure [13, 25]. Dynamics models for cryptocurrencies include [17, 14, 6, 16], while geometric Brownian motion is used, e.g., in [7, 8]. To the best of our knowledge, existing work has not isolated the reserve hand-off problem studied here: a finite native-token reserve, strategic validator participation, stochastic token price and demand, and a state-dependent threshold for sustaining a target security level until fee-only operation becomes possible. Earlier work models the effects of fee-only rewards in Bitcoin: First, higher variance induces “forking” strategies of miners competing for transactions with high tips [5]. [24] further show that mining gaps can emerge even before block rewards disappear, with adverse effects on security and decentralization. Moreover, attackers can bribe honest participants by offering tips for mining a forked chain and increase the chances of successfully performing double spending [19].

Contribution Relative to the Literature

The present paper differs from these strands by focusing on a dynamic reserve-funded transition problem. The object of interest is not simply equilibrium rewards at a point in time, nor fee formation in isolation, but the reserve threshold required to sustain security until fee funding becomes sufficient on its own. That is the role of the security-runway concept developed below.

Our model has some similarities with the validator side in [17]: there, validators decide the amount of blockspace (or security) to provide, and their utility depends on the per-unit fee of the system as equilibrium. That paper, however, does not consider reserve constraints, and it models users demand differently.

In our model, both token prices and demand are exogenous quantities. In [14, 16] instead, the system value is determined by the protocol evolution (at some cost function) and the token price is an endogenous quantity which is part of the equilibrium (the price at which tokens are exchanged in a spot market involving validators and users). In this respect, our model exhibits greater flexibility and generality, while additionally incorporating the aforementioned reserve constraint, which is absent in [14, 16].

Deferred Proofs

All proofs are available in the appendix of the full version of the paper [22].

2 Model

Time is discrete and indexed by $t = 0, 1, 2, \dots$. All random variables are defined on a filtered probability space $(\Omega, \mathcal{F}, (\mathcal{F}_t)_{t \geq 0}, \mathbb{P})$. The protocol state at time t is given by a triple $X_t = (R_t, Price_t, FeeMax_t)$, where:

- $R_t \in \mathbb{R}_+$ is the reserve stock measured in native-token units.
- $Price_t \in \mathbb{R}_{++}$ is the token price in an external numeraire. Validator costs are paid in this numeraire.
- $FeeMax_t \in \mathbb{R}_{++}$ is a demand state governing willingness to pay for block space. Intuitively, this is the largest fee that at least some user is willing to pay (any higher fee results in zero demand).

The policy parameters are fixed ex ante:

$$\theta \in [0, 1), \quad \rho \in (0, 1),$$

where θ is the share of fee revenue diverted into the reserve and ρ is the fraction of the reserve paid out each period. Token price and demand are stochastic, exogenous, and positive quantities:

► **Assumption 1** (State dynamics). *The token price process $(Price_t)_{t \geq 0}$ and the demand process $(FeeMax_t)_{t \geq 0}$ are strictly positive and adapted to the filtration $(\mathcal{F}_t)_{t \geq 0}$.*

2.1 Game Form and Economic Scope

The players are the N validators. Their date- t actions are participation levels, interpreted as active stake or security supply. Users are not modeled as strategic players; instead, user behavior is summarized by an inverse-demand schedule for block space. Likewise, token-market trading is not modeled strategically; it enters through the stochastic price process $Price_t$. This separation isolates the security-runway question: how the reserve interacts with validator incentives when current fee conditions and the external token price fluctuate over time.

This modeling choice has two implications that should be kept distinct throughout the paper. First, the token price $Price_t$ is externally given to validators at date t , so the model does not claim to explain token valuation. Second, the transaction fee f_t is endogenous, because it is the market-clearing fee induced by validator participation and the current demand state $FeeMax_t$. Thus demand and token price are state variables, while validator participation and fees are equilibrium outcomes.

2.2 A Reduced-Form Security Technology

There are $N \geq 2$ validators. Validator i chooses an amount of active stake, or security supply,

$$a_{i,t} \in \mathbb{R}_+,$$

and aggregate security and blockspace is

$$s_t := \sum_{i=1}^N a_{i,t}. \tag{1}$$

The baseline model is symmetric across validators: every validator has the same feasible action set, the same cost function, and the same proportional reward-sharing rule.¹

The action $a_{i,t}$ should be interpreted as validator i 's effective supply of secure blockspace at date t . It combines the validator's active participation with the amount of transaction-processing capacity that the validator makes available to users. Thus, a validator who is

¹ The model therefore suppresses ex ante heterogeneity in stake, scale, or operating cost.

online but produces empty or only partially filled blocks is represented as supplying a lower effective $a_{i,t}$. The aggregate s_t is therefore the quantity of secure blockspace supplied in equilibrium. It enters the security side of the model because larger effective participation raises the security level, and it enters the demand side because it is the quantity of blockspace cleared at the market fee. This convention keeps the model focused on the funding problem: how much effective secure capacity validators are willing to supply when rewards come from fees and reserve payouts.

Users are described by an *inverse demand curve* expressed in terms of the external numeraire (for example, in dollars):

$$P(s_t, FeeMax_t) := FeeMax_t - Sec \cdot s_t . \quad (2)$$

Intuitively, given the available blockspace provided by the validators (1), Equation 2 gives the *market-clearing* fee:

- Users demand equals the supplied blockspace s if the protocol charges $P(s, FeeMax)$ dollars per unit of blockspace. Note that we are assuming each user to consume one unit of blockspace – say one transaction. Hence, the demand corresponding to (2) is of the form $D(p) = (FeeMax - p)/Sec$ for $p \in [0, FeeMax]$, and $D(p) = 0$ for $p > FeeMax$.
- The parameter $Sec > 0$ determines the rate at which market-clearing fees decline as total blockspace supply increases (intuitively, the protocol must lower fees to accommodate higher user demand).

As already mentioned above, while $FeeMax_t$ is an exogenous parameter controlling users demand, blockspace s_t is strategic and results from validators aiming at maximizing their own utilities. If the supplied blockspace is s_t , the protocol posts a token-denominated fee

$$f_t(s_t, FeeMax_t, Price_t) := \frac{P(s_t, FeeMax_t)}{Price_t} = \frac{FeeMax_t - Sec \cdot s_t}{Price_t} . \quad (3)$$

The resulting fee revenue, in external numeraire, is thus

$$Price_t \cdot f_t \cdot s_t = (FeeMax_t - Sec \cdot s_t) \cdot s_t .$$

The model restricts attention to action profiles that yield nonnegative clearing fees, reflecting the natural requirement that the system does not subsidize users for consuming blockspace. Since $P(s_t, FeeMax_t) \geq 0$ requires $s_t \leq FeeMax_t/Sec$, we impose the symmetric feasibility bound

$$a_{i,t} \in \left[0, \frac{FeeMax_t}{N \cdot Sec} \right] . \quad (4)$$

This condition implies that $0 \leq s_t \leq FeeMax_t/Sec$, ensuring that clearing fees are always nonnegative. Furthermore, in the context of symmetric equilibria – where all validators select an identical strategy – these two conditions are equivalent.

