Cryptoeconomic Systems

Competitive Equilibria Between Staking and On- chain Lending

Tarun Chitra¹

¹Gauntlet

Published on: Nov 30, 2020

License: [Creative Commons Attribution 4.0 International License \(CC-BY 4.0\)](https://creativecommons.org/licenses/by/4.0/)

Abstract

Proof of Stake (PoS) is a burgeoning Sybil resistance mechanism that aims to have a digital asset (“token”) serve as security collateral in crypto networks. However, PoS has so far eluded a comprehensive threat model that encompasses both Byzantine attacks from distributed systems and financial attacks that arise from the dual usage of the token as a means of payment and a Sybil resistance mechanism. In particular, the existence of derivatives markets makes malicious coordination among validators easier to execute than in Proof of Work systems. We demonstrate that it is also possible for on-chain lending smart contracts to cannibalize network security in PoS systems. When the yield provided by these contracts is more attractive than the inflation rate provided from staking, stakers will tend to remove their staked tokens and lend them out, thus reducing network security. In this paper, we provide a simple stochastic model that describes how rational validators with varying risk preferences react to changes in staking and lending returns. For a particular configuration of this model, we provide a formal proof of a phase transition between equilibria in which tokens are predominantly staked and those in which they are predominantly lent. We further validate this emergent adversarial behavior (*e.g.*, reduced staked token supply) with agent-based simulations that sample transitions under more realistic conditions. Our results illustrate that rational, non-adversarial actors can dramatically reduce PoS network security if block rewards are not calibrated appropriately above the expected yields of on-chain lending.

1. Introduction

There is currently an intense effort to improve the scalability of blockchains and other decentralized value systems known as crypto networks. These networks use cryptographic proofs and game-theoretic constructions to provide tamper-resistant updates to a global ledger. While there are a variety of research and engineering challenges in setting up these systems, one of the major bottlenecks to network throughput is the cost of Sybil resistance mechanisms within a decentralized consensus protocol. Proof of Work (PoW) networks achieve Sybil resistance by requiring consensus-participating nodes to provably burn energy to compute many iterations of a particular cryptographic hash function. PoW, while effective and permissionless, expends a large amount of natural resources and has resulted in concentrated ownership of the underlying digital assets (*e.g.*, Bitcoin). Proof of Stake (PoS) was first introduced as an alternative in a 2012 BitcoinTalk post [1] that showed the equivalence between a PoW miner who could immediately reinvest her block rewards into hash power within the network and a PoS validator who can reinvest their validation earnings into network security. PoS works by instead allowing users to ‘lock’ a digital asset, known as a token, into a smart contract that provides them with token-denominated returns in exchange for validating transactions and providing network security. Using shared, verifiable randomness, all network participants can use

a multi-party computation protocol to sample the distribution of asset ownership locked into the contract, and choose participant(s) who receive the block reward emitted by the network. This is analogous to how PoW can be thought of as a protocol that samples the distribution of hash power to choose the next block producer [2]. One of the main benefits of PoS is that one does not have to commit a costly natural resource to participate in the network. Instead, a purely digital asset is used as collateral for the network and the network can control its supply to provide the desired properties. For an introductory background on PoS protocols and their complex security models, please see [3][4][5][6][7].

In this paper, we show that these purported benefits do not come for free. As PoS algorithms inherently connect a decentralized network's security with the capital cost of a digital asset, PoS protocols tie their security to the cost of capital rather than to the cost of a natural resource. Volatility in the cost of capital, which is usually higher than that of natural resources [8], can have adverse effects on capital commitments to PoS networks. The main result we show is that alternative sources of yield can drive staking token capital allocators to collectively drain a network's security, akin to a bank run. In particular, we find that PoS in deflationary systems is unstable and unlikely to work, and that for more reasonable inflation rates, the effectiveness of PoS depends on the relationship between staking and lending rates. This relationship should inform the further design of PoS systems, especially as a large number of networks are launching in 2020.

Transitioning Security from PoW to PoS

The move from PoW to PoS presents a plethora of challenges. In a PoS system, the network relies on participants who are staked in the system to stay online in order to achieve liveness. In practice, this is implemented by slashing participants—redistributing or burning a participant's stake that is committed for validation rewards when they perform a malicious act—who go offline or miss a block that they are supposed to produce. Moreover, there are attacks that are unique to PoS such as the *nothing-at-stake* and *long-range* attacks [9][10]. These attacks are impossible in PoW, as resource costs of digital assets are practically zero, especially when compared to costs of natural resources [11]. Lastly, as the asset used for staking is also the medium of exchange, a malicious validator needs to only aggregate 33% of the token to perform a Byzantine attack.¹

However, if there exist physically settled futures contracts on PoS tokens, then it is possible for an attacker to buy futures that allow for staking participants to sell their staked token in the future. This attacker can aggregate this stake and upon reaching an attack threshold, begin to perform a double spend or other malicious attack [12]. As these derivatives can be settled off-chain (*e.g.*, using a centralized exchange like BitMEX or Deribit), monitoring of this type of attack can be difficult. In PoW, one would need to aggregate the hash power needed to produce 50% of the network's hashrate, which is a much harder task that relies on aggregating data centers, specialized hardware, cheap electricity,

and a favorable country of residence. Moreover, PoS systems are vulnerable to *financial attacks*, or attacks that utilize the fact that PoS tokens serve as both the instrument of security and as the medium of exchange. Such attacks are often feasible due to emergent and unexpected coordination between participants in a PoS network and an external market.

Given the vulnerability of PoS to cartel-like behavior that can be coordinated via an external market, one might naturally ask if there are also any endogenous financial risks on PoS protocols that support smart contracts. Recently, there has been an uptick in interest in Decentralized Finance (‘DeFi’), which uses smart contracts to implement standard financial primitives in a purely on-chain manner [13]. These primitives, such as exchanges [14], lending [15], and stable reserve currencies [16][17][18], decentralize banking functions by creating incentives that encourage rational participants to receive arbitrage profits for maintaining the system’s security, while also meting out financial punishments for misbehavior. Instead of explicitly punishing fraud via legal recourse, these protocols use purely financial modes of recourse to encourage network participation and growth. One of the biggest sectors within DeFi is the on-chain lending market, in which the largest single platform is Compound [15], an Ethereum smart contract that allows users to lend and borrow assets that conform to the ERC-20 token standard. The Compound smart contract has held up to US \$175 million of assets, has had over 40% of the float of the Dai stablecoin [18], and saw double digit asset growth during 2019. Given that Ethereum is likely to transition to PoS soon, one must evaluate: are there any financial attacks against chain security that result from an on-chain lending system? A simple Gedankenexperiment (a thought-experiment) for answering this question from the view of on-chain lending might be of the form:

Suppose that we assume that validators are rational financial agents. Would they not simply move their assets between staking and on-chain lending, depending on which has a higher yield?

In particular, it is clear that there is a relationship between the price of capital availability and participant’s willingness to stake, as stakers have to earn more than a risk-adjusted market rate on their staked capital. However, unlike Proof of Work, there are no physical limits that prevent validators in staking networks from rapidly moving their assets into higher-yielding activities. On-chain lending, such as Compound, makes this particularly efficient, as validators simply have to post a single Ethereum transaction to unbond their tokens and begin earning yield within a single block time. High lending yields would likely lead to a reduction in network security (*e.g.*, a financial attack) as these yields would encourage rational actors to unexpectedly coordinate and reduce network security by optimizing for financial gain.

We answer these questions with a stochastic model that can be theoretically solved in certain situations and is easily simulated via Monte Carlo methods. Our aim is not to model realistic networks parameters perfectly, but rather to show that even in the most simplified model of agents optimizing portfolios composed of staked and lent tokens, on-chain lending can cause dramatic volatility in

network security. In particular, we construct an agent-based model, where each network participant is represented via an agent with a utility and decision function. Agent-based modeling has been previously used for modeling censorship properties in sharded PoW chains [19] and can serve as a conduit for comparing theoretical results to empirical data in a statistically rigorous manner. Our model for rational network agents involves having each participant view their total token wealth as a two-component portfolio of tokens staked and tokens lent. We assume that the agents have different risk preferences and locally optimize their token portfolios using mean-variance optimization [20], which allows for agents to adjust their portfolios based on observed returns and risk preferences. [Figure 2](#) illustrates that there is a phase transition in inflation rates that leads to tokens going from being predominantly lent in deflationary regimes to predominantly staked in inflationary regimes. Moreover, [Figure 3](#) shows that the spread between borrowing and lending rates is significantly worse for deflationary PoS assets, further confirming the existence of a phase transition. We explicitly state our simplifying assumptions in [Section 2](#) and note that agent-based modeling allows one to relax these assumptions and analyze how these results carry over to real PoS protocols, such as Cosmos [21] and Tezos [22]. In [Section 4](#), we prove properties of this model that match the observed phase transition from [Figure 2](#) and [Figure 5](#). A single sample path is depicted in [Figure 1](#), which visually show that agents participating in on-chain lending and staking can cause a “flipping,” where there are more assets lent out than staked. The observed volatility in the amount of assets staked is tantamount to dramatically reducing the cost of taking over a staking network, which implies that the security model of PoS networks needs to account for attacks on the network that stem from reduced cost of capital. The combination of theoretical and simulation-based results demonstrate that the threat model for PoS networks needs to be expanded to include financial attacks that result from yield competition with on-chain financial products.

For provenance, note that the mathematical notation used in this paper is documented in [Appendix A](#).

Figure 1: Staking and lending with Compound's smart contract parameters (**top**) and more aggressive parameters with yields of up to 75%. The y-axis in both figures is the percentage of the total outstanding token supply that is allocated to staking versus that which is allocated to lending. This simulation had 512 agents whose parameters were sampled as in [Equation 3](#) and [Equation 4](#). Note that the two figures have different time-scales to illustrate how the parameters of the model (see [Section 3.2.1](#)) affect the oscillations between staked and lent supply.

Figure 2: Plots of relative amount of staking to lending, with inflation increasing in a clockwise manner. When these plots are negative, there is, on average, more lending than staking and vice-versa, when positive. When we are deflationary (**top two rows**), we see that regardless of the lending rate parameters (the x-axis is β_0 and the y-axis is β_1 from [Equation 1](#)), tokens on average lent. This empirically demonstrates that there is a phase transition from a lending-dominated regime (which is deflationary) to one that is staking-dominated and inflationary. This simulation also uses 512 agents and the titles state the inflation rate (which is deflationary if it is less than 1), block reward, half-life, and borrow threshold (see [Section 5.3](#)) and each sample involves an average over a hundred random seeds.

Figure 3: Plots of the relative spread between borrowing and lending rates. If the value is higher, then borrowing costs more than lending. Note that in the deflationary regime, as the lending rates increase (e.g., as β_0, β_1 from Equation 1 increase), the spread also increases, whereas this effect is far more muted. The simulation parameters are the same as Figure 2.

2. Protocol Assumptions

In order to simplify our model, we will make some simplifying assumptions. These assumptions hold for all models analyzed in this paper and they focus on properties of the underlying distributed ledger rather than properties about the economic behavior of participants, which is described in the proceeding section.

Assumption 1 (Cryptographic)

All sampling processes use true randomness and not pseudorandomness.

This differs from many standard cryptographic threat models that assume pseudorandomness and provide an ϵ -approximation to true random sampling [23]. In the Appendix, we use this property to ensure that our PoS algorithm is non-anticipating (e.g., adapted to a suitably chosen filtration), allowing for conditional probabilities to be computed without having to aggregate the ϵ error terms. One can remove this assumption, in a similar manner to [24], with a significant increase in proof complexity.

Assumption 2 (Distributed Systems)

All communication between participants is synchronous.

This can be relaxed to partial synchrony, with an increase in complexity in the proofs and simulations described in the following section.

Assumption 3 (Identity)

All pseudonymous identities are known by all participants and the number of participants (measured by unique addresses), n , is fixed.

Assumption 4 (Time)

The entire system will update at discrete time intervals, with each tick to be thought of as a block update.

While participants will likely execute strategies in continuous time, assuming discrete time evolution ensures that participants only respond to event updates that are received on-chain. One can also remove this assumption at the cost of increased variance using ‘Poissonization’ techniques [25] and ‘sleepy consensus’ assumptions [26].

Assumption 5 (No External Markets)

Only on-chain lending, borrowing, and staking will be considered.

We are ignoring off-chain lending (e.g., OTC desks or lending businesses, such as Galaxy and Tagomi) and are assuming that there always exists block space for any participant’s action to succeed. We are omitting this to reduce the complexity of our model, as off-chain lending has varied pricing models and term structures.

Assumption 6 (Money Supply is Deterministic)

The block reward at time t , R_t and the money supply at time t , $S_t = \sum_{s=0}^t R_s$ are deterministic and known to all participants.

