cfbenchmarks
a kraken company

Understanding the Dynamics of Ethereum Staking Returns: Risk and Reward Framework

Author: James Flamant, Senior Product Manager, CF Benchmarks

October 2024

Contents

1 Introduction

Paper Objective - Understanding Staking from a Risk-Reward Perspective

The primary objective of this paper is to explore Ethereum staking through a probabilistic framework, with a focus on understanding both the potential rewards and the inherent risks involved. Ethereum's Proof of Stake (PoS) mechanism provides validators with rewards for securing the network, but the outcomes are not uniform. Instead, they are influenced by a mix of random and deterministic factors, creating variability in the rewards a validator or a group of validators can expect to receive. By modelling these rewards probabilistically, the paper aims to offer a clear understanding of staking returns and how to assess the reward risk profile when participating in Ethereum staking. Other risks associated with staking, such as operational risks, are not the subject of this paper.

To achieve this, the paper will first break down the staking reward framework, offering a detailed explanation of how various types of rewards are distributed. These include rewards from attestation, block proposing, sync committees, and tips, each contributing differently to the overall reward structure. Next, the paper will demonstrate how the probabilistic nature of these rewards can be modeled, allowing readers to understand the range of possible outcomes for individual validators as well as for groups of validators, or pools.

In addition to explaining the structure of rewards, the paper will analyze how key parameters—such as the number of validators, or the level of tips —impact staking outcomes and risk profiles. This analysis will provide insights into how changes in these parameters affect both the overall reward distribution and the associated risks. Finally, the paper will interpret these findings through financial metrics, such as return distributions, expected value, and confidence intervals, offering readers a way to understand staking in financial terms.

By the end of the paper, readers will have a comprehensive understanding of the probabilistic structure that governs Ethereum staking rewards. They will also gain an appreciation for how changes in key parameters, such as validator numbers and tip levels, affect both the average rewards and the risks of staking. With this framework, readers will be equipped to view staking as a financial investment, complete with tools to evaluate the risk-return profile of their staking activities. By its very nature, this paper will utilize terminology that is native to blockchains and the Ethereum network, and while we assume that the reader has a functioning understanding of blockchains, we define some of the terms that are specific to the staking process below:

Consensus Layer

The consensus layer of Ethereum is responsible for ensuring the security and agreement of all participants in the network. It governs how validators propose, validate, and confirm blocks, relying on Ethereum's Proof of Stake (PoS) system. This layer handles validator responsibilities, including attesting to block proposals, ensuring no double spends, and maintaining consensus on the blockchain's state by confirming that the same blocks are added to the chain across all nodes.

Execution Layer

The execution layer processes and executes transactions and smart contracts. It handles the computation of smart contracts, the state transitions that result from transactions, and records the resulting state changes in blocks. While the consensus layer governs block validation and inclusion, the execution layer ensures that the transactions contained within those blocks are properly executed. Together, the consensus and execution layers ensure the validity and correctness of Ethereum's blockchain.

Proposing Blocks

Proposing blocks is one of the key duties of an Ethereum validator. Validators are randomly selected to propose a block for each new slot (every $\sim$12 seconds). When chosen, the validator assembles a block by including pending transactions and broadcasting it to other validators for attestation. Validators receive rewards for proposing blocks and can also capture additional fees such as tips from included transactions and rewards from MEV (Maximal Extractable Value).

Attesting to Blocks

Attesting is the process by which validators vote on blocks proposed by other validators, confirming that the block follows the network's rules and should be added to the chain. Validators attest to multiple blocks per epoch, and their votes are key to determining finality—the point at which a block is permanently added to the chain. Timely and accurate attestations are essential to maintaining consensus and are rewarded accordingly.

Sync Committee

The sync committee is a special group of validators randomly selected for a set period (typically a day and a half) to assist light clients—nodes that do not store the entire blockchain—by signing block headers. This ensures that these light clients can remain synced with the network without needing to process every block. Validators in the sync committee receive rewards for their participation.

Double Signing

Double signing occurs when a validator signs and broadcasts two conflicting blocks for the same slot. This is a serious violation of consensus rules and results in a slashing penalty. Slashing not only removes a portion of the validator's staked ETH but also ejects the validator from the active set, significantly reducing their future staking rewards and harming the network's security.

2 The Staking Reward System

2.1 How Ethereum Staking Rewards Work

Ethereum's Proof of Stake (PoS) system operates by requiring validators to stake ETH as collateral in order to participate in securing the network and validating transactions. Validators, who are responsible for proposing and attesting to blocks, earn rewards for their participation but are also subject to penalties for misbehavior or inactivity. This staking mechanism ensures that validators are financially incentivized to act in the best interest of the network, while discouraging malicious or negligent actions.

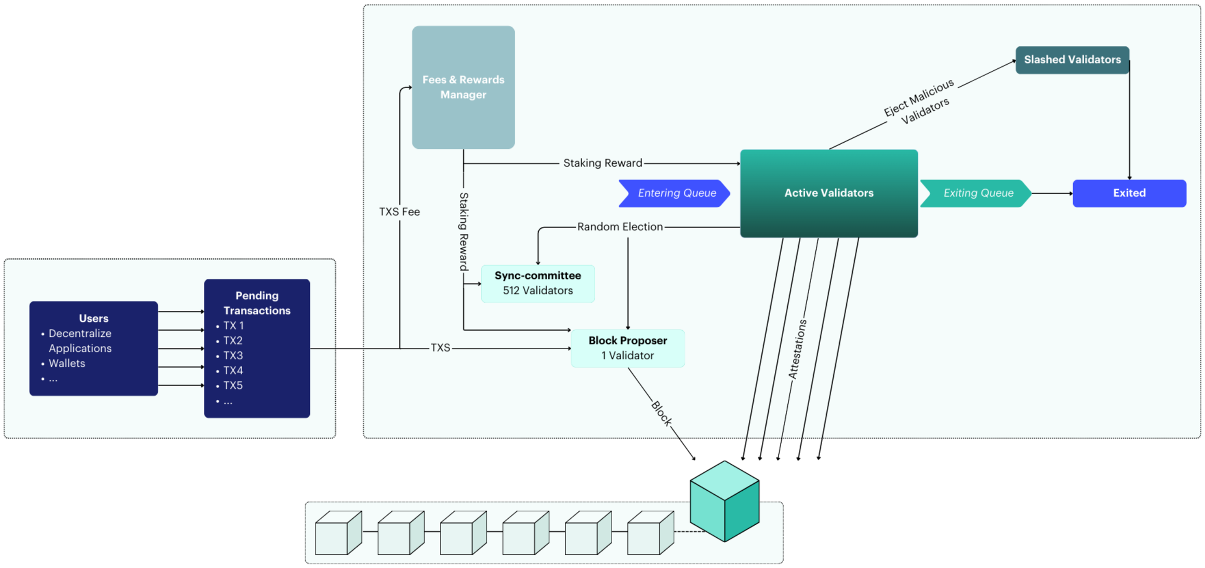

Figure 1: Infographic: Ethereum

To become a validator, a participant must meet both technical and capital requirements. From a technical perspective, validators are required to run specialized software that allows them to perform tasks such as proposing and attesting to blocks. This involves maintaining a well-connected validator node that is always synchronized with the Ethereum network. Validators must remain online and actively participate in proposing and attesting to maximize their rewards. Any downtime or technical issues can lead to missed rewards and potential penalties.