2.3 Reserve Dynamics and Validator Rewards

The reserve absorbs a fraction θ of current fee revenue and a fraction ρ of the reserve is paid out each period. Because the reserve is held in tokens, its law of motion is

$$R_{t+1} = (1 - \rho) \cdot R_t + \theta \cdot \frac{(FeeMax_t - Sec \cdot s_t) \cdot s_t}{Price_t} . \quad (5)$$

The total reward pool paid to validators in period t , measured in the external numeraire, is

$$W_t := (1 - \theta) \cdot (FeeMax_t - Sec \cdot s_t) \cdot s_t + \rho \cdot Price_t \cdot R_t. \quad (6)$$

The first term is the portion of fees not diverted to the reserve; the second term is the external value of the reserve payout.

Each validator has quadratic operating cost

$$\frac{\kappa}{2} a_{i,t}^2, \quad \kappa > 0. \quad (7)$$

The quadratic cost reflects increasing marginal costs: providing additional secure blockspace becomes more costly as a validator's supply grows. Kiayias et al. [17] use a general convex cost for blockspace provision, while Houy [12] uses a quadratic cost. The quadratic form also keeps the first-order condition tractable whenever the capacity bound does not bind.

2.4 Stage-Game Equilibrium

In this section, we consider a single-stage game which is fully specified by the current state:² the current reserve R_t , token price $Price_t$, and demand $FeeMax_t$. As we focus only on a generic stage t , we suppress t and consider a generic state $x = (R, Price, FeeMax)$. The stage game at state x is played only by the N validators. A pure action for validator i is a feasible action a_i satisfying (4), that is, $a_i \in [0, FeeMax/(N \cdot Sec)]$. Actions are chosen simultaneously. If the aggregate action profile is $\mathbf{a} = (a_1, \dots, a_N)$ and $s = \sum_i a_i > 0$, validator i receives its pro rata share of the reward pool (6) given the current state and action profile,

$$W = W(\mathbf{a}; x) := \frac{a_i}{s} \cdot ((1 - \theta) \cdot (FeeMax - Sec \cdot s) \cdot s + \rho \cdot Price \cdot R).$$

The resulting utility (reward minus incurred cost) is thus

$$u_i(\mathbf{a}; x) := \frac{a_i}{s} \cdot W - \frac{\kappa}{2} a_i^2 = \frac{a_i}{s} ((1 - \theta) \cdot (FeeMax - Sec \cdot s) \cdot s + \rho \cdot Price \cdot R) - \frac{\kappa}{2} a_i^2. \quad (8)$$

When all validators choose zero, we set $u_i(0, \dots, 0; x) = 0$.

► **Definition 2** (Symmetric Nash equilibrium). *Fix a state $x = (R, Price, FeeMax)$. A symmetric Nash equilibrium is a profile $\mathbf{a}^*(x) = (a^*(x), \dots, a^*(x))$ such that*

$$u_i(a^*(x), \dots, a^*(x); x) \geq u_i(a_i, \mathbf{a}_{-i}^*(x); x)$$

for every validator i and every feasible deviation $a_i \in [0, FeeMax/(N \cdot Sec)]$, where $(a_i, \mathbf{a}_{-i}^(x))$ denotes the vector obtained by replacing the i^{th} entry in profile $\mathbf{a}^*(x)$ with a_i .*

The induced equilibrium security, fee, and next-period reserve are given by:

$$\begin{aligned} s^*(x) &= N \cdot a^*(x), & f^*(x) &\stackrel{(3)}{=} \frac{FeeMax - Sec \cdot s^*(x)}{Price}, \\ R^+(x) &= (1 - \rho) \cdot R + \theta \cdot f^*(x) \cdot s^*(x) = (1 - \rho) \cdot R + \theta \cdot \frac{(FeeMax - Sec \cdot s^*(x)) \cdot s^*(x)}{Price}. \end{aligned}$$

² Later, we extend this game and its analysis to a dynamic “multi-stage” version.

Symmetric equilibria arise naturally in the context of blockchain systems (see, e.g., [14, 16]), as they constitute the intended design objective of many protocols: they correspond to a regime of maximal decentralization [21, 20].

The definition above is for a generic state and yields a state by state evolution in the natural way: Given the current state $x_t = (R_t, Price_t, FeeMax_t)$, the reserve at the next state is $R_{t+1} = R^+(x_t)$, and therefore the next state is $x_{t+1} = (R_{t+1}, Price_{t+1}, FeeMax_{t+1})$. Dynamic incentives are encoded only through the reserve stock carried to the next date; validators do not solve an intertemporal control problem. The analysis focuses on whether the current state provides enough funding to sustain current security, not on dynamic reputation effects.

Observe that the state is composed of three variables, of which two are exogenous: The reserve stock R is an endogenous state variable inherited from past protocol outcomes. The token price $Price$ and the demand shifter $FeeMax$ are exogenous state variables from the perspective of the validator game. Conditional on a state x , equilibrium participation $a^*(x)$ and the clearing fee $f^*(x)$ are endogenous objects solved within the model.

3 Static Equilibrium Analysis

This section solves the stage game exactly. The main result is a closed-form symmetric Nash equilibrium, which then becomes the building block for the dynamic runway analysis.

3.1 Existence and Uniqueness

Define the constants

$$A_N := \frac{\kappa}{N} + \frac{N+1}{N} \cdot (1-\theta) \cdot Sec, \quad B_N := \frac{N-1}{N} \cdot \rho. \quad (9)$$

The next theorem shows the uniqueness of symmetric equilibria in the stage game. This result carries over to the multi-round dynamic game, as we prove below (Proposition 23). Note that the number of validators N does affect the equilibrium strategies through the two quantities defined above.

► **Theorem 3** (Unique symmetric Nash equilibrium). *For every state $x = (R, Price, FeeMax) \in \mathbb{R}_+ \times \mathbb{R}_{++}^2$, the stage game admits a unique symmetric Nash equilibrium. Its aggregate security level is*

$$s^*(x) = \min \left\{ \frac{(1-\theta) \cdot FeeMax + \sqrt{(1-\theta)^2 \cdot FeeMax^2 + 4A_N B_N \cdot Price \cdot R}}{2A_N}, \frac{FeeMax}{Sec} \right\}. \quad (10)$$

The corresponding equilibrium fee is $f^(x) = \frac{FeeMax - Sec \cdot s^*(x)}{Price}$.*

► **Remark 4.** The equilibrium formula shows exactly how the reserve enters incentives. Current fee funding depends on demand $FeeMax$ and congestion Sec , while reserve funding enters only through the product $Price \cdot R$, the reserve's value in the external numeraire. The same reserve stock therefore has different security consequences at different token prices. Hence, the reserve balance alone does not determine how much security it can support; its value must be considered together with the token price.

Theorem 3 implies that, in any symmetric equilibrium, the validators' utilities are given by the following expression.