In particular, we avoid assuming that there exist governance mechanisms for changing the block reward, which have been proposed for protocols such as Algorand [27] and Celo [16].

Assumption 7 (No Transaction Fee Revenue)

The only revenue that validators receive from staking comes from the block reward.

This is a model assumption that removes the complexity of modeling crypto network fee markets, which currently have unstable dynamics and are poorly understood.

Assumption 8 (PoS)

The following properties hold for our idealized PoS algorithm:

- *No compounding:* The PoS mechanism uses epoch-based sampling and does not immediately reinvest block rewards. This is to avoid the concentration behavior described in [28]. One can relax our model to handle Pólya urn processes, at the cost of significantly worse variance.
- *Single validator per block:* To simplify the model, we avoid using committees [27][29][16][22][21][30][31][7] and verifiers [32][33] as they add more variance and make both formal and simulation methods more difficult.

- *Constant slashing probability*: We assume that each validator has a fixed probability p_{slash} of being slashed on a block that they produce. This is simplistic as it assumes that all validators have the same chance of being slashed, regardless of stake and validation history. However, in practice, we have seen very few live slashes and this model encourages simpler formal and simulation analysis.
- *Sharded state is synchronously traversed*: Any sharded state in our PoS blockchain is read synchronously and we assume linearizability in our blockchain. This can be relaxed to non-linearizable protocols like Avalanche [29], but will depend on the scoring function used for each branch.
- *No unbonding period*: We do not assume that the PoS protocol has an unbonding period like that of Cosmos [21] or Tezos [22].

3. Model

3.1 PoS Model

Let S_t be the total outstanding token supply of a PoS protocol at time t and let $S_t = \zeta_t + \ell_t$, where ζ_t is the number of staked tokens and ℓ_t is the number of lent tokens. We use a simple model of a PoS system that samples a single block producer from a discrete time-series of stake distributions, denoted by $\pi_{\text{stake}}(t) \in \zeta_t \Delta^n$, where Δ^n is the n -dimensional probability simplex (see [Appendix A](#)). The main input parameters are:

- $\pi_{\text{stake}}(0)$: Initial asset distribution
- R_t : Staking block reward at block height t
- $S(i, \pi_{\text{stake}})$: Slash that validator i receives if the staking distribution is π_{stake}

The formal specification of the algorithm can be found in [Algorithm 1](#) in [Appendix B](#). The algorithm state includes the current stake distribution, the current epoch's reward set, the current epoch's slash set, and the current block time. At a high-level, for each block, we select a validator who should receive a block reward and decide if they are to be slashed by flipping a coin with probability p_{slash} . If they are slashed, we add their *id* and the amount that they are to be slashed to the current epoch's slash set. Otherwise, we add them to the block reward set. The algorithm updates the stake distribution on a per-epoch basis.

3.2 Lending Model

We will study two models for lending, one involving an explicit model for borrowing demand and one that is implicit. The implicit model, which is simpler, provides a prototype for constant borrowing demand and is amenable to theoretical results. On the other hand, the model involving an explicit description of borrowing demand is more realistic and amenable to being fit by historical data.

3.2.1 Two-State Lending Model

In the two-state lending model, we track the time evolution of two token distributions, $\pi_{\text{stake}}(t), \pi_{\text{lend}}(t) \in S_t \Delta^n$, which respectively represent the distribution of staked tokens and those locked in a lending contract. The i th component of the stake distribution, $\pi_{\text{stake}}(t)_i$, corresponds to the amount of tokens that the i th agent has staked. Each agent has a wealth $W_i(t)$ at time t that is equal to the sum of their portfolio of staked and lent tokens, i.e., $W_i(t) = \pi_{\text{stake}}(t)_i + \pi_{\text{lend}}(t)_i$. By definition, the total lending supply is equal to the sum of all lent portfolios, i.e., $\ell_t = \|\pi_{\text{lend}}(t)\|_1$, and the total money supply is equal to the sum of all portfolios, i.e., $S_t = \|\pi_{\text{stake}}(t)\|_1 + \|\pi_{\text{lend}}(t)\|_1$. At each time step, agents update their portfolios based on returns accrued from the previous time step and, after portfolios are updated, the lending rate, γ_t , is updated based on the total amount lent. This means that we are making two implicit assumptions:

1. *Constant relative borrowing demand*: We are assuming that the ratio of borrowing demand (represented via a quantity of tokens) to the total token supply stays constant, as the rate only depends on the lending and staking supplies. Formally, this means that the demand at time t is equal to kS_t
2. *Flows are the only determining factor*: Participants who move tokens from staking to lending or vice versa are the only causes for changes to the lending rate

We draw inspiration from Compound [15], which provides a simple formula for the borrow lending rate, β_t , and the lending rate, γ_t . The Compound model computes a utilization rate U_t at block height t , which is the ratio of the borrowing demand to the token supply locked in the contract, and uses that to update formulas for β_t, γ_t . Mathematically, they define the utilization rate as

$$U_t = \frac{kS_t}{\ell_t + kS_t}$$

We compute the borrow and lend rates using the following formulas, where $\beta_0, \beta_1 \in (0, 1)$ are interest-rate parameters and $\gamma_0 \in (0, 1)$ is a measure of the spread between lending and borrowing (i.e., $1 - \gamma_0$ is the relative spread).

$$\beta_t = U_t(\beta_0 + \beta_1 U_t) \tag{1}$$

$$\gamma_t = (1 - \gamma_0)\beta_t \tag{2}$$

For reference, the Compound V2 contract uses the values $\beta_0 = 5\%$ and $\beta_1 = 45\%$. As depicted in Figure 2, there can be an enormous amount of volatility in the fraction of the token supply that is lent, $\frac{\ell_t}{S_t}$. In Section 4, we prove tail bounds on the inflows and outflows of lent tokens over a time step, e.g., $\Pr[|\ell_t - \ell_{t-1}| > \epsilon S_t]$, that explicitly depend on the block reward, R_t , and the interest rate parameters. These bounds suggest that even for the overly simplistic setting of the two-state model,

PoS protocols need to carefully choose their block rewards if they desire to have a large fraction of the outstanding token supply staked at all times.²

3.2.2 Three-State Model

Instead of assuming that there is constant relative borrowing demand, we can relax this by specifying an additional distribution, $\pi_{\text{borrow}}(t)$, such that $\pi_{\text{token}} = \pi_{\text{stake}} + \pi_{\text{lend}} + \pi_{\text{borrow}}$ and $S_t = \zeta_t + \ell_t + \xi_t$, where $\xi_t = \|\pi_{\text{borrow}}(t)\|_1$ is the total amount borrowed at time t . In this world, the utilization ratio is now defined as:

$$U_t = \frac{\xi_t}{\ell_t + \xi_t}$$

Formally analyzing this model has a variety of difficulties that stem from that fact that we have to explicitly model the borrow demand distribution and disentangle how it couples to each participant's local model of risk, which is defined in the next section. In [Section 5](#), we simulate this model with a variety of different borrowing demand distributions, but all formal proofs that follow only analyze the two-state lending model. Since the pool of borrowers (e.g., arbitrageurs) is often disjoint from the pool of stakers and lenders, we will make the following assumption:

Assumption 9 (Borrowing demand is independent of staking and lending)

The borrowing demand distribution π_{borrow} is probabilistically independent of the lending and staking distributions.

3.3 State Transition

The final task needed to completely specify this model is to define the state transition rule that sends $\pi_{\text{lend}}(t), \pi_{\text{borrow}}(t), \pi_{\text{token}}(t)$ to $\pi_{\text{lend}}(t+1), \pi_{\text{borrow}}(t+1), \pi_{\text{token}}(t+1)$. As per [Assumption 9](#), the evolution of $\pi_{\text{borrow}}(t)$ is independent of staking and lending, and will be specified separately in [Section 5](#). We therefore need to specify state transition rules on a per-agent basis, where each agent's state is their token portfolio $(\pi_{\text{stake}}(t)_i, \pi_{\text{lend}}(t)_i)$. Traditionally, the strategy space of rational actors is described via an expected utility function that an agent aims to maximize by taking various allowable actions. Before specifying the strategy space that we will sample, let's consider a few examples to motivate the need for agents who have varying risk preferences. If at time t , staking is returning more than lending and every agent moves their entire portfolio from lending to staking, then we will observe a correlated spike in lending rates ℓ_t will go to zero and $U_t = 1$. Moreover, the relative staking reward to any agent will decrease as S_t will increase by the amount of tokens that flow from lending to staking as the expected return on an epoch for an agent is their staked tokens divided by the token supply, $\frac{\pi_{\text{stake}}(t)_i}{S_t - \ell_t}$. Thus the greedy strategy of moving all of one's assets to the higher yielding activity and causing drastic swings in the relative yields of staking and lending is unstable and doesn't accurately represent reality, where token holders have differing risk preferences and will not

immediately move their entire portfolio from staking to lending (or vice-versa). Furthermore, cryptocurrency holders are often looking for returns that are multiples of their initial investment and have a long time-preference [34]. In our evolution of each agent's portfolio, we also assume that an agent making a decision at time t can only use the information about all portfolios up to time t , $\{\pi_i(t) : i \in [n]\}$, and the implied rate γ_t . Since we are dealing with on-chain lending only, this assumption says that players cannot use strategies that look into the future and that all agent portfolios are public. In order to capture strategies that are independent of front-running and latency arbitrage, we will make the following assumption:

Assumption 10 (Martingale Ordering)

There exists an $[n]$ -valued martingale Z_t that chooses the ordering in which participants are allowed to update their portfolios at time t to those at time $t + 1$.

Under this assumption, agents receive no advantage in expected returns by trying to predict when their strategy is executed (e.g., is agent 1's strategy executed before agent 2's strategy because agent 1 has more staked than agent 2?). This assumption is reasonable as our goal is to figure out if rational, but non-Byzantine, agents will cause PoS network security to decrease when on-chain lending activity is sizeable.

3.3.1 Optimal Portfolio Construction

Mean-variance methods, pioneered by Markowitz's Nobel Prize winning work on portfolio theory [20], provide a way for rational traders of risky assets to construct portfolios that trade-off individual preferences for maximizing returns with those of risk minimization. These methods, which are the backbone of the majority of trillions of dollars of passive portfolios and statistical arbitrage strategies, provide a simple, easy-to-solve model that involves two parameters for constructing portfolios of n assets: an expected return vector $\mu \in \mathbb{R}^n$ and a positive-definite covariance matrix $\Sigma \in S_+^{n \times n}$. Given these parameters, one solves a strongly convex program that aims to compute the fraction of an agent's wealth that should be allocated to each asset while ensuring that the sum of the allocations is 1, and each entry is positive. In particular, the seminal work of Markowitz aimed to optimize the quadratic form $f : \mathbb{R}^n \times \mathbb{R}^n \times \mathbb{R} \times S_+^{n \times n} \rightarrow \mathbb{R}$, $f(w, \mu, \lambda, \Sigma) = w^T \Sigma w - \lambda \mu$, where λ is a parameter that controls the riskiness of the output portfolio and w is the portfolio allocation. As λ is varied, the *efficient frontier* of admissible portfolios is defined as the surface $S(\mu, \Sigma) = \{w \in \mathbb{R}^n : \exists \lambda \text{ such that } w \in \arg \min f(w, \mu, \Sigma, \lambda)\}$. The original work of Markowitz [20] focused on the single-period allocation problem, where an investor aims to find the optimal portfolio over a single time-period, which corresponds to assuming that μ and Σ do not change over time. Further work on multiple period [35] and continuous-time methods [36] for mean-variance optimization allow for μ and Σ to vary as functions of time, with the continuous-time methodology drawing μ from an Itô process, such as a solution to the Black-Scholes equation for options pricing. As

blockchain systems have incremental updates with independent games per update (e.g., transaction fee markets can differ wildly from block to block), we will necessarily have to consider the multiple period model and define how each agent's mean vector and risk-preferences evolve over time. Finally, we note that our methodology is directly comparable to that of multistrategy backtesting in quantitative trading [37].