Due to the fact that Ethereum is a pseudonymous blockchain network, individuals or institutions are not limited in the number of validators they can operate. In terms of capital requirements, each validator must stake 32 ETH as a deposit. This capital allocation serves as collateral, ensuring that validators have a financial stake in the network's security. The staked ETH is locked for the duration of the validator's participation and cannot be accessed until they exit the staking process, which is likely to not be instantaneous and is dependent on the network conditions. This mechanism ensures that validators have ”skin in the game” and are financially motivated to behave honestly and efficiently.

The rewards earned by validators come from two main sources: the consensus layer and the execution layer. Consensus layer rewards are the primary source and are distributed for tasks such as attesting to blocks, proposing new blocks, and participating in sync committees. These rewards are predictable and deterministic, as they are part of Ethereum's native issuance of ETH. The amount distributed per block is fixed, but the share received by each validator depends on the total number of active validators and their participation.

In addition to consensus rewards, validators can also earn rewards from the execution layer in the form of tips. These tips are paid by users who wish to prioritize their transactions within the block. Since tips are part of the transaction fees and fluctuate based on user demand and network activity, they add an element of variability to a validator's total rewards. When the network is congested or during periods of high transaction volume, tips can significantly increase the total rewards for validators, while in quieter periods, the rewards from tips may be minimal.

While the reward system incentivizes active participation, validators are also subject to penalties if they fail to fulfill their duties or if they engage in malicious behavior. Validators who become inactive—whether due to going offline, failing to propose blocks, or missing attestation duties—face penalties that gradually reduce their staked ETH. Although these inactivity penalties are relatively small in the short term, they can accumulate over time if the validator remains offline, leading to a significant reduction in their capital.

More serious misbehavior, such as attempting to propose conflicting blocks or double-signing, results in slashing.

Balancing Act: How Transaction Fees and Staking Rewards Influence Ethereum's Monetary Policy

In addition to consensus and execution rewards, Ethereum's transaction fee mechanism plays a significant role in the network's economics and the overall supply dynamics of ETH. Since the implementation of EIP-1559 (the London hard fork), Ethereum's transaction fees consist of two components: the base fee and the priority fee (also known as tips). The base fee is a mandatory fee per unit of gas that dynamically adjusts based on network congestion and demand. Importantly, this base fee is **burned**, meaning it is permanently removed from circulation. The priority fee is an optional amount that users can include to incentivize validators to prioritize their transactions; these tips are awarded directly to validators as part of their execution layer rewards.

Conversely, the staking rewards paid out to validators as compensation for performing their duties—block proposing, sync committee participation, and attestations—are newly created ETH (in blockchain terminology, they are minted). The protocol burns the base fees collected from transactions but mints new ETH to pay out staking rewards. The net balance between the ETH burned through base fees and the ETH minted as staking rewards determines whether ETH's overall supply is inflationary or deflationary at any given time. When the amount of ETH burned exceeds the amount minted, the total supply decreases, leading to deflation. Conversely, if more ETH is minted than burned, the supply increases, resulting in inflation.

Slashing

Slashing is a punitive mechanism that permanently removes a portion of the validator's staked ETH. It must be noted that Ethereum does not operate a judicial system; there are no courts to determine whether any given outcome was the result of malicious behavior intended to negatively impact the integrity of the ledger. The slashing mechanism is activated once certain phenomena are observed. The cause of this can be malicious or mere incompetence, but both would trigger a slashing event, with the severity depending on the nature of the infraction. In extreme cases, slashing can lead to a validator losing their entire stake and being ejected from the validator set. This ensures that validators not only have a strong disincentive to act against the network's interests but also forces validators to act in the best interests of the network, such as ensuring that validator nodes are running the latest software and correctly maintained, as the financial consequences of slashing can be severe. Ethereum's staking system is designed to align the interests of validators with the security of the network. By offering rewards for participation and penalties for inactivity or misbehavior, it creates a balance that motivates validators to act honestly and efficiently, ensuring the ongoing stability and security of Ethereum.

2.2 Random vs. Deterministic Rewards (Overview of Reward Types)

Ethereum's Proof of Stake (PoS) reward system offers a combination of random and deterministic rewards, each designed to incentivize validator participation and maintain the network's security. The rewards can be broken down based on the roles validators play and the probabilistic nature of the outcomes for each role.

Deterministic

These rewards are deterministic, meaning every active validator can fulfill the associated job and receives a predictable portion of the reward.

Attestation Rewards:

Validators attest at each epoch, confirming the legitimacy of transactions and the state of the blockchain, helping to maintain consensus. Each active validator is assigned to one of the 32 slots of the epoch to attest to the block and provide the state of the blockchain. Attestation rewards make up the majority of the consensus layer rewards, accounting for 84.375% of the consensus layer staking rewards distributed to validators.

Random

Block Proposer Rewards:

For every slot, one validator is randomly selected from the network's active validator set to be the block proposer. Block proposers receive a reward for their role, which is subject to chance due to the randomness of the selection process. The selection process is a random sampling with replacement, where all active validators in the network have an equal chance to be selected for each block, regardless of past selections. Block proposer rewards constitute around 12.55% of consensus rewards. Since only one validator per block is chosen, this reward introduces a significant level of variance in individual validator returns over time.

Note: Slots and blocks are one-for-one provided the block is correctly proposed. In this paper, we assume all blocks are correctly proposed and therefore slot and block can be used interchangeably.

Sync Committee Rewards:

Sync committees are groups of 512 randomly chosen validators tasked with ensuring that Ethereum's light clients (like wallets or mobile devices) can verify blocks efficiently. A newly elected sync committee comes into play every 256 epochs ($\sim$24 hours). Validators within the sync committee earn additional rewards for this role. Similar to the block proposer selection, the sync committee selection process is a random sampling with replacement. Sync committee rewards represent 3.125% of consensus rewards and are another random component that introduces variability in returns, as participation depends on being randomly selected for the committee.