► **Corollary 5.** *For every state $x = (R, Price, FeeMax) \in \mathbb{R}_+ \times \mathbb{R}_{++}^2$, the equilibrium utilities of the validators in the corresponding stage game are*

$$u_i^*(x) := \frac{1}{N} \cdot ((1 - \theta) \cdot (FeeMax - Sec \cdot s^*(x)) \cdot s^*(x) + \rho \cdot Price \cdot R) - \frac{\kappa}{2} \cdot \left(\frac{s^*(x)}{N} \right)^2 \quad (11)$$

$$= s^*(x) \cdot \frac{(1 - \theta) \cdot FeeMax + \rho \cdot Price \cdot R}{N} - (s^*(x))^2 \cdot \left(\frac{(1 - \theta) \cdot Sec}{N} + \frac{\kappa}{2N^2} \right). \quad (12)$$

In particular, in any non-interior equilibrium, i.e. $s^*(x) = \frac{FeeMax}{Sec}$, the utilities are equal to

$$u_i^*(x) = \frac{1}{N} \cdot (\rho \cdot Price \cdot R) - \frac{\kappa}{2} \cdot \left(\frac{FeeMax}{N \cdot Sec} \right)^2. \quad (13)$$

The second part of the corollary above implies that, for sufficiently large R , the utility of each validator is entirely reserve-driven. This is because the resulting equilibrium must be non-interior, and the corresponding market-clearing fee (3) is zero. The next corollary deals with the opposite case of small reserve (intuitively, when R is not large enough to make the utility in (13) nonnegative).

► **Corollary 6.** *For sufficiently small R , the unique symmetric equilibrium must be interior, that is,*

$$s^*(R, Price, FeeMax) < \frac{FeeMax}{Sec}.$$

In particular, this holds true for any $R < R_{\min} := \left(\frac{FeeMax}{N \cdot Sec} \right)^2 \cdot \frac{\kappa N}{2\rho \cdot Price} = \frac{\kappa FeeMax^2}{2\rho \cdot Price \cdot N \cdot Sec^2}$.

Proof. By contradiction, if $s^*(R, Price, FeeMax) \geq \frac{FeeMax}{Sec}$, then $u_i^*(x) = \frac{1}{N} (0 + \rho \cdot Price \cdot R) - \frac{\kappa}{2} \cdot \left(\frac{FeeMax}{N \cdot Sec} \right)^2 < 0$. But validator i can deviate to $a_i = 0$, yielding utility 0. Hence the boundary profile cannot be a Nash equilibrium. Therefore the equilibrium must be interior. ◀

Note that for any interior equilibria and for $Q := N \cdot (1 - \theta) \cdot FeeMax$, we have

$$s^*(R, Price, FeeMax) = \frac{Q + \sqrt{Q^2 + 4 \cdot (\kappa + (N + 1) \cdot (1 - \theta) \cdot Sec) \cdot (N - 1) \cdot \rho \cdot Price \cdot R}}{2(\kappa + (N + 1) \cdot (1 - \theta) \cdot Sec)}. \quad (14)$$

► **Example 7 (No reserve).** Consider $R = 0$, so validator rewards are funded entirely by fees. By Corollary 6, the equilibrium is interior. Hence, for $x = (0, Price, FeeMax)$ we have

$$s^*(x) \stackrel{(14)}{=} \frac{N \cdot (1 - \theta) \cdot FeeMax}{\kappa + (N + 1) \cdot (1 - \theta) \cdot Sec} \stackrel{(9)}{=} \frac{(1 - \theta) \cdot FeeMax}{A_N}. \quad (15)$$

Next, we calculate the utilities of the validators. By Corollary 5, with $R = 0$

$$\begin{aligned}
 u_i^*(x) &= \frac{1}{N} \cdot (1 - \theta) \cdot (FeeMax - Sec \cdot s^*(x)) \cdot s^*(x) - \frac{\kappa}{2} \cdot \left(\frac{s^*(x)}{N} \right)^2 \\
 &\stackrel{(15)}{=} \frac{1}{N} \cdot s^*(x) \cdot (A_N - (1 - \theta) \cdot Sec) \cdot s^*(x) - \frac{\kappa}{2} \cdot \left(\frac{s^*(x)}{N} \right)^2 \\
 &= \left(\frac{s^*(x)}{N} \right)^2 \cdot (N \cdot A_N - N \cdot (1 - \theta) \cdot Sec) - \frac{\kappa}{2} \cdot \left(\frac{s^*(x)}{N} \right)^2 \\
 &\stackrel{(9)}{=} \left(\frac{s^*(x)}{N} \right)^2 \cdot (\kappa + (N + 1) \cdot (1 - \theta) \cdot Sec - N \cdot (1 - \theta) \cdot Sec) - \frac{\kappa}{2} \cdot \left(\frac{s^*(x)}{N} \right)^2 \\
 &= \left(\frac{s^*(x)}{N} \right)^2 \cdot \left(\frac{\kappa}{2} + (1 - \theta) \cdot Sec \right) > 0.
 \end{aligned}$$

Thus, when $R = 0$, the equilibrium is interior and each validator obtains strictly positive utility.

The next result shows that equilibrium security increases when validator rewards become easier to finance. Higher demand raises current fee revenue, and a higher token price makes reserve payouts more valuable in external terms. By contrast, higher operating costs, stronger fee compression, or diverting a larger fraction of the fees into the reserve reduce the validators' current incentives to supply security and blockspace.

► **Proposition 8** (Monotonicity). *The equilibrium security level $s^*(R, Price, FeeMax)$ from Theorem 3 is weakly increasing in R , in $Price$, and in $FeeMax$, and weakly decreasing in κ , in Sec , and in θ .*

3.2 Security Threshold

In this section, we analyze the conditions under which the system sustains an equilibrium level of security (blockspace) above a threshold. Let $\underline{s} > 0$ denote the target security requirement.

► **Definition 9** (Security-feasible state). *A state $(R, Price, FeeMax)$ is security feasible for target $\underline{s}$ if*

$$s^*(R, Price, FeeMax) \geq \underline{s}.$$

It is fee-only feasible if the same inequality holds at $R = 0$.

The next theorem is the static core of the paper. It converts the equilibrium formula into a reserve threshold.

► **Theorem 10** (Exact reserve threshold). *For every target $\underline{s} > 0$, define*

$$\text{SecureR}(Price, FeeMax; \underline{s}) := \begin{cases} \frac{[A_N \cdot \underline{s}^2 - (1 - \theta) \cdot FeeMax \cdot \underline{s}]_+}{B_N \cdot Price}, & \text{if } \underline{s} \leq FeeMax/Sec, \\ \infty, & \text{if } \underline{s} > FeeMax/Sec. \end{cases} \quad (16)$$

Then the following are equivalent:

- (i) $(R, Price, FeeMax)$ is security feasible for target $\underline{s}$.
- (ii) $R \geq \text{SecureR}(Price, FeeMax; \underline{s})$.

Moreover, $(R, Price, FeeMax)$ is fee-only feasible if and only if $\text{SecureR}(Price, FeeMax; \underline{s}) = 0$.




Figure 1 Exact reserve threshold as a function of token price $Price$ and demand state $FeeMax$. For any point in the $(Price, FeeMax)$ plane, the color gives the minimal reserve stock required to sustain the target security level $\underline{s}$: darker regions require larger reserves. The gray lower region lies below the capacity cutoff $FeeMax = Sec \cdot \underline{s}$; there the target is infeasible at any finite reserve level because aggregate supply cannot be increased enough while keeping fees nonnegative. Above the fee-only cutoff $FeeMax = A_N \cdot \underline{s} / (1 - \theta)$, the reserve threshold is zero because current fees alone can sustain the target. The figure uses the illustrative parameter values $N = 5$, $\theta = 0.25$, $\rho = 0.04$, $\kappa = 0.6$, $Sec = 0.8$, and $\underline{s} = 1$.