We will assume that each agent treats their token wealth, $(\pi_{\text{lend}}(t)_i, \pi_{\text{stake}}(t)_i)$, as a Markowitz portfolio and updates, on receiving tokens from staking and lending, an estimate for a time-dependent return vector $\mu_i(t) \in \mathbb{R}^2$. We also assume that each agent has a different, time-independent covariance matrix Σ_i that is drawn from a random matrix ensemble.³ In other words, the expected return vector adjusts with time (it depends on the staked and lent quantities), while the covariance stays fixed in time. As long as there is some variance in the chosen random matrix ensemble (e.g., $\exists i, j \in [n]$ such that $\Pr[\Sigma_i \neq \Sigma_j] > 0$), then the dynamics will not deadlock into a state where all participants end up with the same portfolio (e.g., all outstanding tokens are 100% allocated to staking or lending). If agents simply move all of their assets from one pool to another, as opposed to some risk-adjusted proportion, then the system can deadlock quickly when borrowing demand is constant. Using the notation of the prequel, we define $\mu_i(t) = \mu_i(\pi_{\text{stake}}(t), \pi_{\text{lend}}(t))$ as:

$$\mu_i(t) = \begin{bmatrix} \mu_{\text{stake}}(t)_i \\ \mu_{\text{lend}}(t)_i \end{bmatrix} = \begin{bmatrix} \frac{\pi_{\text{stake}}(t)_i}{S_t - \ell_t} \\ \gamma_t \end{bmatrix} = \begin{bmatrix} \frac{\pi_{\text{stake}}(t)_i}{\zeta_t} \\ \gamma_t \end{bmatrix} \quad (3)$$

Let us motivate this choice of expected return vector. Recall that the return vector is supposed to represent the relative rate of return (usually referred to as an alpha in the quantitative finance literature) over a riskless asset. In this situation, the riskless asset is holding our tokens (as bearer instruments) and we expect to earn yields⁴ of $\frac{R_t}{S_t - \ell_t}$ and γ_t for staking and lending, respectively. However, note that the former yield can be greater than 1 (when $S_t - \ell_t < R_t$), whereas the latter yield cannot as per Equation 1. In order for Markowitz optimization to be well-defined, we need to choose yields that are directly comparable (e.g., have the same range). Moreover, we note that the naïvely calculated staking yield loses the dependence on the i th party's current wealth. Since our system allows for slashing and each validator incurs a variance in reward proportional to the inverse square root of the epoch length, it is preferable to find a yield that depends on the i th parties wealth to reflect this variance. The simplest estimator for a staking validator's yield that is in $[0, 1]$ (comparable to lending yield) and has a component of the variance is the probability of winning the block reward, $\frac{\pi_{\text{stake}}(t)_i}{S_t - \ell_t}$, which is exactly what Equation 3 describes.

In order to describe the covariance matrix, we will first state an assumption that aims to connect the variances in the system with the expected staked and lent times:

Assumption 11

The i th agent's covariance matrix Σ_i will be static (e.g., does not vary with block height), diagonal, and will have variances connected to the expected stake time τ_{stake} and expected lent time τ_{lend} .

Formally we use the following model for the covariance matrix:

$$\Sigma_i^{-1} = \begin{bmatrix} \alpha_i & 0 \\ 0 & \beta_i \end{bmatrix} \quad (4)$$

$$\alpha_i \sim \text{Exp}(\tau_{\text{stake}}) \quad (5)$$

$$\beta_i \sim \text{Exp}(\tau_{\text{lend}}) \quad (6)$$

One can interpret τ_{stake} as the expected epoch length (see [Algorithm 1](#)), while τ_{lend} represents the withdrawal window under which a lender can remove their tokens from the lent pool.⁵ By connecting the covariance matrix to these quantities, we encode the connection between a validator's expected risk preference and the time that capital is locked into either staking or lending. [Equation 4](#) represents the fact that risk and time preferences are, to first-order, inversely correlated (e.g., you are willing to lock up capital for the longest duration in the least risky assets) [\[38\]](#). To explain these choices of random variables, first recall that the Markowitz objective function is of the form

$f(\mu, \Sigma, w, \lambda) = w^T \Sigma w - \lambda \mu$, where λ parametrizes the agent's preference for return maximization over risk minimization. If we divide this objective function by λ , provided that $\lambda > 0$, we get an equivalent optimization problem for $\lambda^{-1} f$. Note that λ can be thought of as encoding 'collective' risk-preferences for all agents (e.g., a 'bull market' when λ is large, a 'bear market' when λ is small). As such we need a methodology for choosing λ that depends on the network size and/or number of agents. We make $\lambda \sim \chi^2(n)$ in order to capture the fact that as the number of participants increases, so should the expected number that are risk-seeking. Once we do this, we can directly interpret the variables α_i, β_i as encoding both this risk preference and the inherent duration difference between lending risk and staking risk. The static nature of the covariance matrix in [Assumption 11](#) implies that each agent has an 'equilibrium' risk-preference. Moreover, this implies that the distribution of risk-preferences amongst changes is not changing—which one can expect in the limit as the number of agents goes to infinity.⁶ Making this assumption allows for significantly faster simulation (e.g., less compute needed to reduce the variance of the estimator below an error threshold ϵ).

To summarize, we evolve the system by having each agent update their Markowitz estimate at each time step, which changes the lending and staking distributions for the next time step. In simulation, agents will only be able to migrate their staking tokens at times $k\tau_{\text{stake}}, k \in \mathbb{N}$. Explicitly, we evolve the system via the following loop:

- Initialize distributions $\{\pi_{\text{stake}}(0)_i\}_{i \in [n]}, \{\pi_{\text{lend}}(0)_i\}_{i \in [n]}$

- Initialize empirical distributions $\{\hat{\pi}_{\text{stake}}(0)_i\}_{i \in [n]}, \{\hat{\pi}_{\text{lend}}(0)_i\}_{i \in [n]}$
- For $t = 1, 2, 3, \dots$
 - Observe empirical distributions $\hat{\pi}_{\text{stake}}(t)_i, \hat{\pi}_{\text{lend}}(t)_i$
(generated via actual staking / lending rewards accrued in the epoch)
 - Compute Markowitz weights $w_{i,t} = (p, 1 - p), p = p(\hat{\pi}_{\text{stake}}(t)_i, \hat{\pi}_{\text{lend}}(t)_i, \alpha_i, \beta_i) \in [0, 1]$
 - Define new portfolio as $p_i(t) = p(p_i(t-1) + \ell_i(t-1)), \ell_i(t) = (1 - p)(p_i(t-1) + \ell_i(t-1))$

4. Formal Properties

We will describe a few formal properties (proved in the [Appendix](#)) of the two-state model, as it is feasible to analytically analyze this model. Our goal in this section is to bound the amount of turnover in the stake distribution that is caused by on-chain lending becoming more attractive than staking rewards to rational stakers. A secondary goal is to understand what properties of the stochastic processes that represent the evolution of the lending and staking distributions are necessary and/or sufficient for ensuring that we do not have a lot of volatility in agent token portfolios. This is important as volatility in these portfolios implies that PoS networks have volatile security, which is a distinct defect when compared to PoW. We will make a few additional assumptions that are necessary to provide analytical results:

Assumption 12 (Bounded Size)

There is a minimum fraction $\delta > 0$ of the money supply that needs to be staked, e.g., $S_t - \ell_t > \delta S_t$.

If no one is staking, then the on-chain lending contract has no value (as it doesn't have any security), so this is a realistic assumption that matches the practical parameters chosen in live networks such as Tezos [\[22\]](#) and Cosmos [\[21\]](#). Note that one can directly interpret δ as the fraction of altruistic validators who will never reallocate or rebalance their assets.

Assumption 13 (Number of Agents)

The number of agents n is larger than a constant multiple of the product of the exponential parameters, e.g., $n = \Omega(\tau_{\text{stake}}\tau_{\text{lend}})$.

This assumption is required for purely technical reasons that are explained in the proofs in the [Appendix](#).⁷ If we solve the *unconstrained* problem⁸ for the Markowitz objective function, $f(\mu, \Sigma, w, \lambda) = w^T \Sigma w - \lambda \mu$, then w needs to satisfy the first-order condition, $\nabla_w f(\mu, \Sigma, w, \lambda) = 0$. This yields,

$$w = \lambda \Sigma^{-1} \mu \tag{7}$$

Given that we are in the multi-period Markowitz setting, this means that we can estimate the i th participant's portfolio weights via $w_i(t) = \lambda \Sigma_i^{-1} \mu_i(t)$. As a measure of volatility in the security of the

underlying staking mechanism, we will first look at how $w_i(t)$ changes in time. Intuitively, this change corresponds to how large rebalancing events (e.g., moving tokens from staking to lending, or vice-versa) are between subsequent blocks. If this rebalancing is large, then the network could dramatically reduce its security as holders move their assets from staking to lending. On the other hand, if this rebalancing is small and decreases over time, then we know that the staked token supply is stable. Using the Markowitz update rule ([Equation 7](#)), we have the following bound:

$$\|w_i(t+1) - w_i(t)\|_1 = \|\Sigma^{-1}(\mu_i(t+1) - \mu_i(t))\|_1 \leq \|\Sigma^{-1}\|_{1 \rightarrow 1} \|\mu_i(t+1) - \mu_i(t)\|_1$$

where $\|A\|_{1 \rightarrow 1} = \max_{v \in \mathbb{R}^n, \|v\|_1=1} \|Av\|_1$ is the L^1 operator norm. This simple inequality implies that the volatility in portfolio weights, represented by the single block difference in weights, is controlled by the difference in the mean vector, as $\mathbb{E}[\|\Sigma^{-1}\|_{1 \rightarrow 1}] = \mathbb{E}[\alpha \vee \beta] \leq \tau_{\text{stake}} \vee \tau_{\text{lend}}$. Therefore, we focus on trying to bound the difference in expected returns as a function of the block reward at time t , R_t , the total token supply S_t , the lent supply ℓ_t , and the Compound lending curve parameters β_0, β_1 .

Let $\pi_{\text{stake}}(t+1)_i - \pi_{\text{stake}}(t)_i = \Delta_i^{\text{stake}}(t)$ and $\ell_{t+1} - \ell_t = \Delta_{\text{lend}}(t)$. Since the PoS algorithm is adapted to a filtration $\mathcal{F}_t$ on $S_t \Delta^n$ and the covariance matrices are constant as time evolves,² $\pi_{\text{stake}}(t), \pi_{\text{lend}}(t)$ are also $\mathcal{F}_t$ -adapted random variables. Note that by definition, $\Delta_i^{\text{stake}}(t) \leq S_t$ and $\Delta_{\text{lend}}(t) \leq S_t$, since we cannot change the amount staked or lent by more than the outstanding money supply. We will first bound the difference in expected returns as a function of $\Delta_i^{\text{stake}}(t), \Delta_{\text{lend}}(t)$, and S_t :

Claim 1

There exist constants $C, C' > 0$ such that

$$\|\mu_i(t+1) - \mu_i(t)\|_1 < \frac{C|\Delta_i^{\text{stake}}(t)|}{S_{t+1}} + C' \Delta_{\text{lend}}^2(t)$$

Note that these constants are allowed to depend on δ from [Assumption 12](#). If we have an compounding and inflationary rewards schedule, e.g., $S_t = e^{\lambda t}$, then this claim implies the following:

Claim 2

If there exists $\lambda > \tau_{\text{stake}}$ such that $S_t = \Omega(e^{\lambda t})$ and $\min_i W_i(t) > 0$, then as $t \rightarrow \infty$, the maximum change in stake is bounded above by the lending volatility, $\Delta_{\text{lend}}^2(t)$, with high probability.

This claim implies that if there is not much variance in the lending rate, either due to choosing small parameters β_0, β_1 or because borrowing demand is minimal, then we should not expect portfolios to rebalance regularly and rational stakers will tend to keep their tokens locked in a staking contract. Another natural quantity to look at is the variance of the lent assets. We show that the money supply and time-preference for lending, τ_{lend} , control the variance of lent assets.

Claim 3

Let $\mathcal{F}_t$ be the filtration such that the lending process ℓ_t is adapted. Then we have:

$$\text{Var}[\ell_{t+1} | \mathcal{F}_t] = \frac{\gamma_t}{\tau_{\text{lend}}^2} \|W(t)\|_2^2$$

Moreover, we have the following bounds:

$$\frac{\gamma_t^2 S_t^2}{\tau_{\text{lend}}^2 \sqrt{n}} \leq \text{Var}[\ell_{t+1} | \mathcal{F}_t] \leq \frac{\gamma_t^2 S_t^2}{\tau_{\text{lend}}^2}$$

Note that if k is the constant representing the ratio of borrowing demand to S_t (cf. [Section 3.2.1](#)), there exists a constant¹⁰ $\aleph$ that depends on k such that $\gamma_t \geq \aleph$. As such, [Claim 3](#) implies that as long as $S_t = \Omega(n^{\frac{1}{4}})$, we have reallocation from staking to lending. Thus, any monetary policy that grows sufficiently quickly with the number of users of the network will *always* have assets moving into and out of lending. If we place constraints on the demand k , we can strengthen this result into a bound on how much Δ_{lend}^2 oscillates:

Claim 4

Let $\eta_t = 1 + \frac{\|W(t)\|_2^2}{S_t^2} \geq 1 + \frac{1}{\sqrt{n}}$ and $\alpha_t = \frac{\ell_t \tau_{\text{lend}}}{S_t \eta_t}$. If for all t , $k \geq \frac{\alpha_t}{1-\alpha_t}$ where and the hypotheses of [Claim 2](#) hold, then $\Delta_{\text{lend}}^2(t)$ is a submartingale and we have

$$\Pr [\max_t \Delta_{\text{lend}}^2(t) > \lambda] < \frac{1}{\lambda^2} \text{Var}[\Delta_{\text{lend}}^2(t)]$$

and subsequently,

$$\Pr [\max_t \|\mu_i(t+1) - \mu_i(t)\|_1 > \lambda] < \frac{1}{\lambda^2} \text{Var}[\Delta_{\text{lend}}^2(t)] \quad (9)$$

In words, this claim says that as long as we have inflation and there is enough borrowing demand, then we can be sure that the *worst-case* rebalancing is bounded by the variance of lending volatility. If we add another constraint on the behavior of the increments $\Delta_{\text{lend}}^2(t)$, then we can strengthen this claim to get a phase transition that resembles the Galton-Watson phase transition [\[39\]](#).