Tips:

Tips are additional rewards that the block proposer can earn from users, incentivizing the validator to prioritize certain transactions. Tips are derived from the execution layer, where users pay them to validators to ensure either (i) faster processing of their transactions in most cases or (ii) specific transaction processing order for the special case of Maximal Extractable Value (MEV) transactions. The variability in tips is driven by market demand and network congestion, making them one of the most unpredictable components of the reward structure. On average over the first half of 2024, tips account for around 15% of the total rewards, but this can fluctuate significantly depending on network activity.

Note: The tips reward will be considered stable in Chapters 3 and 4, at 15% of the total consensus reward. Analysis of the impact of considering the tips reward stochastic will be explored in Chapter 5.

This combination of random and deterministic rewards creates a probabilistic return profile for a validator. While attestation rewards provide a stable, predictable income, the randomness of block proposer selection, sync committees, and tips introduces variability. Over time, the balance of these factors determines the overall expected return and the risk level associated with staking.

2.3 Role of Model Parameters (Validator Count, Tips)

The rewards that validators earn in Ethereum's staking system are deeply influenced by several key model parameters, primarily the number of active validators and the amount of tips in the system. These parameters shape both the size of rewards and the variability in the outcomes.

Number of Active Validators

The total number of active validators directly affects how rewards are distributed. As the number of validators increases, the per-validator reward decreases, since rewards are shared among a larger group.

Currently, there are approximately one million validators on the network. As more validators join, the deterministic portion of rewards (e.g., attestation rewards) becomes smaller on a per-validator basis. This introduces a form of diminishing returns for validators, as the larger the validator pool, the lower the expected reward per unit of effort.

This relationship can be modelled probabilistically, showing that the expected reward for any individual validator decreases as more validators join, leading to a narrower reward distribution but also reducing the potential for outsized gains.

Tips

Tips are an essential source of randomness in the staking reward structure. Unlike the consensus layer rewards (attestations, block proposals), which are fixed per block, tips are variable and depend on the state of the network.

When the network is congested, users are more willing to pay higher tips to have their transactions prioritized. This can lead to a temporary increase in validator income. Conversely, in periods of low network demand, tips may be minimal, leading to a drop in rewards from the execution layer.

The tips parameter can vary significantly based on Ethereum's market environment. For instance, during high periods of decentralized finance (DeFi) activity or network congestion, tips tend to spike, while in more stable or quieter times, tips can decline, impacting the total reward a validator can expect to receive.

The variability introduced by these parameters means that staking rewards cannot be viewed as static. Instead, they must be understood probabilistically, as fluctuating with network conditions and the behavior of other validators. By accounting for these factors, validators can better anticipate their potential returns and adjust their strategies to optimize their participation in the network.

3 Validator Return Profile

The return profile of a validator can be obtained by first aggregating the staking rewards over the period of interest (yearly) for each reward type, and then aggregating the different rewards to obtain the global distribution.

Information about the Model

Hyper-parameters

The exogenous parameters that impact the quantity of rewards between slots (e.g., tips level, number of active validators, number of slashed validators, etc.) are considered fixed in Chapters 3 and 4. This also implies that the amount of rewards is constant across slots. The hyper-parameters are defined as:

- Number of active validators: S (e.g., 1,000,000)
- Tips level: TS (for Tips Share), e.g., 15% of all consensus rewards at each block
- Slashed validators: 0

Units

Staking rewards on Ethereum are paid in Ether. This paper doesn't cover the returns in other currencies such as USD, which would require taking into account an ETH/USD exchange rate.

Reward Categories

The four reward types can be aggregated into three relevant categories when determining the return profile. Indeed, the recipient of the consensus reward for proposing a block coincides with the tips rewards' recipient, with both of them going to only the block proposer. This helps to simplify the complexity of the model. The category covering the sum of the consensus reward for proposing a block and tips rewards will be called Block Proposer Rewards. The two other categories were described in previous sections: Attestation Rewards and Sync Committee Rewards.

Scaling Factor

We decide in Chapters 3 and 4 to use a scaling factor such that the total annual reward amount is 100 times the number of validators, allowing us to express the different scenario payouts as a percentage of the average validator's reward.

3.1 Validator Payout Return Profile by Reward Type

The objective of this section is to calculate the yearly reward distribution for each of the three staking reward categories independently.

3.1.1 Sync Committee Rewards Payout

An elected sync committee is effective for a period of 256 epochs, with each epoch containing 32 slots, totaling 256 epochs $\times$ 32 slots per epoch $= 8,192$ slots.

The different scenarios can be visualized using a probability tree diagram:

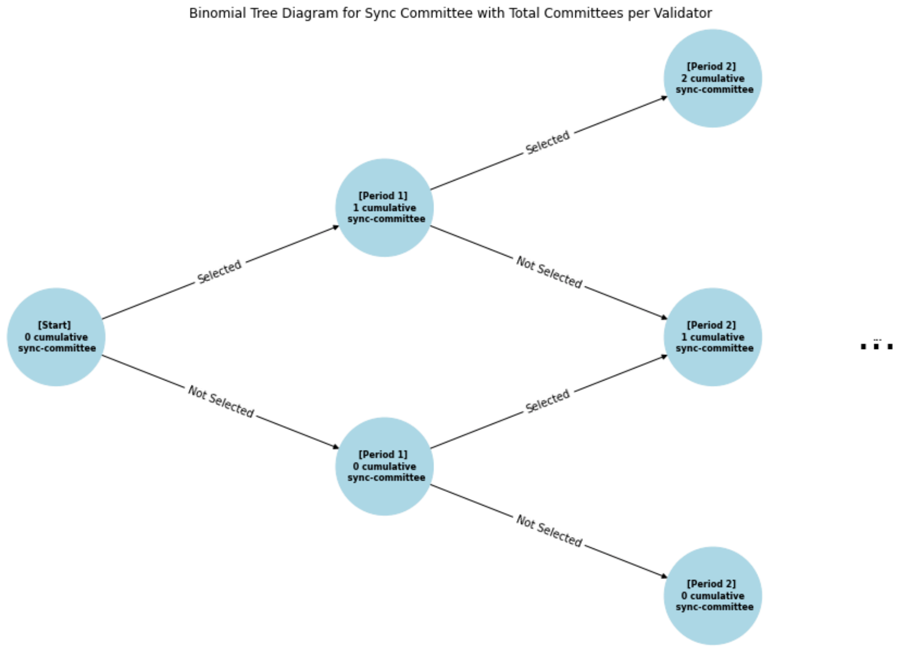

Figure 2: Probability Tree Diagram for Sync Committee Rewards

Mathematically, the reward payout can be modelled by a binomial distribution. The binomial distribution gives the probability of obtaining a specific number of successes in a fixed number of independent trials, each with the same probability of success.