Theorem 10 isolates the object that governance actually needs to monitor: not the reserve level alone, but the reserve level relative to a state-dependent threshold $SecureR(\cdot)$. When demand is strong or the token price is high, the threshold falls because either current fees or reserve purchasing power are more favorable. When demand is too weak to support the target capacity, no reserve level can repair the shortfall. The threshold becomes infinite when the target security level $\underline{s}$ exceeds current demand capacity $FeeMax / Sec$ (recall that by (4) this ratio is the largest feasible s , i.e., the protocol cannot buy more secure throughput at any nonnegative fee).

► **Remark 11.** The threshold $SecureR(\cdot)$ is easiest to interpret geometrically. Figure 1 plots the exact threshold over the $(Price, FeeMax)$ plane for an illustrative parameterization. The figure makes visible the three economically distinct regions identified by Theorem 10: (i) an infeasible low-demand region, (ii) a strictly reserve-dependent transition region, and (iii) a fee-only region where the threshold collapses to zero.

4 Dynamic Security Runway

The static threshold becomes dynamic once price and demand evolve over time and the reserve follows Equation (5). This section formalizes the hand-off problem.

4.1 Failure Time and Hand-Off Time

In this section, we study the hand-off time for myopic players, whose equilibrium is characterized at every step t by Theorem 3. In particular, the reserve evolution in (5) becomes

$$R_{t+1} = (1-\rho)R_t + \theta \frac{(FeeMax_t - Sec \cdot s^*(R_t, Price_t, FeeMax_t))s^*(R_t, Price_t, FeeMax_t)}{Price_t}. \quad (17)$$

► **Definition 12** (Runway stopping times). *Fix a target security level $\underline{s} > 0$. Define*

$$\tau_{\text{fail}} := \inf\{t \geq 0 : R_t < \text{SecureR}(Price_t, FeeMax_t; \underline{s})\}, \quad (18)$$

$$\tau_{\text{hand}} := \inf\{t \geq 0 : \text{SecureR}(Price_t, FeeMax_t; \underline{s}) = 0\}. \quad (19)$$

We say that the protocol achieves a successful hand-off if $\tau_{\text{hand}} < \tau_{\text{fail}}$.

Next, we turn the hand-off problem into a hitting-time problem. Security lasts exactly as long as the reserve process remains inside the safe region, and the transition succeeds exactly when the protocol reaches the fee-only region before leaving that safe region. While a successful handoff is a necessary condition, it does not provide any guarantee for the system to be self-sustainable (i.e. stay in the fee-only region) after the handoff. As we prove below, we shall provide sufficient conditions under which the system does not fail before any target period T (Theorem 15).

► **Proposition 13** (Exact runway criterion). *For every horizon $T \in \mathbb{N}$ and every sample path:*

(i) *Security is maintained at all dates $t = 0, \dots, T$ if and only if*

$$R_t \geq \text{SecureR}(Price_t, FeeMax_t; \underline{s}) \quad \text{for all } t = 0, \dots, T.$$

(ii) *A successful hand-off occurs if and only if the process reaches the fee-only region before the failure region, that is,*

$$\tau_{\text{hand}} < \tau_{\text{fail}}.$$

Proof. Part (i) follows directly from Theorem 10 applied state by state. Part (ii) is the definition of successful hand-off expressed in terms of the stopping times in Definition 12. ◀

4.2 A Conservative Stress-Test Bound

The exact criterion is pathwise. For protocol design one often wants a finite-horizon sufficient condition that can be checked before the system is launched. The next result gives such a bound.

► **Lemma 14** (Pure-decay lower bound). *Along every sample path and for every $t \geq 0$,*

$$R_t \geq (1 - \rho)^t \cdot R_0.$$

Proof. The reserve recursion in Equation (17) has a nonnegative inflow term. Hence $R_{t+1} \geq (1 - \rho) \cdot R_t$. Iterating yields the claim. ◀

Lemma 14 is the worst-case reserve benchmark: even if future fee inflows are ignored entirely, the reserve cannot decay faster than pure payout at rate ρ . This simple lower bound is what makes conservative stress testing possible.

6:12 Reserve Depletion and Security Runway in Proof-of-Stake Systems

The next theorem provides a robust sufficient condition: if the reserve is large enough to survive a pessimistic lower envelope for price and demand, then the protocol is safe on every path inside that envelope. This form is useful for ex ante certification and governance stress tests.

► **Theorem 15** (Finite-horizon stress-test guarantee). *Fix a horizon $T \in \mathbb{N}$ and deterministic lower envelopes $(\underline{Price}_t)_{t=0}^T$ and $(\underline{FeeMax}_t)_{t=0}^T$. Define the event*

$$E_T := \{Price_t \geq \underline{Price}_t, FeeMax_t \geq \underline{FeeMax}_t \quad \text{for all } t = 0, \dots, T\}.$$

If

$$(1 - \rho)^t R_0 \geq \text{SecureR}(\underline{Price}_t, \underline{FeeMax}_t; \underline{s}) \quad \text{for all } t = 0, \dots, T, \quad (20)$$

then event E_T implies that failure occurs after T , that is, $\tau_{\text{fail}} > T$. Consequently, for any $\alpha \in [0, 1]$, if $\mathbb{P}(E_T) \geq 1 - \alpha$, then $\mathbb{P}(\tau_{\text{fail}} > T) \geq 1 - \alpha$.

► **Remark 16.** Theorem 15 is intended for protocol governance. One may estimate lower confidence bands for $Price_t$ and $FeeMax_t$, plug them into Equation (20), and obtain a conservative reserve requirement for a target horizon T . The bound is conservative because it ignores future reserve inflows from fees; it treats the reserve as if it only decayed. That makes it suitable for adverse-scenario certification.

4.3 Fee-Only Region and Long-Run Design

We identify the economic moment at which the reserve becomes redundant, that is, current demand is enough to finance the target security level. In this case, reserve policy no longer determines feasibility.

► **Corollary 17** (Fee-only region). *Suppose $\underline{s} \leq FeeMax/Sec$. Then fee-only feasibility at target $\underline{s}$ holds if and only if*

$$(1 - \theta) \cdot FeeMax \geq A_N \cdot \underline{s}. \quad (21)$$

Proof. By Theorem 10, fee-only feasibility is equivalent to $\text{SecureR}(Price, FeeMax; \underline{s}) = 0$. Under $\underline{s} \leq FeeMax/Sec$, this is equivalent to

$$A_N \cdot \underline{s}^2 - (1 - \theta)FeeMax \cdot \underline{s} \leq 0.$$

Because $\underline{s} > 0$, division by $\underline{s}$ yields Equation (21). ◀

► **Remark 18.** Equation (21) says that fee-only sustainability is governed by demand, *not* by token price. This is because current transaction fees paid by users are already measured in the external numeraire. Token-price risk matters only because the reserve is a stock of tokens (carried from the past).

5 Dynamic Validator Participation: A Markov Extension

The baseline model in considered so far is myopic: At state X_t , validators choose current participation to maximize current payoff only. Intuitively, this is what makes the reserve threshold explicit. In this section, we consider a fully strategic dynamic extension that endogenizes continuation values while preserving a precise equilibrium concept.

There are two technical issues. First, once validators become forward looking, a deviation at time t changes the future reserve stock and therefore future incentives. Second, proving equilibrium existence in the original continuous-state model requires a substantial stochastic-game apparatus that would obscure the runway results. We therefore proceed in two steps. We first formulate a finite-horizon finite-state Markov game and prove existence of a Markov perfect equilibrium. We then return to the continuous model and derive the exact first-order condition that any differentiable pure Markov equilibrium must satisfy.