Claim 5

Suppose that for all $t > 0$, $\Delta_{\text{lend}}(t) < \frac{\ell_t^2}{2\ell_{t-1}\eta_t}$, $\mathbb{E}[\Delta_{\text{lend}}(t)] > 0$, and let η_t be as in [Claim 4](#). Define $r_{\pm}$ as follows:

$$r_{\pm} = \frac{\ell_t \tau_{\text{lend}}}{S_t \eta_t} \left(1 \pm \sqrt{1 + \frac{\eta_t \ell_{t-1} (\ell_{t-1} - 2\ell_t)}{\ell_t^2}} \right)$$

Then ℓ_t^2 is a supermartingale when $\gamma_t \in (r_-, r_+)$, a submartingale when $\gamma_t \in [0, r_-) \cup (r_+, 1]$, and a martingale when $\gamma_t \in \{r_-, r_+\}$.

The intuition for this is as follows:

- When there is either too little or too much borrowing demand (e.g., $\gamma_t < r_-$ or $\gamma_t > r_+$), then the expected lent supply either increases (on average) to one or decreases to zero. This is analogous to a gambler's wealth after playing a game with probability $p < 1/2$ for n rounds—the wealth concentrates into either the house (staked supply) or the gambler (lent supply).
- When there is a moderate amount of borrowing demand, $\gamma_t \in [r_-, r_+]$, then we have stable, potentially oscillatory behavior. Doob's Supermartingale Convergence Theorem [40] intimates that the distribution is stationary as $t \rightarrow \infty$. This corresponds to a gambler playing a game in which their chance of winning is $p > 1/2$.

These results for the simpler two-level model suggest that our simulated phase transition results convey the existence of a deeper phase transition.

Agent-level Behavior

In order to get a stronger understanding of what is going on at the agent level and for different monetary policies, one needs an understanding of the probability that a single agent has a large rebalancing event that affects the staked portion of their portfolio. This involves studying how the staking components of $\mu_i(t)$ change over time. Let $\mu_{\text{stake}}(t)_i$ be the difference in the staking component of $\mu_i(t)$ between times t and $t + 1$. We can bound the probability that an agent rebalances their portfolio by an ϵ fraction via the following claim:

Claim 6

Let $\mu_{\text{stake}}(t)_i = \frac{\pi_{\text{stake}}(t+1)_i}{S_{t+1} - \ell_{t+1}} - \frac{\pi_{\text{stake}}(t)_i}{S_t - \ell_t}$. Then for all $t > 0$, we have

$$\Pr \left[|\mu_{\text{stake}}(t)_i| < \epsilon \mid \mathcal{F}_{t-1} \right] = \Omega \left(1 - \left(\left(\frac{\epsilon S_{t+1}}{\gamma_t} \right)^n e^{-\tau_{\text{stake}} \tau_{\text{end}} \frac{\epsilon S_{t+1}}{\gamma_t}} \right) \right) \quad (10)$$

Bounds like Claim 6 are of the form $\Pr[X_t < \epsilon] = Y_t$, which imply that the extrema of Y_t provide a guaranteed bound of how large X_t can be when Y_t is minimized. This means that we can try to bound the first hitting time t^* for the maxima of $|\mu_{\text{stake}}(t)_i|$ by analyzing the minima of the right-hand side of Equation 10. Note that the function $1 - (kx)^n e^{-k'kx}$ has a minima at $x = \frac{nk^{n-2}}{k'}$, which means that we can estimate when, as a function of S_t, γ_t, ϵ , we have maximal deviations in stake. In the following section, we examine this claim for different monetary policies S_t .

4.1 Deflationary Monetary Policy

In this case, $R_t = kr^{-t}$, $r < 1$ and $S_t = \frac{k(1-r^{-t-1})}{1-r} = C_\infty - C'r^{-t-1}$, where C_∞ is the final money supply (e.g., 21 million, for Bitcoin). Letting $\epsilon = \delta S_{t+1}$ and plugging this into the right hand side of Equation 10 and optimizing for t using the chain rule gives the condition

$t^* = -\log \left(C - \sqrt{\frac{n\delta}{\gamma_t \tau_{\text{stake}} \tau_{\text{lend}}}} \right)$. Requiring that $t^* > 0$ gives the condition $C - \sqrt{n\delta} \gamma_t \tau_{\text{stake}} \tau_{\text{lend}} \in (0, 1)$ or

$$\delta > \frac{(C_\infty - 1) \gamma_t \tau_{\text{stake}} \tau_{\text{lend}}}{n}$$

The above condition implies that at the time with the highest expected amount of removed stake, we expect that quantity to be on the order of $\frac{C_\infty}{n}$. If there are fewer terminal coins than participants, $C_\infty \ll n$, then we should expect large rebalances for every staker. This formalizes the intuition that totally deflationary currencies are vulnerable to large rebalances that depend strongly on the number of coins in the system. It also confirms the intuition that if there are fewer coins than participants, then borrowing demand will be high and we should again expect large rebalances.

4.2 Inflationary Monetary Policy

Next, we consider what [Claim 6](#) implies about inflationary monetary policies. We will consider two types of inflationary policies: Polynomial ($S_t = \Omega(t^k)$) and Exponential ($S_t = \Omega(e^{\lambda t})$).

4.2.1 Polynomial Inflation

For polynomial inflation, we have $R_t = ct^{k-1}$ and $S_t = \Omega(t^k)$. Optimizing [Equation 10](#) with this form of S_t gives $t^* = \left(\frac{n\delta}{\gamma_t \tau_{\text{stake}} \tau_{\text{lend}}} \right)^{1/k}$. Plugging this into the right-hand side of [Equation 10](#) gives a lower bound of

$$1 - \frac{\delta^2 n}{\gamma_t \tau_{\text{stake}} \tau_{\text{lend}}} e^{-\tau_{\text{stake}} \tau_{\text{lend}} \delta^2 n / \gamma_t}$$

which means that if $\delta = O(n^{-1/2})$, individual agents will not have rebalances of size δS_t with high probability. These guarantees do not depend on the polynomial degree k , which means that even simple linear policies ($k = 1$) are sufficient to get low stake turnover.

4.2.2 Exponential Inflation

In the case of exponential inflation, $S_t = C_0 e^{\lambda t}$ for some $\lambda > 0$ and initial token distribution C_0 . Following the logic of [Section 4.1](#), we arrive at a conclusion of,

$$\delta > \frac{C_0 \gamma_t \tau_{\text{stake}} \tau_{\text{lend}}}{n}$$

This states that the worst-case rebalancing is bounded by the *initial token distribution* as opposed to the final token distribution. Given that most networks assume that there will be more users than the *initial* distribution, C_0 , this implies that large rebalancing events should become rarer once the network achieves a scale of $n \gg C_0$.

4.3 Conclusions

The results of this section show that the model of [Section 3.3](#) has a volatility that is mainly dependent on lending volatility. This volatility, if demand is sufficiently shallow, is small enough to help bound the worst-case portfolio rebalancing for any validator in the system. However, the turnover in staked quantity, measured by how each agent’s expected reward changes over each epoch period, is sensitive to the precise nature of the monetary policy, S_t . In particular, deflationary policies cannot support on-chain lending, as the worst-case rebalancing rate depends on the terminal money supply, whereas the rebalancing rate for exponentially inflationary monetary policies only depends on the initial money supply (e.g., Ethereum’s pre-mine). Finally, we note that polynomial inflation provides particularly good rebalancing guarantees that are independent of the rate of growth of the money supply. Many of these results rely on asymptotic behavior of agents in this system and all of these results depend on simple models for borrowing demand. The remainder of the paper will focus on using simulation to provide a more realistic picture of [Section 3.3](#).

5. Agent-Based Simulations

In order to test the three-state model in a quantitatively rigorous and realistic fashion, we turn to agent-based simulations. One of the main reasons to use agent-based simulation over formal methods is that it is hard to formally prove what block reward growth rate is ideal for mitigating volatility in portfolio allocation. For instance, it is difficult to evaluate whether $R_t = \Omega(t^2)$ or $R_t = \Omega(t)$ provide ideal mitigation for on-chain lending with parameters β_0, β_1 . Moreover, we can test various realistic demand distributions, including ones that are atomic and do not have a probability density. Prior work on agent-based simulations of blockchain systems [\[19\]](#) has focused on analyzing consensus protocols via event-based simulation. We follow a similar event-based framework for simulating staking and lending, albeit ignoring the details of peer-to-peer networking and consensus. Our goal is to sample as many trajectories $X_t | \mathcal{F}_t = (\pi_{\text{stake}}(t), \pi_{\text{lend}}(t), \gamma_t)$ as possible for different combinations of parameters β_0, β_1 , and block reward schedules R_t .

5.1 Initialization

In order to generate sample paths X_t , we need to specify the following variables:

- Initial distributions: $\pi_{\text{stake}}(0), \pi_{\text{lend}}(0)$
- Initial interest rate: $\gamma_0 \in (0, 1)$
- Interest rate parameters: $\beta_0, \beta_1 \in (0, 1)$
- Demand-generating distribution parameters (vary depending on the demand generating distributions chosen): $\mathcal{P}_\xi$
- Staking and lending time scale: $\tau_{\text{stake}}, \tau_{\text{lend}}$
- Slashing probability: p_{slash}

- PRNG seed: $s \in \mathbb{Z}_{2^{64}}$

We chose to model $\pi_{\text{stake}}(0)_i \sim \text{Exp}(\lambda_{\text{stake}})$ and $\pi_{\text{lend}}(0)_i \sim \text{Exp}(\lambda_{\text{lend}})$, as these distributions exemplify the extreme concentration of wealth that accompanies most token distributions.

Exponential distributions are also useful in that the order statistics are also exponential,¹¹ which represents the idea that the k th entrant to a network should have a decaying fraction (in this case $\frac{1}{k}$) of the total token supply. We note that we do not use power law distributions as there are conflicting reports of the wealth distribution in Bitcoin actually following a power law [41][42]. Given the statistical difficulty of discerning if one has a power law versus an exponential decay [43][44] and the extra parameter in a power law (e.g., $p(x = t) \propto x_{\min} t^{-a}$), we decided to use exponential initial distributions.

In our simulations, we kept the initial interest rate at 10%, as we saw very little sensitivity to the initial choice. In particular, agents quickly rebalance their portfolio into lending, if the rate is high enough, and this appeared to equilibrate within a small number of time steps. We swept the other parameters, $\beta_0, \beta_1, \tau_{\text{stake}}, \tau_{\text{lend}}, p_{\text{slash}}$, over realistic ranges and for each choice of parameter, we sampled trajectories for a variety of seeds s to get an ensemble average.

5.2 Simulation Loop

The main simulation loop follows that as described at the end of [Section 3.3](#). We exactly evaluate the Markowitz update rule, with constraints, via the exact solution for optimal portfolios [37]. Given that we are solving n independent, two-dimensional Markowitz problems, we evaluated the constraints analytically (as opposed to using a convex solver like Gurobi or CVX). The main event loop has the following causal ordering:

1. If $t = k\tau_{\text{stake}}$ for some $k \in \mathbb{N}$, allow agents to update their Markowitz portfolios¹²
2. Sample the borrowing demand distribution
3. Update γ_t as a function of the new borrowing demand
4. Run [Algorithm 1](#) to determine who wins the reward R_t and/or if they get slashed

5.3 Borrowing Demand Distributions

We sampled a variety of borrowing demand distributions for deflationary and inflationary block rewards, as illustrated via the sample paths in [Figure 4](#). These stochastic borrowing demand paths each have four parameters: mean, variance, maximum, and minimum demand. We define the maximum and minimum demand parameters as multipliers of the token supply, so that the minimum is $(1 - \eta_0)S_t$ and the maximum is $(1 + \eta_1)S_t$ (illustrated via the lower and upper bounds in [Figure 4](#), respectively).