In our case, the success is for the validator to be one of the 512 validators to be elected in the sync committee, failure as not being elected in the sync committee, the probability of success as the probability of being in the 512 validators chosen among all active validators, the probability of failure being one minus the probability of success, and finally the number of trials is the number of sync committee periods during one year.

All active validators have an equal probability of being elected as part of a sync committee. The probability for a validator to be in the 512 validators elected as the sync committee is:

$$p = \frac{512}{S}$$

The number of sync committees during a year can be calculated as the number of seconds in 365 days divided by the multiplication of the number of slots in a sync committee period and the slot time length in seconds. This implies that the number of sync committees per year is:

$$\text{Number of Sync Committees per Year} = \frac{365 \times 24 \times 60 \times 60}{8,192 \times 12} \approx 321$$

The reward amount must consider both the scaling factor and the sync committee rewards in percentage of the total rewards. The sync committee rewards represent 3.125% of the consensus rewards and the tips rewards add an additional TS on top of the consensus rewards. The sync committee rewards percentage of the total staking reward is therefore:

$$\text{Sync Committee Rewards Percentage} = \frac{3.125\%}{1 + TS}$$

Given all the parameters, we can calculate the probability for each number of successes and obtain the following distribution:

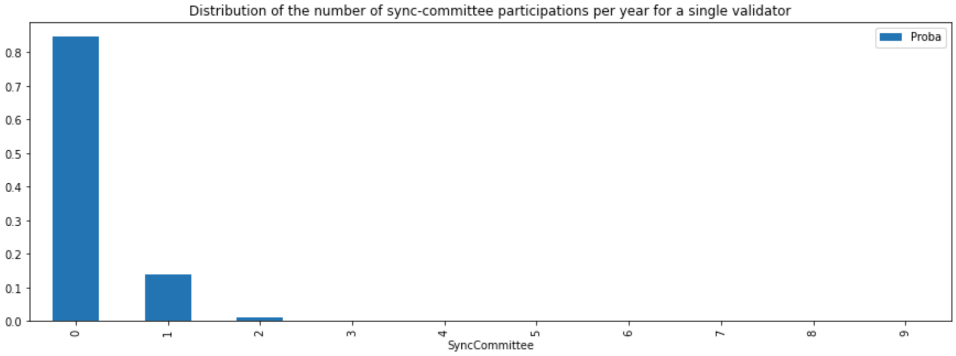

Figure 3: Sync Committee for a single validator

3.1.2 Block Proposer Rewards Payout

An elected block proposer is effective for only one slot.

The different scenarios can be visualized with the probability tree diagram below:

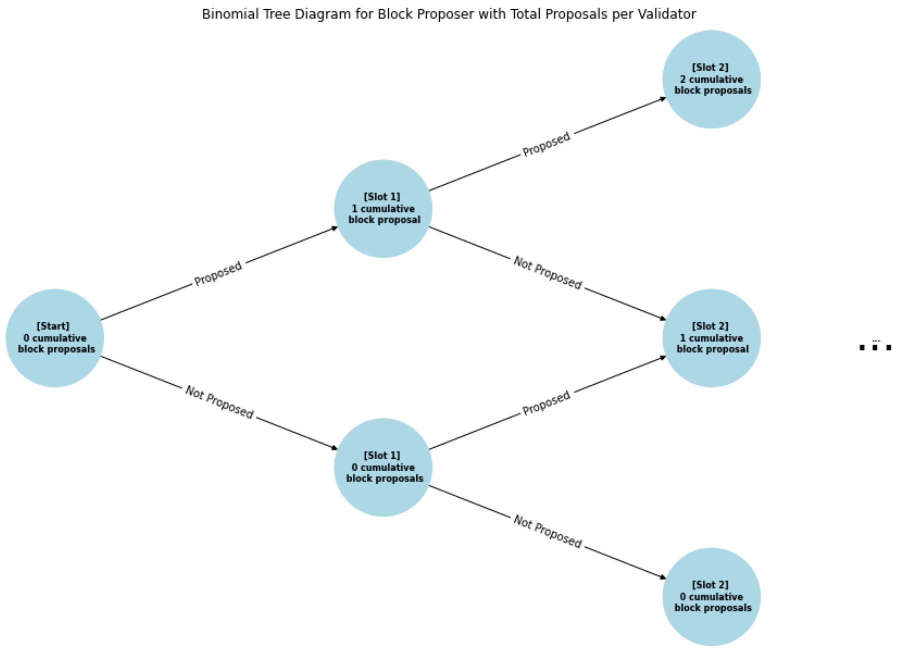
Binomial Tree Diagram for Block Proposer with Total Proposals per Validator

Figure 4: Probability Tree Diagram for Block Proposer Rewards

Similar to the sync committee rewards, the block proposer rewards can be modeled using a binomial distribution. In this case, the success event is for the validator to be the block proposer, failure as the validator not being the block proposer, the probability of success as the probability of being elected the block proposer among all the active validators, the probability of failure being one minus the probability of success, and finally the number of trials is the number slots during one year.

All active validators have an equal probability to be elected as the block proposer. The probability for a validator to be elected as the block proposer is:

$$p = \frac{1}{S}$$

The number of slots during a year can be calculated as the number of seconds in 365 days divided by the slot time length in seconds. This implies that the number of slots per year is:

$$\text{Number of Slots per Year} = \frac{365 \times 24 \times 60 \times 60}{12} = 2,628,000$$

The reward amount must consider both the scaling factor and the block proposer rewards in percentage of the total rewards. The block proposer rewards are the sum of the 12.55% of the consensus rewards plus the tips rewards TS. The block proposer rewards percentage of the total staking reward is therefore:

$$\text{Block Proposer Rewards Percentage} = \frac{12.55\% + TS}{1 + TS}$$

Given all the parameters, we can calculate the probability for each number of successes and obtain the following distribution:

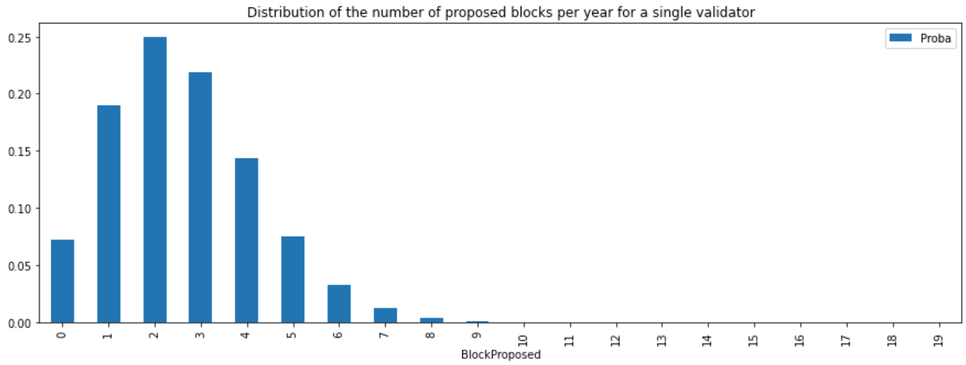

Figure 5: Block Proposition for a single validator

3.1.3 Attestation Rewards Payout

As described in Section 2.2, attestation rewards are distributed evenly across the active validators set, as this paper assumes no operational downtime on either the validator or the blockchain.