5.1 Finite-Horizon Markov Game

Fix a horizon $T \in \mathbb{N}$ and a discount factor $\delta \in (0, 1)$.

► **Assumption 19** (Finite-state Markov extension). *There is a finite set of exogenous states*

$$\mathcal{Z} = \{z^1, \dots, z^M\} \subset \mathbb{R}_{++}^2, \quad z^m = (Price^m, FeeMax^m),$$

and a Markov transition matrix $P = (f_{mn})_{m,n=1}^M$ on $\mathcal{Z}$. There is also a finite reserve grid

$$\mathcal{R} = \{R^1, \dots, R^L\} \subset \mathbb{R}_+, \quad 0 = R^1 < \dots < R^L = \bar{R},$$

with $Price^{\min} := \min_{1 \leq m \leq M} Price^m$, $FeeMax^{\max} := \max_{1 \leq m \leq M} FeeMax^m$ and

$$\bar{R} \geq \max \left\{ R_0, \frac{\theta \cdot (FeeMax^{\max})^2}{4Sec \cdot \rho \cdot Price^{\min}} \right\}.$$

At state $x = (R, z) = (R, Price, FeeMax) \in \mathcal{R} \times \mathcal{Z}$, each validator chooses an action in the compact interval

$$\mathcal{A}(x) := \left[0, \frac{FeeMax}{NSec} \right].$$

The next-period reserve is computed by first determining the intermediate value

$$\hat{R}'(x, a) := (1 - \rho) \cdot R + \theta \cdot \frac{(FeeMax - Sec \cdot s) \cdot s}{Price}, \quad s := \sum_{i=1}^N a_i,$$

and subsequently projecting this value onto the discrete reserve grid via

$$\Pi_{\mathcal{R}}(r) := \max\{R^\ell \in \mathcal{R} : R^\ell \leq r\}.$$

Accordingly, the next-period reserve is given by

$$R'(x, a) := \Pi_{\mathcal{R}}(\hat{R}'(x, a)).$$

► **Remark 20.** Assumption 19 is a computational extension of the baseline model. The exogenous Markov chain is a finite approximation of the original price-demand process, and the reserve grid is the standard discretization used in dynamic policy computation. The bound on $\bar{R}$ guarantees that the reserve state is invariant: because $(FeeMax - Sec \cdot s) \cdot s \leq FeeMax^2/(4Sec)$ and $Price \geq Price^{\min}$,

$$\hat{R}'(x, a) \leq (1 - \rho) \cdot \bar{R} + \frac{\theta \cdot (FeeMax^{\max})^2}{4Sec \cdot Price^{\min}} \leq \bar{R}.$$

► **Definition 21** (Markov strategy and Markov perfect equilibrium). *For each date $t \in \{0, \dots, T\}$ and each validator i , a (mixed) Markov strategy is a mapping*

$$\sigma_t^i : \mathcal{R} \times \mathcal{Z} \rightarrow \Delta(\mathcal{A}(x)),$$

where $\Delta(\mathcal{A}(x))$ denotes the set of Borel probability measures on $\mathcal{A}(x)$.

Given a strategy profile $\sigma = (\sigma_t^i)_{i,t}$, define continuation values recursively by

$$V_{i,T+1}^\sigma(x) := 0,$$

and, for $t = T, T-1, \dots, 0$,

$$V_{i,t}^\sigma(x) = \mathbb{E}_{a_1 \sim \sigma_t^1(\cdot|x), \dots, a_N \sim \sigma_t^N(\cdot|x)} \left[u_i(a; x) + \delta \sum_{z' \in \mathcal{Z}} P(z, z') V_{i,t+1}^\sigma(R'(x, a), z') \right]. \quad (22)$$

A strategy profile σ is a Markov perfect equilibrium (in mixed strategies) if for every date t , every state x , every validator i , and every alternative mixed action $\mu \in \Delta(\mathcal{A}(x))$,

$$\begin{aligned} & \mathbb{E}_{a_1 \sim \sigma_t^1(\cdot|x), \dots, a_N \sim \sigma_t^N(\cdot|x)} \left[u_i(a; x) + \delta \sum_{z' \in \mathcal{Z}} P(z, z') V_{i,t+1}^\sigma(R'(x, a), z') \right] \\ & \geq \mathbb{E}_{\substack{a_i \sim \mu \\ a_j \sim \sigma_t^j(\cdot|x), j \neq i}} \left[u_i(a; x) + \delta \sum_{z' \in \mathcal{Z}} P(z, z') V_{i,t+1}^\sigma(R'(x, a), z') \right]. \end{aligned} \quad (23)$$

The next theorem shows that once the state space is discretized and the horizon is finite, forward-looking validator behavior can still be analyzed with a well-defined equilibrium concept.

► **Theorem 22** (Existence of finite-horizon Markov perfect equilibrium). *Under Assumption 19, the finite-horizon dynamic validator game admits a Markov perfect equilibrium in mixed strategies.*

The proof of the above theorem is a backward-induction argument: each date- t problem becomes an ordinary continuation game once later continuation values are fixed.

The next result says that dynamic multiplicity does not arise if every continuation game is already pinned down locally. Once each state-date problem has a unique symmetric pure best-response fixed point, backward induction propagates that uniqueness through the full dynamic game.

► **Proposition 23** (Purity and uniqueness under state-by-state uniqueness). *Define the date- t continuation-game payoff by*

$$g_{i,t}(x, a) := u_i(a; x) + \delta \sum_{z' \in \mathcal{Z}} P(z, z') V_{i,t+1}^\sigma(R'(x, a), z'). \quad (24)$$

Suppose that, for every date $t \in \{0, \dots, T\}$ and every state $x \in \mathcal{R} \times \mathcal{Z}$, the continuation game with payoff Equation (24) admits a unique symmetric pure Nash equilibrium action $a_t^*(x) \in \mathcal{A}(x)$. Then the dynamic validator game admits a unique symmetric pure Markov perfect equilibrium, namely the profile that assigns action $a_t^*(x)$ to every validator at every state-date pair.

► **Example 24** (Illustrating Proposition 23). We construct a simple finite-horizon example in which the hypothesis of Proposition 23 is satisfied. Consider the finite-horizon Markov game with

$$N = 2, \quad T < \infty, \quad \delta \in (0, 1),$$

and let the exogenous state space Z and transition matrix P be arbitrary but finite, as in Assumption 19. Let the reserve grid $\mathcal{R}$ also be arbitrary and finite. Now choose $\theta = 0$. Under this choice, the reserve transition becomes

$$R'(x, a) = (1 - \rho)R,$$

so the next-period reserve no longer depends on the current action profile a . Hence, for every date t and state $x = (R, z)$, the continuation-game payoff

$$g_{i,t}(x, a) = u_i(a; x) + \delta \sum_{z' \in Z} P(z, z') V_{i,t+1}^\sigma(R'(x, a), z')$$

can be written as

$$g_{i,t}(x, a) = u_i(a; x) + C_t(x),$$

where $C_t(x)$ is a constant with respect to the action profile a . Therefore, the continuation game at any state-date pair (t, x) has exactly the same best responses as the baseline stage game. By Theorem 3, for every state $x = (R, Price, FeeMax)$, the baseline stage game admits a unique symmetric Nash equilibrium, with aggregate security

$$s^*(x) = \min \left\{ \frac{(1 - \theta)FeeMax + \sqrt{(1 - \theta)^2 FeeMax^2 + 4A_N \cdot B_N \cdot Price \cdot R}}{2A_N}, \frac{FeeMax}{Sec} \right\},$$

and corresponding individual action $a^*(x) = \frac{s^*(x)}{N}$. Since here $\theta = 0$, each continuation game therefore admits the same unique symmetric pure Nash equilibrium action $a_t^*(x) = a^*(x)$ for every $t \in \{0, \dots, T\}$, $x \in \mathcal{R} \times Z$. Thus, the hypothesis of Proposition 23 is satisfied, and it follows that the dynamic validator game admits a unique symmetric pure Markov perfect equilibrium.