Figure 4: Borrowing demand sample paths, ξ_t , juxtaposed with the money supply S_t . The sample paths in the upper left (linear, $R_t = c$), upper right (deflationary, $R_t = kr^t$, $r < 1$), and lower left (inflationary, $R_t = kr^t$, $r > 1$) use simple random walks, whereas the figure in the lower right samples a geometric Brownian motion path with reflecting boundary conditions at $(1-\eta_0)S_t$, $(1+\eta_1)S_t$.

5.4 Results

In order to evaluate a variety of realistic conditions, we swept through many parameters illustrated above. The most interesting results come from looking at individual trajectories, such as [Figure 1](#), and heatmaps of how certain scalar functions behave as we vary the lending rate parameters β_0, β_1 . We generated heatmaps for a number of random seeds and took averages over these random instantiations to generate [Figure 6](#) and [Figure 7](#). The two main measurements that we looked at were

$$f(\beta_0, \beta_1) = \mathbb{E}_t \left[\frac{S_t - \ell_t}{S_t + \ell_t} \middle| \beta_0, \beta_1, \tau \right] \text{ and } g(\beta_0, \beta_1) = \mathbb{E}_t[\gamma_t - \beta_t | \beta_0, \beta_1, \tau],$$

where given a non-anticipating stochastic process X_t and a function stopped at time τ , $\mathbb{E}_t[f(X_t) | \tau] = \int_0^\tau f(X_t) dX_t$, for the stochastic measure dX_t . We ran all trajectories for $\tau = 250\,000$ blocks and approximate

$$\mathbb{E}_t[f(X_t) | \tau] \approx \sum_{i=0}^{\tau/\Delta} f(X_i) \xi_i, \text{ where } \xi_i \text{ is sampled from } dX_{t+\Delta} - dX_t.$$

The first quantity, $f(\beta_0, \beta_1)$, is the normalized difference between the staked supply and the lent supply and when it is greater

than zero, there is more quantity staked, on average. It is normalized relative to the total money

supply $S_t + \ell_t$ so that $\forall t, \frac{S_t - \ell_t}{S_t + \ell_t} \in (0, 1)$ and thus we can compare the relative staking to lending

proportions at different times. The second quantity, $g(\beta_0, \beta_1)$, measures the linear spread between

borrowing and lending rates, which varies over time even though $1 - \gamma_0 = \frac{\beta_t - \gamma_t}{\beta_t}$ is constant. If the market for borrowing is efficient and there is less churn, we should expect that this rate should be quite low. Note that in all simulations used, we set the relative spread, $1 - \gamma_0$, to be 0.5%.

In [Figure 5](#), we see plots of $f(\beta_0, \beta_1)$ for different inflation rates $r \in \{0.25, 0.5, 1.0\}$. Note that when $r = 1.0$, then we have a linear polynomial inflation rate. This figure demonstrates that even though the dependence on β_0, β_1 appears to be random, it is clear that the deflationary figures tend to have significantly more lending than staking. On the other hand, the linearly inflating component enjoys a significant advantage with regards to staked supply. Note that the dependence on β_0, β_1 appears mostly random because of the use of a common scale to plot them. The usage of a common scale dampens the correct scale to plot the figures at (which should be on the order of $\text{Var}(f(\beta_0, \beta_1))$), which varies dramatically as a function of the other parameters that we sample, such as the block reward. In some of the other plots that we will look at, we use independent color scales for each plot to emphasize the variation as a function of β_0 and β_1 .

Figure 5: Heat maps of $f(\beta_0, \beta_1) = \mathbb{E}_t[\frac{S_t - \ell_t}{S_t + \ell_t} | \beta_0, \beta_1]$ for different deflation schedules at a half-life equal to 0.25% of the total simulation time (400 Bitcoin-style halvings occur). Each pixel is the average of 100 random initializations with the other parameters fixed. On the left, we have a very high deflationary rate and we can see that the values are negative, implying that $\ell_t > S_t$, whereas on the far right picture, which is non-deflationary, we see that $S_t > \ell_t$.

On the other hand, in [Figure 6](#), we see an array of plots that show how the borrowing spread, $g(\beta_0, \beta_1)$, changes as a function of inflation rate and borrowing threshold. The third column, which is the linearly increasing block rewards regime, demonstrates very tight spreads, suggesting that even changing the borrowing demand threshold causes little variation in spreads. On the other hand, the deflationary regime is much more sensitive to shocks in borrowing demand, as illustrated by the plots in the upper left hand corner of [Figure 6](#). These empirical results validate [Claim 2](#) and [Claim 3](#), as we directly see that lending volatility has a much more muted effect on inflationary systems.

Figure 6: Plots of $g(\beta_0, \beta_1) = \mathbb{E}[\gamma_t - \beta_t | \beta_0, \beta_1]$. As you go from left to right, we increase the inflation rate r (recall that when $r < 1$, we are deflationary) and as we go from top to bottom, the borrow threshold increased. Recall that the borrow threshold is the lower bound line from [Figure 4](#). We see that when the system has linearly increasing block rewards, we receive very tight spreads and a smooth, almost linear $g(\beta_0, \beta_1)$. On the other hand, as we have increasing deflationary policies, we see that we cause the expected spread to become more random and larger.

Finally, in [Figure 7](#), we see heatmaps of $f(\beta_0, \beta_1)$ for different borrowing thresholds and inflation rates. Note that these plots have different scales, unlike [Figure 5](#), and that the deflationary figures are

all mainly negative. From the first two columns, it becomes clear that as borrowing demand increases, $g(\beta_0, \beta_1)$ becomes increasingly monotonically decreasing in β_0, β_1 . This means that as borrowing demand increases, the system's loss of security to lending increases in a more predictable manner. On the other hand, the linearly inflating regime does not have this issue and continues to have relatively random and non-monotonic dependence on β_0, β_1 as we adjust the borrowing threshold.

Figure 7: Plots of $f(\beta_0, \beta_1)$ for differing inflation rates and borrowing thresholds. Note that the color bars vary for each figure. We see that when the policy is deflationary, $f(\beta_0, \beta_1) < 0$. Moreover, when the borrowing threshold is high, we cause the plots to become less random and have a more clear dependence on β_0, β_1 .

6. Conclusions and Future Work

This paper has explored how on-chain lending affects network security in PoS networks by studying how rational multiple period Markowitz optimizing agents optimize their token portfolios. In

particular, we find that bank runs can occur when many agents collectively move their tokens from staking contracts to lending contracts even when agents have independently drawn risk preferences. These attacks, which are coordinated only by rational optimization, show that the strictly Byzantine threat model is insufficient to describe security in PoS networks. Our theoretical and empirical findings show that deflationary PoS tokens are susceptible to attack by non-Byzantine, but rational, participants who constantly turnover and rebalance their staking and lending portfolios to chase yield. Moreover, these findings show that one needs to choose a block reward schedule that increases relative to the yield that these on-chain contracts provide. Our theoretical results give explicit probabilities for the constant relative borrowing demand regime and show that there is a clear and strong dependence on the lending rate γ_t in these probabilities. We buffer these results with agent-based simulations to confirm that these properties hold in more realistic borrowing scenarios. In the future, we aim to look at how unbonding times and transaction fees (*e.g.*, gas costs) affect these protocols and how different lending curves will affect staking security. We also aim to provide more precise theoretical results by relaxing a number of the assumptions that are utilized in this paper.

Acknowledgments

The author would like to thank Yi Sun, Rei Chiang, Hasu, Haseeb Qureshi, Ivan Bogatyy, Georgios Konstantopoulos, John Morrow, Tony Salvatore, Wei Wang, Hsien-Tang Kao, Jonathan Reem, Josh Chen, Tina Zhen, Nader Al-Naji, Michael Jordan, Tim Roughgarden, and Matteo Leibowitz for comments and feedback. The author would also like to thank Barnabé Monnot and Arian Klags-Mundt for correcting an earlier error in Claim 1 and for helpful expository comments. Finally, the author would like to acknowledge the helpful input from the reviewers of both the Stanford Blockchain Conference and the MIT [Cryptoeconomics Systems](#) journal.

Appendix

A. Notation

- Δ^n is n -dimensional probability simplex, $\Delta^n = \{(x_1, \dots, x_n) \in \mathbb{R}^n : \sum_{i=1}^n x_i = 1, \forall i, x_i \geq 0\}$
- For any $x \in \mathbb{R}^n$, we define the p -norm as $\|x\|_p = (\sum_{i=1}^n |x_i|^p)^{1/p}$.
- We turn any nonzero vector $x \in \mathbb{R}^n$ with $x \geq 0$ into a probability distribution by defining $\hat{x} = \frac{x}{\|x\|_1} \in \Delta^n$.
- $S_+^n \subset \mathbb{R}^{n \times n}$ is the cone of positive definite, symmetric matrices
- $\vee, \wedge$ are the standard join and meet of two elements of a lattice. For example if $a, b \in \mathbb{R}$, $a \vee b = \max(a, b), a \wedge b = \min(a, b)$
- We use standard Landau notation [45] on totally ordered sets D : Given functions $f : D \rightarrow \mathbb{R}, g : D \rightarrow \mathbb{R}$, we use the following asymptotic notations:
 - $f \in O(g) \iff \exists C > 0, \forall d \in D, f(d) \leq Cg(d)$

- $f \in \Omega(g) \iff \exists c > 0, \forall d \in D, f(d) \geq cg(d)$
- $f \in o(g) \iff \lim_{d \rightarrow \sup D} \frac{f(d)}{g(d)} = 0$
- $f \in \Theta(g) \iff f \in O(g) \text{ and } f \in \Omega(g)$

B. PoS Algorithm

Algorithm 1 Epoch-based PoS Simulation

```

 $\pi_{\text{stake}} \leftarrow \pi_{\text{stake}}(0) \in \Delta^n$  Initial token distribution
 $R \rightarrow \emptyset$  Initialize current epoch's reward set
 $S \rightarrow \emptyset$  Initialize current epoch's slash set
 $t \leftarrow 0$ 
for  $epoch \in \mathbb{N}$  do
  # Update Staking Distribution
  for  $id \in \{1, \dots, n_{\text{stakers}}\}$  do
    if  $(id, \text{reward}) \in R$  then
       $\pi_{\text{stake}}[id] + = \text{reward}$ 
    end if
    if  $(id, \text{slash}) \in S$  then
       $\pi_{\text{stake}}[id] - = \text{slash}$ 
    end if
   $R \leftarrow \emptyset, S \leftarrow \emptyset$ 

  # Draw next epoch from stake distribution

   $\hat{\pi}_{\text{stake}} \leftarrow \frac{\pi_{\text{stake}}}{\|\pi_{\text{stake}}\|_1}$  Create normalized probability distribution
  for  $block \in \{1, \dots, \tau_{\text{stake}}\}$  do
     $id \sim \hat{\pi}_{\text{stake}}$  # Draw a winner from the stake distribution
     $X \sim \text{Bern}(p_{\text{slash}})$ 
    if  $X == 1$  then
      if  $\exists (id, s) \in S$  then
         $S \leftarrow S \cup \{(id, s - S(id, \pi_{\text{stake}}))\}$ 
         $S \leftarrow S - \{(id, s)\}$ 
      else
        if  $\exists (id, r) \in R$  then
           $R \leftarrow R \cup \{(id, r + R(t))\}$ 
           $R \leftarrow R - \{(id, r)\}$ 
        end if
      end if
    end if
     $t += 1$ 
  end for
end for
end for

```

Algorithm 1

C. Proof of Claim 1

In [Assumption 12](#), we assert that for all t , $S_t - \ell_t > \delta S_t$ or $1 - \frac{\ell_t}{S_t} > \delta$ and $\frac{1}{1 - \frac{\ell_t}{S_t}} < \frac{1}{\delta}$. Since the function $\frac{1}{1-x}$ is convex on $[0, 1]$, we have, for any $x \in [0, 1 - \delta]$, $s \in [0, 1]$, $\frac{1}{1-x} < s \left(\frac{1}{\delta}\right) + (1-s) = s \left(\frac{1-\delta}{\delta}\right) + 1$. Using this bound on the function $\frac{1}{1 - \frac{\ell_t}{S_t}}$, we have, for any $s \in [0, 1]$:

$$\frac{1}{S_t - \ell_t} = \frac{1}{S_t} \frac{1}{1 - \frac{\ell_t}{S_t}} < \frac{1}{S_t} \left(s \left(\frac{1-\delta}{\delta} \right) + 1 \right)$$