The reward amount must consider both the scaling factor and the attestation rewards in percentage of the total rewards. The attestation rewards are 84.375% of the consensus rewards and the tips rewards add an additional TS on top of the consensus rewards. The attestation rewards percentage of the total staking reward is therefore:

$$\text{Attestation Rewards Percentage} = \frac{84.375\%}{1 + TS}$$

3.2 Total Validator Payout Return Profile

A validator's overall annual payout distribution can be obtained by aggregating the three categories' reward distributions and their respective shares of the total reward.

We observe reward payout scenarios concentrated within a tight range, primarily explained by having the only stable reward category, the attestation rewards, representing a majority share of the overall reward.

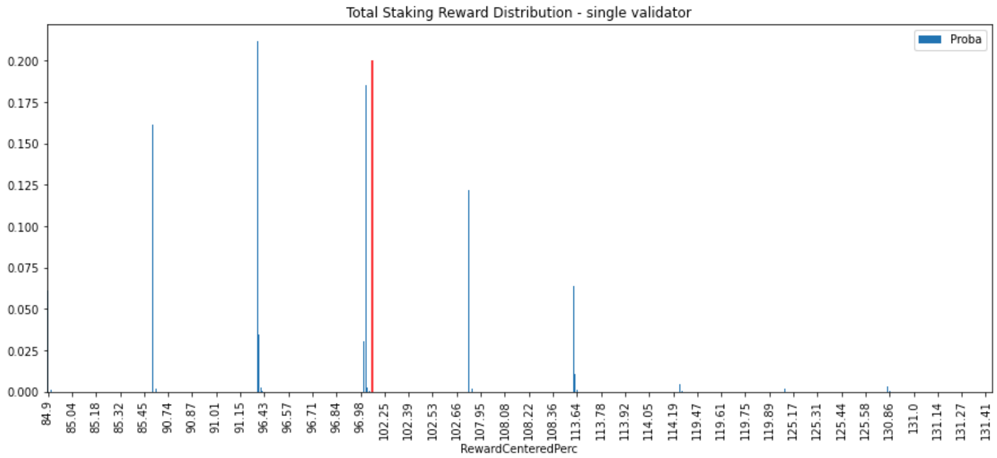

Figure 6: Total Reward Probability Distribution for a single validator

4 Group of Validators Return Profile

4.1 Issues when Scaling to a Large Group of Validators

The Ethereum validator set can be divided into three distinct categories:

1. **Retail Validators**: Validator(s) owned and managed by the same person, usually with one or very few validators.
2. **Institutional Validators**: Validators owned by an institution (e.g., corporation, fund, state), usually managed by a specialized staking service provider.
3. **Pool Validators**: Validators owned by a pool of actors (e.g., liquid staking token treasury), usually managed by multiple specialized staking service providers.

The first category can be satisfied with having only the return profile for a single validator. In the case of up to four or five validators, the payout distribution can be calculated using the resulting scenarios of Chapter 3 and computing all the possible combination scenarios for the group of validators.

However, this approach has exponential complexity and cannot be scaled to a large number of validators. Let's illustrate this with an example considering a group of 4 validators:

To reduce the initial problem size, we consider only scenarios with a probability above 1×10^{-5}. The number of scenarios for the group is the number of permutations with repetition. For example, if the single-validator scenario A has a probability of 6% with an 80% payout (as found in the results of Chapter 3), one permutation can be having all four validators in scenario A. These are permutations with repetition because the order of the scenarios matters. Indeed, the permutation where the first of the four validators is in scenario B and the three others are in scenario A is distinct from the permutation where the first three validators are in scenario A and the last validator is in scenario B. While both scenarios have the same payout structure, they must both be considered. The number of permutations with repetition is the number of single-validator scenarios to the power of the number of validators. In our example, the number of permutations is:

$$40^4 = 2,560,000$$

Below is a table with the number of validators and the associated number of permutations:

In order to obtain the group reward distribution, two metrics must be computed for each permutation:

Number of Validators	Number of Permutations
4	2,560,000
5	102,400,000
6	4,096,000,000
10	$10,485,760,000,000,000$

Table 1: Number of Permutations for Different Validator Group Sizes

- The probability of the permutation, as the product of the probabilities of the single-validator scenarios.
- The reward amount of the permutation, as the sum of the reward amounts of the single-validator scenarios.

The author (using a MacBook Pro) was only able to compute permutations for up to 5 validators within a 10-minute period. While a more powerful computer, optimized code, and faster programming language could result in a performance increase of 100x or 1000x, it would only allow managing groups with one or two additional validators.

The second and third categories of actors within the Ethereum validator set can have up to tens or hundreds of thousands of validators under management or ownership. It is necessary to look beyond the return profile of few validators and enable the computation of the return profile of larger validator groups.

4.2 Mathematics to the Rescue: Using Central Limit Theorem

4.2.1 Staking Rewards Distribution per Reward Category

In this section, we calculate the annual distribution of staking rewards for a group of N validators. We consider each staking reward category separately:

- Attestation Rewards
- Block Proposer Rewards
- Sync Committee Rewards

Each category follows a different probabilistic model. The total yearly reward for the group can be obtained by summing the distributions from each category.

The description of each reward type distribution and how to infer a group of N validators' annual reward rate distribution and its parameters can be found in **Appendix**

A. The results show that the reward rate of a group of validators follows a centered normal distribution, and provide the variance as a function of the hyper-parameters.

As an example where the total number of active validators in the network is 1 million, the tip share is 15%, the variance for a group of 10,000 validators is approximately 0.00114 (rounded to 5 decimal points).