5.2 Continuous-State First-Order Condition

The finite-state extension above gives an existence theorem. To understand economically how forward-looking incentives modify the baseline equilibrium, it is useful to return to the original continuous model and derive the equilibrium condition that must hold whenever a pure differentiable Markov equilibrium exists.

The next result shows exactly how forward-looking incentives perturb the myopic equilibrium.

► **Proposition 25** (Dynamic symmetric first-order condition). *Fix a horizon $T \geq 1$ and suppose that, in the original continuous model, there exists a pure symmetric Markov perfect equilibrium with continuation value functions*

$$V_{t+1}(R, Price, FeeMax), \quad t = 0, \dots, T - 1,$$

that are continuously differentiable in R . Fix a date $t < T$ and a state $x = (R, Price, FeeMax)$. If the equilibrium at that state is interior, with aggregate security $s_t^M(x) \in (0, FeeMax/Sec)$, then it satisfies

$$(1 - \theta)FeeMax - A_N \cdot s + \frac{B_N \cdot Price \cdot R}{s} + \frac{\delta \cdot \theta}{Price} \left(FeeMax - \frac{N + 1}{N} Sec \cdot s \right) M_{t+1}(x; s) = 0, \quad (25)$$

where $s = s_t^M(x)$ and

$$M_{t+1}(x; s) := \mathbb{E} \left[\partial_R V_{t+1} \left((1 - \rho)R + \theta \frac{(FeeMax - Sec \cdot s)s}{Price}, Price_{t+1}, FeeMax_{t+1} \right) \middle| \frac{Price_t = Price, FeeMax_t = FeeMax}{FeeMax_t = FeeMax} \right]. \quad (26)$$

► **Remark 26.** The first three terms in Equation (25) are exactly the myopic first-order condition from the baseline model. The final term is new. It is the marginal value of the effect of current participation on the next reserve stock. If $M_{t+1}(R, Price, FeeMax; s) \geq 0$, then forward-looking validators value reserve accumulation. The sign of the whole term is then governed by

$$FeeMax - \frac{N+1}{N} Sec \cdot s,$$

which is the marginal effect of one validator's additional participation on next period's reserve inflow. When aggregate participation is below the point at which extra participation sharply compresses fees, dynamic incentives push equilibrium security upward relative to the myopic benchmark. When the system is already close to maximum congestion, the intertemporal effect becomes weaker and can eventually reverse.

The next proposition shows that the runway logic survives the Markov extension.

► **Proposition 27** (Runway under a pure Markov equilibrium). *Suppose the pure symmetric Markov perfect equilibrium from Proposition 25 exists and denote its aggregate security policy at date t by*

$$s_t^M(R, Price, FeeMax).$$

Define the date- t security set

$$\mathcal{S}_t(\underline{s}) := \{(R, Price, FeeMax) \in \mathbb{R}_+ \times \mathbb{R}_{++}^2 : s_t^M(R, Price, FeeMax) \geq \underline{s}\}.$$

Then, along any equilibrium sample path, security is maintained through date T if and only if

$$(R_t, Price_t, FeeMax_t) \in \mathcal{S}_t(\underline{s}) \quad \text{for all } t = 0, \dots, T.$$

If, in addition, $s_t^M(R, Price, FeeMax)$ is weakly increasing in R for each fixed $(Price, FeeMax)$, then the dynamic reserve threshold

$$SecureR_t^M(Price, FeeMax; \underline{s}) := \inf\{R \geq 0 : s_t^M(R, Price, FeeMax) \geq \underline{s}\}$$

is well defined, and the runway criterion can be written as

$$R_t \geq SecureR_t^M(Price_t, FeeMax_t; \underline{s}) \quad \text{for all } t = 0, \dots, T.$$

The threshold $SecureR_t^M(\cdot)$ need no longer have a closed form, but once equilibrium participation is monotone in reserves, the safe region is still summarized by a reserve cutoff at each state and date.

6 Failure Probabilities and Expected Hand-Off Times

The main results identify the safe region state by state and path by path without committing to a particular stochastic law for price and demand. For quantitative policy analysis, however, one often wants explicit probabilistic outputs such as failure probabilities over a fixed horizon or expected time to fee-only operation. This section derives such metrics under a parametric specialization.

► **Assumption 28** (Lognormal state dynamics). *There exist constants $\mu_{Price}, \mu_{FeeMax} \in \mathbb{R}$ and $\sigma_{Price}, \sigma_{FeeMax} > 0$, a standard Brownian motion $(W_t)_{t \geq 0}$, and an i.i.d. sequence*

$$(\Delta W_t, \eta_t)_{t \geq 1}$$

of centered bivariate normal random vectors with unit marginal variances, where $\Delta W_t := W_t - W_{t-1}$ for $t \geq 1$, such that, for every $t \geq 0$,

$$\log Price_t = \log Price_0 + \left(\mu_{Price} - \frac{\sigma_{Price}^2}{2} \right) t + \sigma_{Price} \cdot W_t, \quad (27)$$

$$\log FeeMax_t = \log FeeMax_0 + \mu_{FeeMax} t + \sigma_{FeeMax} \sum_{k=1}^t \eta_k. \quad (28)$$

We write Φ for the standard normal cumulative distribution function (cdf).

► **Remark 29** (geometric Brownian motion). Assumption 28 is imposed only in this section. It strengthens Assumption 1 in order to convert the law-free runway criteria into explicit probabilistic metrics. The price process is a geometric Brownian motion observed at integer dates, while demand remains in a discrete-time lognormal specification. Correlation between token-price shocks and demand shocks is allowed through the joint law of $(\Delta W_t, \eta_t)$; the results below use only the normal marginal distributions and a union-bound argument.

Next, we convert the law-free stress-test criterion from Theorem 15 into an explicit statement.