In particular, since $\ell_t < S_t$, we can use this bound at $s = \frac{\ell_t}{S_t}$ to get

$$\frac{1}{S_t - \ell_t} < \frac{1}{S_t} \left(\frac{\ell_t}{S_t} \left(\frac{1-\delta}{\delta} \right) + 1 \right) \quad (11)$$

This implies that

$$\mu_{\text{stake}}(t)_i = \frac{\pi_{\text{stake}}(t+1)_i}{S_t - \ell_t} < \frac{\pi_{\text{stake}}(t+1)_i}{S_t} \left(\frac{\ell_t}{S_t} \left(\frac{1-\delta}{\delta} \right) + 1 \right)$$

Note that the condition $S_t - \ell_t > \delta S_t$ also implies that $\frac{\ell_t}{S_t} < 1 - \delta$, so that we can write this as:

$$\mu_{\text{stake}}(t)_i = \frac{\pi_{\text{stake}}(t+1)_i}{S_t - \ell_t} < \frac{\pi_{\text{stake}}(t+1)_i}{S_t} \left(\frac{(1-\delta)^2}{\delta} + 1 \right)$$

On the other hand, since $S_t > S_t - \ell_t$, we also have $\mu_{\text{stake}}(t)_i > \frac{\pi_{\text{stake}}(t+1)_i}{S_t}$. Thus we have

$$|\mu_{\text{stake}}(t+1)_i - \mu_{\text{stake}}(t)_i| = \Theta \left(\frac{\pi_{\text{stake}}(t+1)_i}{S_{t+1}} - \frac{\pi_{\text{stake}}(t)_i}{S_t} \right)$$

Let $f(\delta)$ be a constant such that

$$|\mu_{\text{stake}}(t+1)_i - \mu_{\text{stake}}(t)_i| \leq f(\delta) \left(\frac{\pi_{\text{stake}}(t+1)_i}{S_{t+1}} - \frac{\pi_{\text{stake}}(t)_i}{S_t} \right)$$

which exists as per the last statement. Given these equations, we can now bound the single time step rebalance as follows:

$$\begin{aligned}
\|\mu_i(t+1) - \mu_i(t)\|_1 &= \left| \frac{\pi_{\text{stake}}(t+1)_i}{S_t - \ell_t} - \frac{\pi_{\text{stake}}(t)_i}{S_t - \ell_t} \right| + |\gamma_{t+1} - \gamma_t| \\
&\leq f(\delta) \left| \frac{\pi_{\text{stake}}(t+1)_i}{S_{t+1}} - \frac{\pi_{\text{stake}}(t)_i}{S_t} \right| + \beta_1 \gamma_0 \left| \left(\frac{\ell_{t+1}}{S_{t+1}} \right) \right| \\
&= \left| \frac{\pi_{\text{stake}}(t+1)_i}{S_t + R_t} - \frac{\pi_{\text{stake}}(t)_i}{S_t} \right| + \beta_1 \gamma_0 \left| \left(\frac{\ell_{t+1}}{S_t + R_t} \right) \right| \\
&\leq f(\delta) \left| \frac{\pi_{\text{stake}}(t)_i + \Delta_i^{\text{stake}}(t)}{S_t + R_t} - \frac{\pi_{\text{stake}}(t)_i}{S_t} \right| + \beta_1 \gamma_0 \\
&\leq \pi_{\text{stake}}(t)_i \left| \frac{f(\delta)}{S_t + R_t} - \frac{f(\delta)}{S_t} \right| + \frac{f(\delta) |\Delta_i^{\text{stake}}(t)|}{S_t + R_t} + \\
&\leq \pi_{\text{stake}}(t)_i \frac{k f(\delta) R_t}{(S_t + R_t)^2} + \frac{f(\delta) |\Delta_i^{\text{stake}}(t)|}{S_t + R_t} \\
&+ \beta_1 \gamma_0 \ell_t \left(\frac{\ell_t (2S_t R_t + R_t^2)}{S_t^2 (S_t + R_t)^2} + \frac{1}{(S_t + R_t)^2} \right) \left((2|\Delta_{\text{lend}}|) \right) \\
&\leq \pi_{\text{stake}}(t)_i \frac{f(\delta) R_t}{(S_t + R_t)^2} + \frac{f(\delta) |\Delta_i^{\text{stake}}(t)|}{S_t + R_t} + \beta_1 \gamma_0 \ell_t \\
&\leq \frac{f(\delta)}{S_t + R_t} \left(|\Delta_i^{\text{stake}}(t)| + \frac{2\beta_1 \gamma_0 \ell_t + R_t}{S_t} \right) + \beta_1 \gamma_0 \ell_t
\end{aligned}$$

D. Proof of Claim 2

Since $R_t = o(2^n)$, $S_t = \sum_{i < t} R_i$, we reach the desired conclusion if for any $\beta \in (0, 1]$, $\Pr[|\Delta_i^{\text{stake}}(t)| > S_t^{\beta}] = O(2^{-S_t^{\beta}})$. Let $W_i(t) = \pi_{\text{stake}}(t)_i + \pi_{\text{lend}}(t)_i$ be the wealth of agent i . Then the unconstrained Markowitz update gives,

$$\pi_{\text{stake}}(t+1)_i = \frac{\alpha_i \pi_{\text{stake}}(t)_i}{S_t - \ell_t} W_i(t)$$

where α_i is the inverse covariance sample for agent i . Thus,

$$|\Delta_i^{\text{stake}}(t)| = |\pi_{\text{stake}}(t+1)_i - \pi_{\text{stake}}(t)_i| = \pi_{\text{stake}}(t)_i \left| \frac{\alpha_i W_i(t)}{S_t - \ell_t} - 1 \right|$$

Thus, we have:

$$\Pr \left[|\Delta_i^{\text{stake}}(t)| > S_t^\beta \right] = \Pr \left[\pi_{\text{stake}}(t)_i \left| \frac{\alpha_i W_i(t)}{S_t - \ell_t} - 1 \right| > S_t^\beta \right] \quad (13)$$

$$= \Pr \left[\alpha_i > \left(\frac{S_t - \ell_t}{W_i(t)} \right) (S_t^\beta + \pi_{\text{stake}}(t)_i) \right] < 2^{-\tau_{\text{stake}} S_t^\beta} \quad (14)$$

where the last step holds since α_i is an exponential random variable with parameter τ_{stake} .

E. Proof of Claim 3

Recall that β_i was defined in [Equation 4](#) and is different than the lending rate β_t . Note the following, which uses $\beta_i \sim_{\text{iid}} \text{Exp}(\tau_{\text{lend}})$:

$$\mathbb{E}[\ell_{t+1}^2 | \mathcal{F}_t] = \mathbb{E} \left[\left(\gamma_t \sum_{i=1}^n \beta_i W_i(t) \right)^2 \middle| \mathcal{F}_t \right] \quad (15)$$

$$= \gamma_t^2 \mathbb{E} \left[\sum_{i=1}^n \beta_i^2 W_i(t)^2 + \sum_{i \neq j} \beta_i \beta_j W_i(t) W_j(t) \right] \quad (16)$$

$$= \gamma_t^2 \left(\frac{2}{\tau_{\text{lend}}^2} \sum_{i=1}^n W_i(t)^2 + \frac{1}{\tau_{\text{lend}}^2} \sum_{i \neq j} W_i(t) W_j(t) \right) \quad (17)$$

$$= \gamma_t^2 \left(\frac{1}{\tau_{\text{lend}}^2} \left(\sum_{i=1}^n W_i(t) \right)^2 + \frac{1}{\tau_{\text{lend}}^2} \sum_{i=1}^n W_i(t)^2 \right) \quad (18)$$

$$= \gamma_t^2 \left(\frac{S_t^2}{\tau_{\text{lend}}^2} + \frac{1}{\tau_{\text{lend}}^2} \sum_{i=1}^n W_i(t)^2 \right) \quad (19)$$

$$= \gamma_t^2 \left(\frac{S_t^2}{\tau_{\text{lend}}^2} + \frac{\|W(t)\|_2^2}{\tau_{\text{lend}}^2} \right) \quad (20)$$

Furthermore, we have the following:

$$\mathbb{E}[\ell_{t+1} | \mathcal{F}_t] = \mathbb{E} \left[\sum_{i=1}^n \gamma_t \beta_i W_i(t) \middle| \mathcal{F}_t \right] = \frac{\gamma_t S_t}{\tau_{\text{lend}}} \quad (21)$$

Combining [Equations 20](#) and [Equation 21](#) yields:

$$\text{Var}[\ell_{t+1} | \mathcal{F}_t] = \mathbb{E}[\ell_{t+1}^2 | \mathcal{F}_t] - (\mathbb{E}[\ell_{t+1} | \mathcal{F}_t])^2 = \frac{\gamma_t^2 \|W(t)\|_2^2}{\tau_{\text{lend}}^2} \quad (22)$$

as claimed. As the Cauchy-Schwarz inequality implies that $\|x\|_2^2 \leq \|x\|_1^2 \leq \sqrt{n} \|x\|_2^2$ and

$S_t = \|W(t)\|_1$, we have:

$$1 \leq \frac{S_t^2}{\|W(t)\|_2^2} = \frac{\|W(t)\|_1^2}{\|W(t)\|_2^2} \leq \sqrt{n} \quad (23)$$

Thus we have the final part of the claim:

$$\text{Var}[\ell_{t+1} | \mathcal{F}_t] = \frac{\gamma_t \|W(t)\|_2^2}{\tau_{\text{lend}}^2} = \frac{\gamma_t S_t^2}{\tau_{\text{lend}}^2} \left(\frac{\|W(t)\|_2^2}{S_t^2} \right) \in \frac{\gamma_t S_t^2}{\tau_{\text{lend}}^2} \left[\frac{1}{\sqrt{n}}, 1 \right] \quad (24)$$

F. Proof of Claim 4

Noting that $\gamma_t = \frac{k S_t}{\ell_t + k S_t} = \frac{1}{1 + \frac{\ell_t}{k S_t}}$ implies that for the choice of k in the claim, [Equation 20](#) is greater than ℓ_t^2 and thus ℓ_t^2 is a non-negative submartingale with bounded differences (e.g., $\Delta_{\text{lend}}^2(t) \leq S_t^2$). In particular, note that $\gamma_t \geq \beta_1 \left(\frac{k}{k+1} \right)^2 + \beta_0 \left(\frac{k}{k+1} \right) = \Omega \left(\frac{k}{k+1} \right)$. Thus via [Equation 20](#), we have:

$$\mathbb{E}[\ell_t^2 | \mathcal{F}_t] = \frac{\gamma_t^2 S_t^2}{\tau_{\text{lend}}^2} \left(1 + \frac{\|W(t)\|_2^2}{S_t^2} \right) = \frac{\gamma_t^2 S_t^2 \eta_t}{\tau_{\text{lend}}^2} \geq C' \left(\frac{k}{k+1} \right)^2 \frac{S_t^2 \eta_t}{\tau_{\text{lend}}^2} \quad (25)$$

The choice $k = \Omega \left(\frac{\alpha_t}{1 - \alpha_t} \right)$ ensures that the right-hand side of [Equation 25](#) is greater than ℓ_t^2 . The norm equivalence between a non-negative submartingale and its quadratic variation (e.g., Burkholder's inequality [\[46\]](#)) applies to ℓ_t^2 and implies that $\Delta_{\text{lend}}^2(t)$ is also a submartingale. Note that this relies on the being able to pass bounds on the quadratic variation process to the individual increments, which only applies to the discrete setting [\[47\]](#). The tail bound [Equation 9](#) comes from Doob's submartingale inequality and Kolmogorov's inequality for submartingales [\[40\]](#).