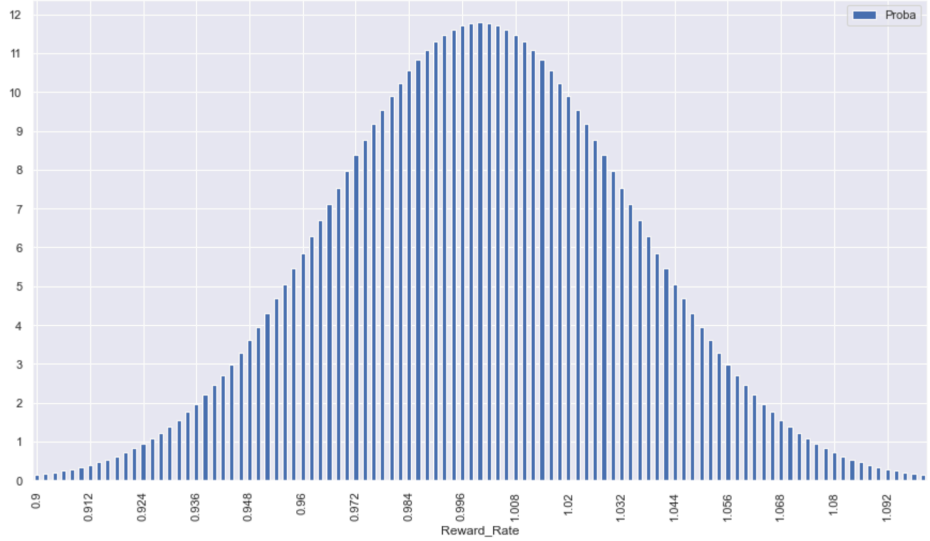

Figure 7: Reward Rate Distribution for a Group of 10,000 Validators and 1M total active validators

We can see below the variance of the staking reward rate as the size of the observed group increases:

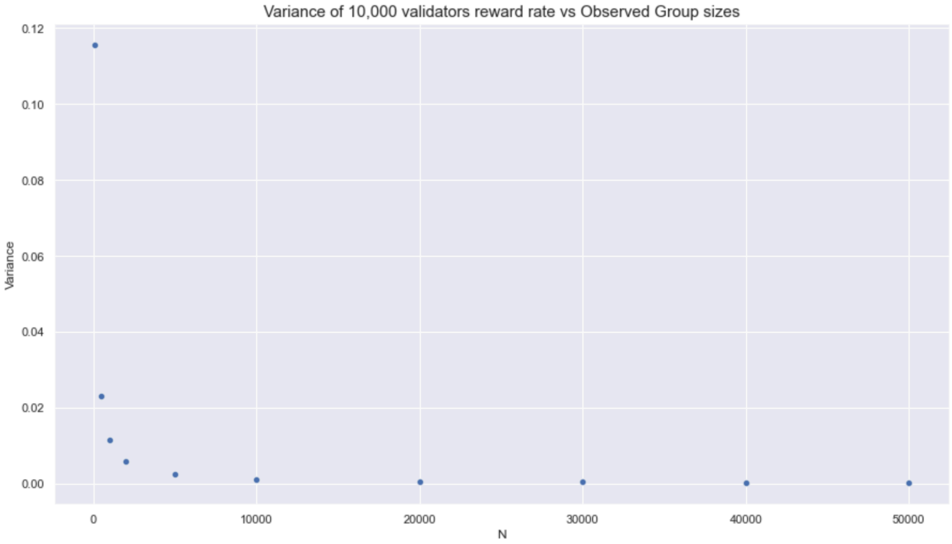

Figure 8: Variance of Staking Reward Rate vs. Observed Group Size

As expected, the variance diminishes exponentially as the size of the observed group increases.

This approximation allows us to model the stochastic nature of staking reward rates and perform risk analysis on the total annual reward rate distribution for a group of validators.

4.3 Staking Reward Rate Confidence Intervals

4.3.1 Confidence Interval Expression

The normal distribution has well-defined confidence intervals, where the $100(1-\alpha)$% confidence interval of the normal distribution is given by

$$\left[\mu - z_{\alpha/2}\sigma,\ \mu + z_{\alpha/2}\sigma\right]$$

where

- μ is the mean.
- σ is the standard deviation.
- $z_{\alpha/2}$ is the critical value from the standard normal distribution corresponding to the desired confidence level.

4.3.2 Calculating the 95% and 99% Confidence Interval

For a 95% confidence level ($\alpha = 0.05$), the critical value $z_{\alpha/2}$ and associated confidence interval are:

$$z_{\alpha/2} = z_{0.025} \approx 1.96$$

For a 99% confidence level ($\alpha = 0.01$), the critical value $z_{\alpha/2}$ and associated confidence interval are:

$$z_{\alpha/2} = z_{0.005} \approx 2.576$$

Taking the example from Section 4.2 with:

- Total number of active validators, $S = 1,000,000$

- Tip share, $TS = 15\%$
- Group size, $N = 10,000$ validators
- Calculated variance of the staking reward rate, $\sigma^2_{\text{SRR}} = 0.00115$

We get the following confidence intervals: We compute the standard deviation:

$$\sigma_{SRR} = \sqrt{0.00114} \approx 0.0338$$

Therefore, the 95% confidence interval for the staking reward rate is:

$$[1 - 1.96 \times 0.0338,\ 1 + 1.96 \times 0.0338] = [0.9338,\ 1.0662]$$

And the 99% confidence interval is:

$$[1 - 2.576 \times 0.0338,\ 1 + 2.576 \times 0.0338] = [0.9129,\ 1.0871]$$

This means that there's a 95% chance that the staking reward rate for the group of 10,000 validators will be between 93.35% and 106.65% of the expected reward rate, and a 99% chance it will be between 91.26% and 108.74

4.4 Impact of a Change in the Hyper Parameters on the Variance

As the staking reward rate distribution is centered, the confidence interval of a group of N validators is only a function of the distribution's variance, itself dependent on two hyper-parameters: the size of Ethereum's active validator set, and the tips share.

4.4.1 Sensitivity of the Variance to the Number of Active Validators

Given a fixed size of the observed sub-group, let's look at the behavior of the variance as a function of the total size of the active validators set:

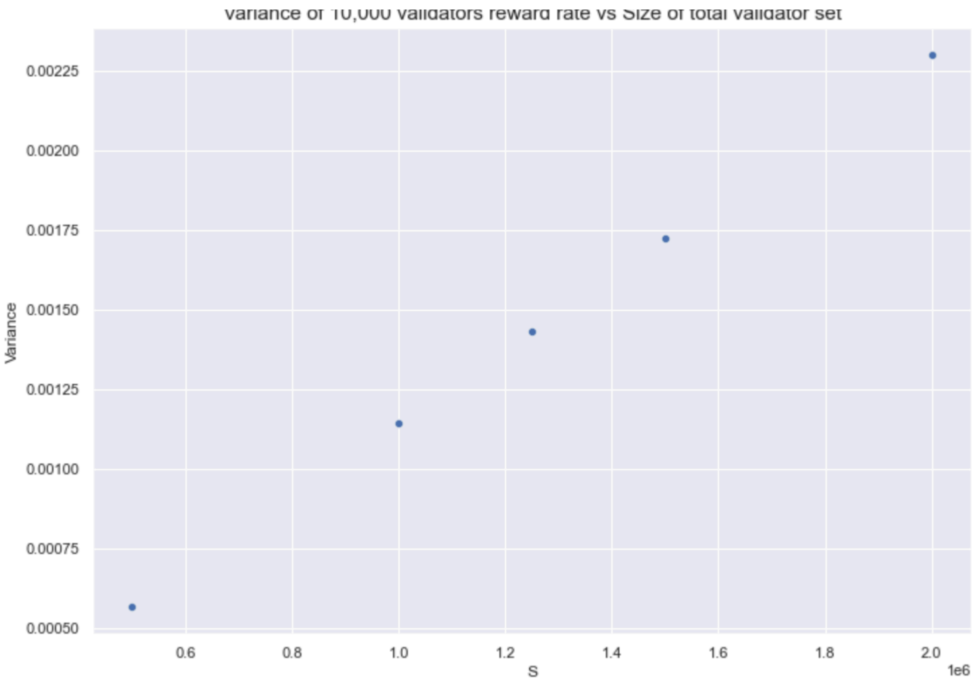

Figure 9: Variance of Staking Reward Rate vs. Validator set size

We can observe a linear increase in the variance as the network's active validator set size increases. This is expected as the share of the sub-group vis-à-vis the overall network size decreases.

Below is the variance as a function of the observed group size for various total active validator set sizes:

Figure 10: Variance of Staking Reward Rate vs. Observed group size for various validators set sizes - y-axis log scale

We observe that, for smaller group sizes, the variance is highly sensitive to changes in the total active validator set size. As group size increases, however, these changes in the validator set have a diminishing effect on variance. This relationship follows an exponential trend, as illustrated by the log-scaled y-axis.

4.4.2 Sensitivity of the Variance to the Tips Share

The tip share represents the portion of execution layer rewards relative to the total consensus layer rewards. Since this reward type is allocated to a single validator per slot, we expect an increase in the tip share to positively influence variance.

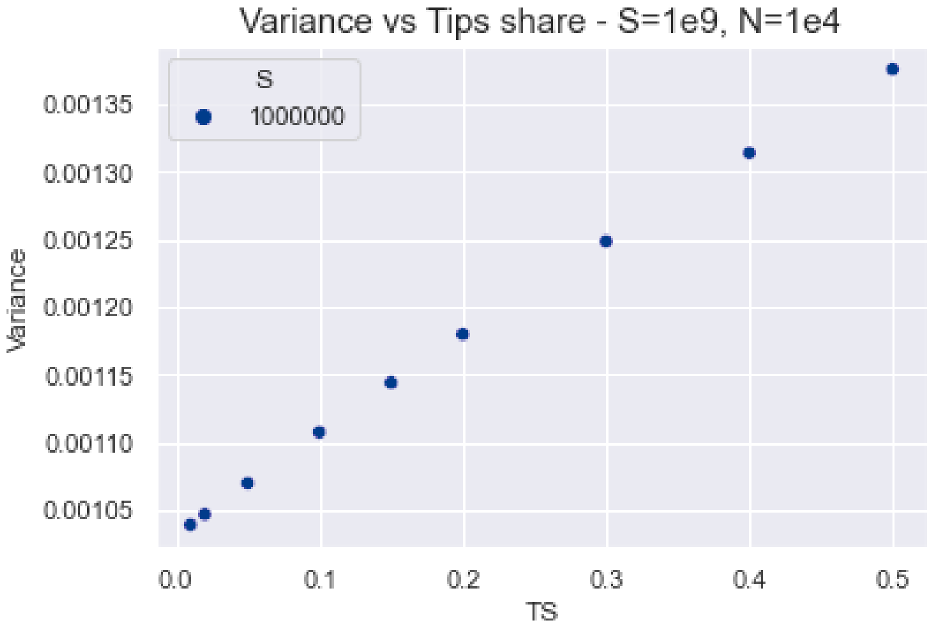

Figure 11: Variance of Staking Reward Rate vs. Tips share

We observe that variance increases linearly with the tip share, at a rate of 0.0007% when the total validator set is 1 million. This means that a 1% increase in the tip share (e.g., from 15% to 16%) leads to a 0.0007 increase in variance, as illustrated in Figure 10.

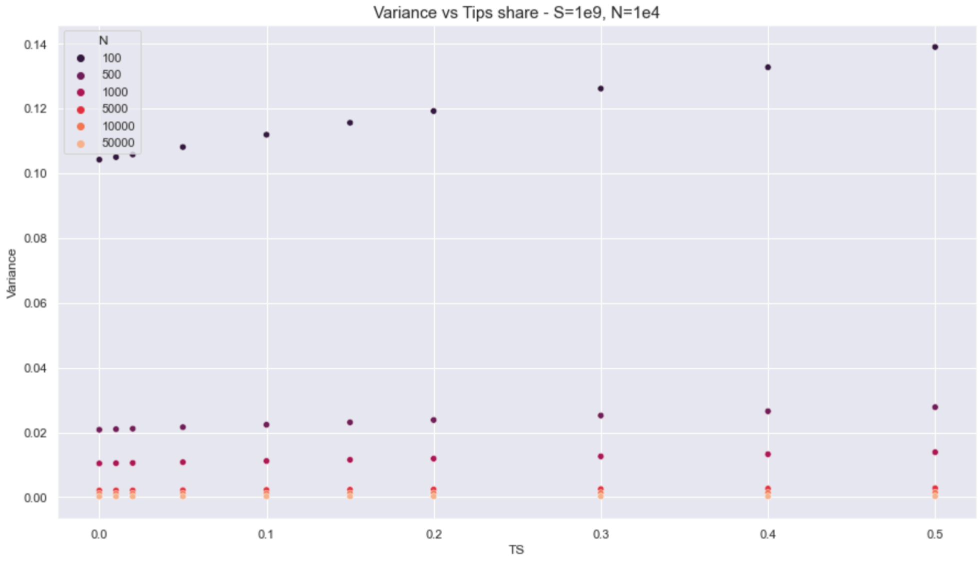

Figure 12: Variance of Staking Reward Rate vs. Tip Share, for Observed Group Sizes