► **Theorem 30** (Finite-horizon failure probability under lognormal dynamics). *Fix a horizon $T \in \mathbb{N}$ and confidence parameters $z_{Price}, z_{FeeMax} \geq 0$. Define deterministic lower envelopes by*

$$\underline{Price}_t(z_{Price}) := Price_0 \cdot \exp\left(\left(\mu_{Price} - \frac{\sigma_{Price}^2}{2}\right)t - \sigma_{Price} \cdot \sqrt{t} \cdot z_{Price}\right), \quad (29)$$

$$\underline{FeeMax}_t(z_{FeeMax}) := FeeMax_0 \cdot \exp\left(\mu_{FeeMax} \cdot t - \sigma_{FeeMax} \cdot \sqrt{t} \cdot z_{FeeMax}\right), \quad (30)$$

for $t = 0, \dots, T$. If

$$(1 - \rho)^t R_0 \geq \text{SecureR}(\underline{Price}_t(z_{Price}), \underline{FeeMax}_t(z_{FeeMax}); \underline{s}) \quad \text{for all } t = 0, \dots, T, \quad (31)$$

then

$$\mathbb{P}(\tau_{\text{fail}} > T) \geq 1 - T \cdot (\Phi(-z_{Price}) + \Phi(-z_{FeeMax})). \quad (32)$$

In particular, if

$$z_{Price} = \Phi^{-1}\left(1 - \frac{\alpha_{Price}}{T}\right), \quad z_{FeeMax} = \Phi^{-1}\left(1 - \frac{\alpha_{FeeMax}}{T}\right)$$

for some $\alpha_{Price}, \alpha_{FeeMax} \in (0, 1)$, then

$$\mathbb{P}(\tau_{\text{fail}} > T) \geq 1 - \alpha_{Price} - \alpha_{FeeMax}.$$

► **Proposition 31** (Hand-off probabilities and expected hand-off time). *Suppose Assumption 28 holds and define the fee-only demand cutoff*

$$FeeMax^{FO} := \frac{A_N \cdot s}{1 - \theta}. \quad (33)$$

Let

$$y_0 := \log FeeMax_0, \quad y^{FO} := \log FeeMax^{FO}.$$

Assume $\mu_{FeeMax} > 0$ and $FeeMax_0 < FeeMax^{FO}$. Then:

(i) The hand-off time satisfies

$$\tau_{\text{hand}} = \inf\{t \geq 0 : FeeMax_t \geq FeeMax^{FO}\} < \infty \quad \text{almost surely.}$$

(ii) For every horizon $T \in \mathbb{N}$,

$$\mathbb{P}(\tau_{\text{hand}} \leq T) \geq 1 - \Phi\left(\frac{y^{FO} - y_0 - \mu_{FeeMax}T}{\sigma_{FeeMax}\sqrt{T}}\right). \quad (34)$$

(iii) The expected hand-off time is finite and satisfies

$$\frac{y^{FO} - y_0}{\mu_{FeeMax}} \leq \mathbb{E}[\tau_{\text{hand}}] \leq 1 + \sum_{t=1}^{\infty} \Phi\left(\frac{y^{FO} - y_0 - \mu_{FeeMax}t}{\sigma_{FeeMax}\sqrt{t}}\right) < \infty. \quad (35)$$

(iv) If, in addition, the condition of Theorem 30 holds for the same horizon T , then

$$\mathbb{P}(\tau_{\text{hand}} < \tau_{\text{fail}}) \geq 1 - T(\Phi(-z_{Price}) + \Phi(-z_{FeeMax})) - \Phi\left(\frac{y^{FO} - y_0 - \mu_{FeeMax}T}{\sigma_{FeeMax}\sqrt{T}}\right). \quad (36)$$

Proposition 31 separates two distinct quantitative questions. The first is how quickly demand alone is likely to carry the protocol into the fee-only region. The second is whether that arrival is likely to happen before reserve failure. Positive demand drift makes eventual hand-off almost sure in this parametric environment, while the finite-horizon bounds show how drift and volatility translate into operational success probabilities.

The analytical results above are complemented by two numerical illustrations. Figure 2 plots three deterministic state paths built from the parametric law: one path fails before hand-off, one reaches fee-only sustainability gradually, and one hands off quickly. Figure 3 then shows how the quantitative runway picture changes with demand drift in the parametric environment. The left panel reports Monte Carlo estimates of survival, hand-off, and successful-transition probabilities at a fixed horizon, while the right panel plots the analytical lower and upper bounds on expected hand-off time from Proposition 31.

7 Conclusion and Future Work

Proof-of-stake systems with finite reserves face a reserve hand-off problem. Early validator rewards may rely on reserve subsidies, while long-run security must eventually be supported by transaction fees. This paper formalized this transition through a stochastic model with token-price risk, demand risk, and strategic validator participation. We solved the symmetric validator participation game and derived a closed-form equilibrium. Using this equilibrium, we obtained an explicit state-dependent reserve threshold. For each token price and demand





Figure 2 Illustrative runway paths under three deterministic state trajectories generated from the parametric law. In each panel, the solid blue curve is the reserve stock R_t and the dashed red curve is the state-dependent threshold $SecureR(\text{Price}_t, \text{FeeMax}_t; s)$. The vertical marker identifies the first economically relevant event along the path: hand-off when the threshold reaches zero, or failure when the reserve falls below the threshold. The adverse path crosses the failure boundary before reaching the fee-only region; the middle path reaches fee-only sustainability only after a prolonged transition; and the favorable path hands off early. The figure uses the same protocol parameters as Figure 1, with initial reserve $R_0 = 9$.



■ **Figure 3** Probabilistic runway metrics under the parametric state dynamics. The horizontal axis in both panels is the demand drift μ_{FeeMax} . In the left panel, the solid curves are Monte Carlo estimates of survival probability by horizon $T = 30$, hand-off probability by horizon $T = 30$, and the probability of successful hand-off before failure. The dashed curve is the analytical lower bound on hand-off by time T from Proposition 31. In the right panel, the two curves are the analytical lower and upper bounds on expected hand-off time from Proposition 31. The figure uses the same protocol parameters as Figure 1, with initial reserve $R_0 = 9$, $Price_0 = 1$, $FeeMax_0 = 1$, price drift $\mu_{Price} = 0.005$, price volatility $\sigma_{Price} = 0.16$, and demand volatility $\sigma_{FeeMax} = 0.10$. Higher demand drift increases survival and hand-off probabilities and narrows the expected time to hand-off.

state, the threshold gives the minimal reserve stock necessary and sufficient to sustain a target security level. This threshold separates states in which security is infeasible, states in which security depends on the reserve, and states in which fees alone are sufficient. We then used the threshold to study the dynamic hand-off problem. Security is maintained along a path exactly while the reserve remains above the state-dependent threshold. Failure occurs when the reserve first falls below this threshold. A successful hand-off occurs when the system reaches the fee-only region before failure. This converts reserve adequacy into a state-dependent hitting-time problem. We also studied forward-looking validator incentives. In the Markov extension, current participation affects both current rewards and the next reserve stock, and therefore changes future reward opportunities. We proved finite-horizon Markov perfect equilibrium existence and derived the dynamic first-order condition that captures this continuation-value effect. Finally, under lognormal price-demand dynamics, we translated the runway analysis into finite-horizon failure-probability bounds, hand-off probabilities, and bounds on expected hand-off time.

The main conclusion is that reserves are not neutral accounting devices. Their security value depends on the current state of the system. A protocol with a large nominal reserve can still be close to failure after an adverse price or demand shock. Conversely, a protocol with a smaller reserve can be safe once demand is strong enough to support the target security level through fees alone. Reserve policy should therefore be evaluated through state-dependent runway analysis rather than through nominal depletion dates or steady-state reward ratios alone.