G. Proof of Claim 5

Recall that $\Delta_{\text{lend}}^2(t)$ is a martingale if $\mathbb{E}[\Delta_{\text{lend}}^2(t)|\mathcal{F}_t] = \Delta_{\text{lend}}^2(t)$, a supermartingale if $\mathbb{E}[\Delta_{\text{lend}}^2(t)|\mathcal{F}_t] \leq \Delta_{\text{lend}}^2(t)$, and a submartingale if $\mathbb{E}[\Delta_{\text{lend}}^2(t)|\mathcal{F}_t] \geq \Delta_{\text{lend}}^2(t)$. The expansion of $\mathbb{E}[\Delta_{\text{lend}}^2(t+1)|\mathcal{F}_t]$ is:

$$\begin{aligned} \mathbb{E}[\Delta_{\text{lend}}^2(t)|\mathcal{F}_t] &= \mathbb{E}[(\ell_{t+1} - \ell_t)^2|\mathcal{F}_t] \\ &= \mathbb{E}[\ell_{t+1}^2|\mathcal{F}_t] - 2\ell_t \mathbb{E}[\ell_{t+1}|\mathcal{F}_t] + \ell_t^2 \\ &= \frac{\gamma_t^2 S_t^2 \eta_t}{\tau_{\text{lend}}^2} - 2\ell_t \frac{\gamma_t S_t}{\tau_{\text{lend}}} + \ell_t^2 \end{aligned} \quad (26)$$

Define the quadratic function $f : \mathbb{R} \rightarrow \mathbb{R}$ as:

$$\begin{aligned} f(\gamma_t) &= \gamma_t^2 \left(\frac{S_t^2 \eta_t}{\tau_{\text{lend}}^2} \right) - \gamma_t \left(\frac{2\ell_t S_t}{\tau_{\text{lend}}} \right) + (\ell_t^2 - (\ell_t^2 - 2\ell_t \ell_{t-1} + \ell_{t-1}^2)) \\ &= \gamma_t^2 \left(\frac{S_t^2 \eta_t}{\tau_{\text{lend}}^2} \right) - \gamma_t \left(\frac{2\ell_t S_t}{\tau_{\text{lend}}} \right) + (2\ell_t \ell_{t-1} - \ell_{t-1}^2) \end{aligned} \quad (27)$$

Note that $f(\gamma_t) = \mathbb{E}[\Delta_{\text{lend}}^2(t)|\mathcal{F}_t] - \Delta_{\text{lend}}^2(t-1)$ so that the roots of f correspond to when Δ_{lend}^2 is a martingale. Moreover, when $f > 0$, $\Delta_{\text{lend}}^2(t)$ is a submartingale, and a supermartingale otherwise. Thus, we only need to find the roots of f and to figure out when $f > 0$ and $f < 0$ to classify this transition. The roots of $f, r_{\pm}$ are:

$$r_{\pm} = \frac{\ell_t \tau_{\text{lend}}}{S_t \eta_t} \left(1 \pm \sqrt{1 + \frac{\eta_t \ell_{t-1} (\ell_{t-1} - 2\ell_t)}{\ell_t^2}} \right).$$

Thus, the roots are real if and only if we have:

$$\begin{aligned} \eta_t (2\ell_t \ell_{t-1} - \ell_{t-1}^2) &< \ell_t^2 \\ \iff \eta_t \ell_{t-1} (\Delta_{\text{lend}}(t-1) + \ell_t) &< \ell_t^2 \end{aligned} \quad (28)$$

$$\iff \Delta_{\text{lend}}(t-1) < \frac{\ell_t^2}{\eta_t \ell_{t-1}} - \ell_t \quad (29)$$

$$\Rightarrow \Delta_{\text{lend}}(t-1) < \ell_t \left(\frac{\ell_t}{2\ell_{t-1}} - 1 \right) \quad (30)$$

Recall that for a quadratic function $f(x) = ax^2 + bx + c$ with real roots $r_{\pm}$, we have

$f^{-1}((-\infty, 0)) = [r_-, r_+]$ if $f(\frac{r_+ + r_-}{2}) = f(-\frac{b}{4a}) < 0$, otherwise

$f^{-1}((-\infty, 0)) = (-\infty, r_-) \cup (r_+, \infty)$. For [Equation 27](#), we have $\frac{-b}{4a} = \frac{\ell_t \tau_{\text{lend}}}{2S_t \eta_t}$, which gives:

$$\begin{aligned} f\left(\frac{\ell_t \tau_{\text{lend}}}{2S_t \eta_t}\right) &= \left(\frac{\ell_t^2 \tau_{\text{lend}}^2}{4S_t^2 \eta_t^2}\right) \left(\frac{S_t^2 \eta_t}{\tau_{\text{lend}}^2}\right) - \left(\frac{\ell_t \tau_{\text{lend}}}{2S_t \eta_t}\right) \left(\frac{2\ell_t S_t}{\tau_{\text{lend}}}\right) + \ell_{t-1}(\Delta_{\text{lend}}) \\ &= -\frac{3\ell_t^2}{4\eta_t} + \ell_{t-1}(\Delta_{\text{lend}}(t-1) + \ell_t) \\ &= -\frac{3\ell_t^2}{4\eta_t} + \ell_{t-1}(\Delta_{\text{lend}}(t-1) + \ell_t) \end{aligned} \quad (31)$$

Note that [Equation 31](#) is less than zero when

$$\Delta_{\text{lend}}(t-1) < \ell_t \left(\frac{3\ell_t}{4\ell_{t-1}} - 1 \right)$$

which is already subsumed by the condition for having real roots. Thus, provided that $r_{\pm} \in (0, 1)$, we have the desired conclusion.

H. Proof of Claim 6

As before, let the i th agent's token wealth be denoted $W_i(t) = \pi_{\text{stake}}(t)_i + \pi_{\text{lend}}(t)_i$. Using the Markowitz update rule, we have the following:

$$\begin{aligned} \mu_{\text{stake}}(t)_i &= \frac{\pi_{\text{stake}}(t+1)_i}{S_{t+1} - \ell_{t+1}} - \frac{\pi_{\text{stake}}(t)_i}{S_t - \ell_t} = \frac{\pi_{\text{stake}}(t)_i}{S_t - \ell_t} \frac{\alpha_i W_i(t)}{S_{t+1} - \ell_{t+1}} - \frac{\pi_{\text{stake}}(t)_i}{S_t - \ell_t} \\ &= \frac{\pi_{\text{stake}}(t)_i}{S_t - \ell_t} \left(\frac{\alpha_i W_i(t)}{S_{t+1} - \ell_{t+1}} - 1 \right) \end{aligned} \quad (32)$$

Applying [Equation 11](#) at $s = \frac{\ell_{t+1}}{S_{t+1}}$ to [Equation 32](#) gives:

$$\mu_{\text{stake}}(t)_i = \frac{\pi_{\text{stake}}(t)_i}{S_t - \ell_t} \left(\frac{\alpha_i W_i(t)}{S_{t+1} - \ell_{t+1}} - 1 \right) = \frac{\pi_{\text{stake}}(t)_i}{S_t - \ell_t} \left(\frac{\alpha_i W_i(t)}{S_{t+1}} \frac{1}{1 - \frac{\ell_t}{S_t}} - 1 \right)$$

$$< \frac{\pi_{\text{stake}}(t)_i}{S_t - \ell_t} \left(\frac{\alpha_i W_i(t)}{S_{t+1}} \left(\frac{\ell_{t+1}}{S_{t+1}} \left(\frac{1 - \delta}{\delta} \right) + 1 \right) - 1 \right) \quad (33)$$

Recall that $\ell_{t+1} = \gamma_t \sum_{i=1}^n \beta_i W_i(t)$ and that for positive functions $f \geq g$, $\Pr[f < \epsilon] \leq \Pr[g < \epsilon]$. Note that $\pi_{\text{stake}}(t)_i$ as adapted to a natural filtration $\mathcal{F}_t$ that is non-anticipating, by construction.¹³ This means that we can compute probabilities conditional on $\mathcal{F}_t$, e.g., $\Pr[\ell_t | \mathcal{F}_t] = \ell_t$. Using this and recalling that the money supply S_t is a deterministic function yields:

$$\begin{aligned} \Pr[|\mu_{\text{stake}}(t)_i| < \epsilon | \mathcal{F}_t] &\geq \Pr \left[\left| \frac{\pi_{\text{stake}}(t)_i}{S_t - \ell_t} \left(\frac{\alpha_i W_i(t)}{S_{t+1}} \left(\frac{\ell_{t+1}}{S_{t+1}} \left(\frac{1 - \delta}{\delta} \right) + 1 \right) - 1 \right) \right| < \epsilon \middle| \mathcal{F}_t \right] \\ &\geq \Pr \left[\left| \frac{\pi_{\text{stake}}(t)_i}{S_t - \ell_t} \left(\frac{\alpha_i W_i(t)}{S_{t+1}} \left(\frac{\ell_{t+1}}{S_{t+1}} \left(\frac{1 - \delta}{\delta} \right) \right) \right) \right| < \epsilon \middle| \mathcal{F}_t \right] (1 - e^{-\frac{S_{t+1}}{W_i(t)}}) \\ &\geq \left(1 - e^{-S_{t+1}/W_i(t)} \right) \Pr \left[\alpha_i \sum_{i=1}^n \beta_i < \epsilon \left(\frac{S_t - \ell_t}{P_i(t)} \right) \left(\frac{S_{t+1}}{\gamma_t} \right) \left(\frac{\delta}{1 - \delta} \right) \right] \end{aligned}$$

where [Equation 35](#) comes from splitting [Equation 34](#) into two terms by conditioning on $A = \left\{ \frac{\alpha_i W_i(t)}{S_{t+1}} - 1 < 0 \right\}$ and A^c and throwing away the term dependent on A^c and the third inequality ([Equation 36](#)) comes from $W_i(t) < S_{t+1}$ and majorizing the sum $\sum_i \beta_i W_i(t)$. Note that [Equation 36](#) is the cumulative distribution function for the product of an exponentially distributed random variable α and an Erlang distributed random variable, $\sum_{i=1}^n \beta_i$. Let $A \sim \text{Exp}(\tau_{\text{stake}})$ and $B \sim \text{Erlang}(n, \tau_{\text{lend}})$. Then the product CDF¹⁴ is:

$$\begin{aligned} \Pr[AB \leq y] &= \int_0^\infty \frac{\tau_{\text{stake}}^n}{(n-1)!} x^{n-1} e^{-\tau_{\text{lend}} x} \Pr \left[A < \frac{y}{x} \right] dx \\ &= 1 - \frac{\tau_{\text{stake}}^n}{(n-1)!} \int_0^\infty x^{n-1} e^{-\tau_{\text{stake}} x - \tau_{\text{lend}} \frac{y}{x}} dx \\ &= 1 - \frac{\tau_{\text{stake}}^n}{(n-1)!} \int_0^\infty e^{-\tau_{\text{stake}} x - \tau_{\text{lend}} \frac{y}{x} - (n-1) \log x} dx \quad (37) \end{aligned}$$

The integrand in [Equation 37](#) is maximized around the saddle point of the large deviation function $\Lambda(x) = \tau_{\text{stake}} x - \frac{\tau_{\text{lend}} y}{x} - (n-1) \log x$. Let $x^* = \arg \max \Lambda(x)$. Using the first-order condition $\nabla \Lambda(x) = 0$ and the quadratic formula, we get:

$$x^* = \frac{n-1}{2\tau_{\text{stake}}} \left(\sqrt{1 + \frac{4\tau_{\text{stake}}\tau_{\text{lend}}y}{(n-1)^2}} - 1 \right)$$

From [Assumption 13](#), we know that $4\tau_{\text{stake}}\tau_{\text{lend}} < (n-1)^2$, so $x^* \approx 2\tau_{\text{stake}}y$. Applying the method of steepest descent [\[48\]](#) to [Equation 37](#) implies that there exists $C > 0$ such that

$\int_0^\infty e^{-\tau_{\text{stake}}x - \tau_{\text{lend}}\frac{y}{x} - (n-1)\log x} dx < Ce^{-\Lambda(x^*)}$ so that we have:

$$\Pr[AB \leq y] \geq 1 - \frac{C\tau_{\text{stake}}^n}{(n-1)!} e^{-\Lambda(X^*)} = 1 - \frac{C\tau_{\text{stake}}^n}{(n-1)!} (2\tau_{\text{stake}}y)^{n-1} e^{-\tau_{\text{stake}}^2 y + \frac{\tau_{\text{lend}}}{2\tau_{\text{stake}}}}$$

Combining [Equation 36](#) with [Equation 38](#) gives the result.

Footnotes

1.

The majority of Proof of Stake systems use traditional Byzantine Fault Tolerant algorithms (Lamport *et al.*; Castro *et al.*) once validators are sampled, which is why they are only resistant to 33% Byzantine actors versus 50% for Nakamoto consensus.

Leslie Lamport, Robert Shostak, and Marshall Pease. 1982. The Byzantine generals problem. *ACM Transactions on Programming Languages and Systems (TOPLAS)* 4, 3 (1982), 382–401.

Miguel Castro, Barbara Liskov, and others. 1999. Practical Byzantine fault tolerance. In *OSDI*, Vol. 99. 173–186. [↗](#)

2.

It should be noted that this has been a stated goal of large pools of validators and the existence of staking services that pool tokens and operate nodes exemplifies this collective desire to have a large portion of the token supply staked (Tezos; Manian).

Tezos (Figment Networks). (Oct 2019). <https://figment.network/networks/tezos/>

Zaki Manian. 2019. Iqlusion: accelerating cryptofinance. (Oct 2019). <https://iqlusion.io> [↗](#)

3.

In graphical models and probability, such as the Ising model or spin glasses, this is known as “quenched disorder” (Mézard *et al.*).