The model made several simplifying assumptions. Token price and demand were treated as exogenous, so the analysis did not model token valuation or user adoption. Validators were symmetric in the baseline model, so the threshold should be interpreted as a benchmark

rather than as a full description of heterogeneous operators. The finite-horizon stress test was deliberately conservative because it lower-bounded the reserve path by ignoring future fee inflows. These assumptions made it possible to isolate the reserve hand-off mechanism and obtain explicit threshold conditions.

A natural next step is adaptive reserve policy. In this paper, the policy parameters (θ, ρ) were fixed. This provided a clean benchmark for analyzing the security-runway mechanism. In principle, a protocol may want payout and fee-retention rules to respond to the current reserve, token price, and demand state. For example, the protocol could reduce reserve payouts after adverse price realizations, increase fee retention when the reserve is low, or change its policy once demand approaches the fee-only region.

Such an extension would turn the analysis from policy evaluation into policy design. Governance would choose a state-dependent rule of the form

$$(\theta_t, \rho_t) = \phi(R_t, Price_t, FeeMax_t)$$

to extend runway, reduce failure risk, or reach fee-only sustainability subject to incentive and credibility constraints. Developing such a theory would require combining the equilibrium framework in this paper with tools from dynamic mechanism design or stochastic control. We view this as an important direction for future work.

References

- 1 Algorand. The Algo, 2025. Accessed: 2nd June 2026. URL: <https://algorand.co/technology/the-algo>.
- 2 All About XRP. XRP Supply & Inflation Explained, 2025. Accessed: 27th May 2026. URL: <https://allaboutxrp.com/learn/xrp-supply-explained>.
- 3 Avalanche. Avax token, 2026. Accessed: 27th May 2026. URL: <https://www.avax.network/about/tokens>.
- 4 Cardano Docs. Cardano monetary policy. <https://docs.cardano.org/about-cardano/explore-more/monetary-policy/>. Accessed March 9, 2026.
- 5 Miles Carlsten, Harry Kalodner, S. Matthew Weinberg, and Arvind Narayanan. On the instability of bitcoin without the block reward. In *Proceedings of the ACM SIGSAC Conference on Computer and Communications Security (CCS)*, pages 154–167, New York, NY, USA, 2016. Association for Computing Machinery. doi:10.1145/2976749.2978408.
- 6 Jonathan Chiu and Thorsten V. Koepl. The economics of cryptocurrency: Bitcoin and beyond. *Canadian Journal of Economics/Revue canadienne d'économique*, 55(4):1762–1798, 2022. doi:10.1111/caje.12625.
- 7 Lin William Cong, Ye Li, and Neng Wang. Tokenomics: Dynamic adoption and valuation. *The Review of Financial Studies*, 34(3):1105–1155, 2021.
- 8 Lin William Cong, Ye Li, and Neng Wang. Token-based platform finance. *Journal of Financial Economics*, 144(3):972–991, 2022.
- 9 Bernardo David, Peter Gaži, Aggelos Kiayias, and Alexander Russell. Ouroboros praos: An adaptively-secure, semi-synchronous proof-of-stake protocol. Cryptology ePrint Archive, Paper 2017/573, 2017. URL: <https://eprint.iacr.org/2017/573>.
- 10 Giulia Fanti, Leonid Kogan, Sewoong Oh, Kathleen Ruan, Pramod Viswanath, and Gerui Wang. Compounding of wealth in proof-of-stake cryptocurrencies. In *Financial Cryptography and Data Security*, 2019.
- 11 Hedera Council. Treasury management, 2025. Accessed: 27th May 2026. URL: <https://hederacouncil.org/treasury>.
- 12 Nicolas Houy. The economics of Bitcoin transaction fees. Working Paper 1407, GATE Lyon Saint-Étienne, 2014. doi:10.2139/ssrn.2400519.

- 13 Gur Huberman, Jacob D. Leshno, and Ciamac C. Moallemi. Monopoly without a monopolist: An economic analysis of the bitcoin payment system. *Review of Economic Studies*, 88(6):3011–3040, 2021. doi:10.1093/restud/rdab014.
- 14 Samuel Häfner. Optimal decentralization and service provision on a blockchain platform with market frictions, September 2023. URL: https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3954773.
- 15 Felix Irresberger and Ruomei Yang. Coin concentration of proof-of-stake blockchains. *Economics Letters*, 229:111219, 2023. doi:10.1016/j.econlet.2023.111219.
- 16 Aggelos Kiayias, Philip Lazos, and Paolo Penna. Single-token vs two-token blockchain tokenomics. In Zeta Avarikioti and Nicolas Christin, editors, *7th Conference on Advances in Financial Technologies (AFT)*, volume 354 of *LIPICs*, pages 22:1–22:22. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2025. doi:10.4230/LIPICs.AFT.2025.22.
- 17 Aggelos Kiayias, Philip Lazos, and Jan Christoph Schlegel. Would friedman burn your tokens? In *28th International Conference on Financial Cryptography and Data Security (FC)*, pages 231–247, Berlin, Heidelberg, 2024. Springer-Verlag. doi:10.1007/978-3-031-78676-1_13.
- 18 Aggelos Kiayias, Alexander Russell, Bernardo David, and Roman Oliynykov. Ouroboros: A provably secure proof-of-stake blockchain protocol. Cryptology ePrint Archive, Paper 2016/889, 2016. URL: <https://eprint.iacr.org/2016/889>.
- 19 Kevin Liao and Jonathan Katz. Incentivizing blockchain forks via whale transactions. In Michael Brenner, Kurt Rohloff, Joseph Bonneau, Andrew Miller, Peter Y.A. Ryan, Vanessa Teague, Andrea Bracciali, Massimiliano Sala, Federico Pintore, and Markus Jakobsson, editors, *Financial Cryptography and Data Security*, pages 264–279, Cham, 2017. Springer International Publishing. doi:10.1007/978-3-319-70278-0_17.
- 20 Shashank Motepalli and Hans-Arno Jacobsen. Decentralization in PoS Blockchain Consensus: Quantification and Advancement. *arXiv preprint arXiv:2504.14351*, 2025. doi:10.48550/arXiv.2504.14351.
- 21 Christina Ovezik, Dimitris Karakostas, Mary Milad, Daniel W. Woods, and Aggelos Kiayias. Sok: Measuring blockchain decentralization. In Marc Fischlin and Veelasha Moonsamy, editors, *23rd International Conference on Applied Cryptography and Network Security (ACNS)*, Lecture Notes in Computer Science, pages 184–214. Springer, 2025. doi:10.1007/978-3-031-95761-1_7.
- 22 Paolo Penna and Manvir Schneider. Reserve depletion and security runway in proof-of-stake systems. *arXiv preprint arXiv:2606.03587*, 2026. doi:10.48550/arXiv.2606.03587.
- 23 Fahad Saleh. Blockchain without waste: Proof-of-stake. *Review of Financial Studies*, 34(3): 1156–1190, 2021. doi:10.1093/rfs/hhaa075.
- 24 Itay Tsabary and Ittay Eyal. The gap game. In *Proceedings of the 2018 ACM SIGSAC Conference on Computer and Communications Security, CCS '18*, pages 713–728, New York, NY, USA, 2018. Association for Computing Machinery. doi:10.1145/3243734.3243737.
- 25 Zhilin Wang, Qin Hu, Yawei Wang, and Yinhao Xiao. Transaction pricing mechanism design and assessment for blockchain. *High-Confidence Computing*, 2(1):100044, 2022. doi:10.1016/j.hcc.2021.100044.