Marc Mézard, and Andrea Montanari. 2009. *Information, physics, and computation*. Oxford University Press. [↗](#)

4.

For staking, the relative yield is defined as future expected staked wealth divided by current stake wealth minus 1, which is

$$\frac{\pi_{stake}(t)_i + R_t \frac{\pi_{stake}(t)_i}{s_t - \ell_t}}{\pi_{stake}(t)_i} - 1 = \frac{R_t}{S_t - \ell_t}$$

This is the standard form of how yields are specified in traditional Markowitz portfolio theory (Markowitz).

Harry Markowitz. 1952. Portfolio selection. *The journal of finance* 7, 1 (1952), 77–91.

⌞

5. This is implemented in a manner analogous to how hedge funds only let you recall your invested capital in small chunks over a window of time. ⌞

6.

The easiest way to see this is to note that one expects a central limit theorem (e.g., Marchenko-Pastur law (Tao)) for the distribution of $\frac{1}{\sqrt{n}} \sum_{i=1}^n \Sigma_i$. Thus, as the number of participants goes to infinity, we get closer and closer to having the staticity assumption hold true.

Terence Tao. 2012. *Topics in random matrix theory*. Vol. 132. American Mathematical Soc. ⌞

7. Also note that the equilibrium assumption ([Assumption 11](#)) implicitly also requires a lower bound on n that is a function of τ_{stake} and τ_{lend} in order for the covariance matrix distribution to be ϵ away from the limit distribution. ⌞

8.

In the constrained version of the problem needs to satisfy Karush-Kuhn-Tucker conditions (Mangasarian) that enforce $w^T 1 = 1$ and $w_i > 0, \forall i$. For simplicity of exposition, we elided these during this section (even though these conditions are satisfied during simulation).

Olvi L Mangasarian. 1994. *Nonlinear programming*. SIAM. ⌞

9.

In particular, they are random at time $t = 0$ and defined in $\mathcal{F}_0$, but are constant throughout. This is a measure theoretic way of defining quenched disorder (Talagrand).

Michel Talagrand. 2010. *Mean field models for spin glasses: Volume I: Basic examples*. Vol. 54. Springer Science & Business Media. ⌞

10. Since $U_t = \frac{kS_t}{\ell_t + kS_t}$ and [Assumption 12](#) implies that $\ell_t < (1 - \delta)S_t$, we have $U_t > \frac{k}{k+1-\delta}$ and $\gamma_t > \beta_1 \left(\frac{k}{k+1}\right)^2 + \beta_0 \frac{k}{k+1}$. ⌞

11. If $\pi_{stake}(0)_{(1)} < \dots < \pi_{stake}(0)_{(n)}$ are the order statistics for the stake distribution, then $\pi_{stake}(0)_{(i+1)} - \pi_{stake}(0)_{(i)} \sim \text{Exp}(\frac{\lambda_{stake}}{n-i})$ and by the tower property for telescoping sums, $E[\pi_{stake}(0)_{(i)}] = \Theta(\frac{\log i}{\lambda_{stake}})$. [↵](#)
12. Note that [Assumption 10](#) allows us to sample the agent portfolio updates in any order. [↵](#)
13. As we sample the PoS algorithm and Markowitz weights randomly (and not psuedorandomly), there is no advantage to looking into the future. [↵](#)
14. Recall that the density of an exponential random variable with parameter α is $\alpha e^{-\alpha x}$ and the PDF of an Erlang distribution with parameters n, λ is $\frac{\lambda^n}{(n-1)!} x^{n-1} e^{-\lambda x}$. [↵](#)

Citations

1. QuantumMechanic. (2011). *Proof of stake instead of proof of work*. BitcoinTalk. Retrieved from <https://bitcointalk.org/index.php?topic=27787.0> [↵](#)
2. Pass, R., Seeman, L., & Shelat, A. (2017). Analysis of the blockchain protocol in asynchronous networks. In *Annual international conference on the theory and applications of cryptographic techniques* (pp. 643–673). Springer. [↵](#)
3. Kiayias, A., Russell, A., David, B., & Oliynykov, R. (2017). Ouroboros: A provably secure proof-of-stake blockchain protocol. In *Annual international cryptology conference* (pp. 357–388). Springer. [↵](#)
4. King, S., & Nadal, S. (2012). Ppcoin: Peer-to-peer crypto-currency with proof-of-stake. *Self-Published Paper, August, 19*. [↵](#)
5. Saleh, F. (2019). Blockchain without waste: Proof-of-stake. *Available at SSRN 3183935*. [↵](#)
6. Daian, P., Pass, R., & Shi, E. (2019). Snow white: Robustly reconfigurable consensus and applications to provably secure proof of stake. In *International conference on financial cryptography and data security* (pp. 23–41). Springer. [↵](#)
7. Neuder, M., Moroz, D. J., Rao, R., & Parkes, D. C. (2019). Selfish behavior in the tezos proof-of-stake protocol. *ArXiv Preprint ArXiv:1912.02954*. [↵](#)
8. Poelhekke, S., & Ploeg, van der R. (2007). Volatility, financial development and the natural resource curse. [↵](#)
9. Houy, N. (2014). It will cost you nothing to 'Kill' a proof-of-stake crypto-currency. *Available at SSRN 2393940*. [↵](#)

10. Gaži, P., Kiayias, A., & Russell, A. (2018). Stake-bleeding attacks on proof-of-stake blockchains. In *2018 crypto valley conference on blockchain technology (cvcbt)* (pp. 85–92). IEEE. [↵](#)
11. Brown-Cohen, J., Narayanan, A., Psomas, A., & Weinberg, S. M. (2019). Formal barriers to longest-chain proof-of-stake protocols. In *Proceedings of the 2019 acm conference on economics and computation* (pp. 459–473). ACM. [↵](#)
12. Chitra, T., Chiang, R., Morrow, J., & Sood, A. (2019). *Comment for cftc rfi on crypto-asset mechanics and markets. Comment for Orders and Other Announcements* 83 FR 64563. Commodities; Futures Trading Commission. Retrieved from <https://comments.cftc.gov/PublicComments/ViewComment.aspx?id=61960> [↵](#)
13. Klages-Mundt, A., & Minca, A. (2019). (In) stability for the blockchain: Deleveraging spirals and stablecoin attacks. *ArXiv Preprint ArXiv:1906.02152*. [↵](#)
14. Warren, W., & Bandeau, A. (2017). Ox: An open protocol for decentralized exchange on the ethereum blockchain. URL: <https://Github.Com/OxProject/Whitepaper>. [↵](#)
15. Leshner, R., & Hayes, G. (2019). *Compound: The money market protocol*. [↵](#)
16. Kamvar, S., Olszewski, M., & Reinsberg, R. (2019). *Celo: A multi-asset cryptographic protocol for decentralized social payments*. [↵](#)
17. Al-Naji, N., Chen, J., & Diao, L. (2017). *Basis: A price-stable cryptocurrency with an algorithmic central bank*. Basis team. [↵](#)
18. MakerDAO. (2017). The dai stablecoin system. URL: <https://Makerdao.Com/Whitepaper/DaiDec17WP.Pdf>. [↵](#)
19. Chitra, T., Quaintance, M., Haber, S., & Martino, W. (2019). Agent-based simulations of blockchain protocols illustrated via kadena’s chainweb. In *2019 ieee european symposium on security and privacy workshops (euros pw)* (pp. 386–395). <https://doi.org/10.1109/EuroSPW.2019.00049> [↵](#)
20. Markowitz, H. (1952). Portfolio selection. *The Journal of Finance*, 7(1), 77–91. [↵](#)
21. Buchman, E. (2019). Byzantine fault tolerant state machine replication in any programming language. In *Proceedings of the 2019 acm symposium on principles of distributed computing* (pp. 546–546). ACM. [↵](#)
22. Goodman, L. (2014). Tezos—a self-amending crypto-ledger white paper. URL: https://www.Tezos.Com/Static/Papers/White\textbackslash_paper.Pdf. [↵](#)

23. Koblitz, N., & Menezes, A. J. (2015). The random oracle model: A twenty-year retrospective. *Designs, Codes and Cryptography*, 77(2–3), 587–610. [↵](#)
24. Chan, T.-H. H., Pass, R., & Shi, E. (2018). PaLa: A simple partially synchronous blockchain. *IACR Cryptology EPrint Archive*, 2018, 981. [↵](#)
25. Kane, D. M. (2016). A new approach to distribution testing. [↵](#)
26. Pass, R., & Shi, E. (2017). The sleepy model of consensus. In *International conference on the theory and application of cryptography and information security* (pp. 380–409). Springer. [↵](#)
27. Gilad, Y., Hemo, R., Micali, S., Vlachos, G., & Zeldovich, N. (2017). Algorand: Scaling byzantine agreements for cryptocurrencies. In *Proceedings of the 26th symposium on operating systems principles* (pp. 51–68). ACM. [↵](#)
28. Fanti, G., Kogan, L., Oh, S., Ruan, K., Viswanath, P., & Wang, G. (2019). Compounding of wealth in proof-of-stake cryptocurrencies. In *International conference on financial cryptography and data security* (pp. 42–61). Springer. [↵](#)
29. Rocket, T., Yin, M., Sekniqi, K., Renesse, van R., & Sirer, E. G. (2019). Scalable and probabilistic leaderless bft consensus through metastability. *ArXiv Preprint ArXiv:1906.08936*. [↵](#)
30. Chitra, T., & Chitra, U. (2019). Committee selection is more similar than you think: Evidence from avalanche and stellar. *ArXiv Preprint ArXiv:1904.09839*. [↵](#)
31. Pass, R., & Shi, E. (2018). Thunderella: Blockchains with optimistic instant confirmation. In *Annual international conference on the theory and applications of cryptographic techniques* (pp. 3–33). Springer. [↵](#)
32. Bagaria, V., Kannan, S., Tse, D., Fanti, G., & Viswanath, P. (2018). Deconstructing the blockchain to approach physical limits. *ArXiv Preprint ArXiv:1810.08092*. [↵](#)
33. Baudet, M., Ching, A., Chursin, A., Danezis, G., Garillot, F., Li, Z., ... Sonnino, A. (2018). *State machine replication in the libra blockchain*. Technical report, The Libra Association. [↵](#)
34. Weber, W. E. (2016). *A bitcoin standard: Lessons from the gold standard*. Bank of Canada Staff Working Paper. [↵](#)
35. Li, D., & Ng, W.-L. (2000). Optimal dynamic portfolio selection: Multiperiod mean-variance formulation. *Mathematical Finance*, 10(3), 387–406. [↵](#)

36. Zhou, X. Y., & Yin, G. (2003). Markowitz's mean-variance portfolio selection with regime switching: A continuous-time model. *SIAM Journal on Control and Optimization*, 42(4), 1466–1482. [↵](#)
37. Boyd, S., Busseti, E., Diamond, S., Kahn, R. N., Koh, K., Nystrup, P., ... others. (2017). Multi-period trading via convex optimization. *Foundations and Trends in Optimization*, 3(1), 1–76. [↵](#)
38. De Santis, G., & Gerard, B. (1997). International asset pricing and portfolio diversification with time-varying risk. *The Journal of Finance*, 52(5), 1881–1912. [↵](#)
39. Asmussen, S., & Hering, H. (1983). *Branching processes* (Vol. 3). Springer. [↵](#)
40. Billingsley, P. (2008). *Probability and measure*. John Wiley & Sons. [↵](#)
41. Kondor, D., Pósfai, M., Csabai, I., & Vattay, G. (2014). Do the rich get richer? An empirical analysis of the bitcoin transaction network. *PloS One*, 9(2), e86197. [↵](#)
42. Alvarez-Ramirez, J., Rodriguez, E., & Ibarra-Valdez, C. (2018). Long-range correlations and asymmetry in the bitcoin market. *Physica A: Statistical Mechanics and Its Applications*, 492, 948–955. [↵](#)
43. Shalizi, C. (2007). So you think you have a power law—well isn't that special. *Three-Toed Sloth*. [↵](#)
44. Stumpf, M. P., & Porter, M. A. (2012). Critical truths about power laws. *Science*, 335(6069), 665–666. [↵](#)
45. Sipser, M. (2012). *Introduction to the theory of computation*. Cengage learning. [↵](#)
46. Burkholder, D. L. (1973). Distribution function inequalities for martingales. *The Annals of Probability*, 19–42. [↵](#)
47. Lowther, G. (2010). *The burkholder-davis-gundy inequality*. *Almost Sure*. Retrieved from <https://almostsure.wordpress.com/2010/04/06/the-burkholder-davis-gundy-inequality/> [↵](#)
48. Fedoryuk, M. (1987). *Asymptotic: Integrals and series*. Nauka. Phys.-math. lit. [↵](#)