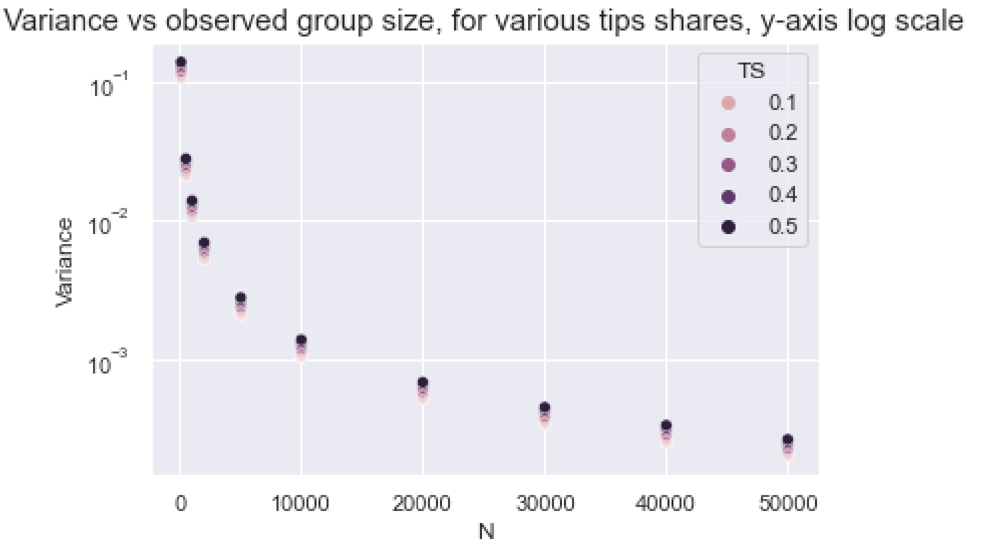

Figure 13: Variance of Staking Reward Rate vs. Observed Group Size, for various Tips share

In figure 11 and figure 12, we see that as the observed group size grows, this linear relationship remains, but the linear coefficient decreases exponentially. Thus, while variance continues to increase with the tip share, the rate of this increase diminishes with larger group sizes.

5 Conclusion

5.1 Summary of Findings

This paper provides a comprehensive probabilistic framework for understanding Ethereum staking from a risk-reward perspective. By dissecting the staking reward system into its constituent components—attestation rewards, block proposer rewards, sync committee rewards, and tips—we were able to model the variability and distribution of staking returns for both individual validators and groups of validators.

The analysis revealed that attestation rewards are deterministic and form the bulk of the staking rewards, ensuring a stable income stream for validators. In contrast, block proposer rewards, sync committee rewards, and tips are subject to randomness, introducing variability into the reward structure. The number of active validators and the level of tips significantly affect both the expected rewards and the variance in returns. Specifically, as the total number of validators increases, per-validator rewards decrease, and smaller validator groups experience higher variance in returns.

Individual validators experience a higher variance in rewards due to the random nature of block proposals and sync committee selections. The probabilistic models show that while the expected returns are positive, the actual returns can fluctuate significantly over time. For larger groups of validators, the Central Limit Theorem allows us to approximate the reward distribution as a normal distribution. This results in a reduced variance, meaning larger validator pools can expect more stable returns compared to individual validators. Moreover, the variance in staking rewards decreases exponentially with the size of the validator group and is linearly affected by changes in the tip share and total number of active validators.

5.2 Implications

The probabilistic framework developed in this paper has several important implications. Validators and staking service providers can use this model to assess the risk associated with staking activities. Understanding the variance and confidence intervals of expected returns enables better financial planning and risk mitigation strategies. The insights into how key parameters affect rewards allow validators to make informed decisions about scaling their operations. For instance, joining or forming larger validator pools can reduce reward variability, offering more predictable income streams. By framing staking rewards in probabilistic and financial terms, stakeholders can integrate staking activities into broader investment portfolios and compare them with traditional financial instruments based on expected returns and associated risks.

5.3 Recommendations for Future Research

While this paper provides a foundational framework, several areas warrant further investigation:

- **Inclusion of Operational Risks:** Future models could incorporate operational risks such as downtime, misconfigurations, and security breaches, which can lead to penalties or slashing, affecting the overall return profile.
- **Stochastic Modeling of Tips:** Tips were considered stable in the main analysis. Including the unpredictable nature of tips, particularly the influence of Maximal Extractable Value (MEV), could lead to a more comprehensive risk assessment.
- **Economic Impact of Slashing Events:** Analyzing the probability and financial impact of slashing events on validators' returns could further improve the accuracy of the expected returns.

5.4 Final Remarks

This paper contributes to a deeper understanding of the financial dynamics inherent in Proof of Stake mechanisms by developing a probabilistic model of Ethereum staking rewards. The blend of deterministic and random reward components creates a unique risk-reward profile that necessitates careful analysis. The framework presented equips validators, staking service providers, and financial professionals with the tools to evaluate staking not just as a technical necessity for network security but as a financial investment with quantifiable risks and returns. As Ethereum and other Proof of Stake networks continue to evolve, such analyses will be crucial in guiding stakeholders toward optimal participation strategies.

Appendix A

In this section, we detail how the variance and the mean of the annual reward rate are calculated.

Block Proposer Rewards

Each block in the blockchain has exactly one proposer. The probability that the proposer is from our group of N validators is:

$$P_{pb} = \frac{N}{S} \tag{1}$$

Given that there are approximately 2,628,000 blocks per year, the number of blocks proposed by the group can be modeled as a binomial distribution. The expected value and variance of this distribution are:

Expected Value:

$$\mu_{pb} = n * P_{pb} = 2{,}628{,}000 \times \frac{N}{S} \tag{2}$$

Variance:

$$\sigma_{pb}^2 = n * P_{pb}(1 - P_{pb}) = 2{,}628{,}000 \times \frac{N}{S}\left(1 - \frac{N}{S}\right) \tag{3}$$

Applying the Central Limit Theorem, we approximate the binomial distribution with a normal distribution. Therefore, the block proposer rewards over a year follow:

$$X_{pb} \sim \mathcal{N}\left(\mu_{pb}, \sigma_{pb}^2\right) \tag{4}$$

Sync Committee Rewards

The probability distribution of the number of validators within a group of N validators chosen to be part of a sync committee can be modeled using a hypergeometric distribution. This distribution models the probability of obtaining a certain number of successes in a sequence of draws from a finite population without replacement.

For each sync committee selection, the number of your validators selected follows:

$$X_{sc} \sim \text{Hypergeometric}(S, N, 512) \tag{5}$$

where:

- S: Total population size (number of active validators)
- N: Number of successes in the population (size of the validator group)
- $n = 512$: Sample size (validators selected for the sync committee)

We bound the size N of the validator group to be below 10% of the total population size S, allowing us to approximate the hypergeometric distribution with a binomial distribution for each sync committee.

Applying the Central Limit Theorem on the sum of the binomial distributions, the sync committee rewards over a year follow:

$$X_{sc} \sim \mathcal{N}\left(\mu_{sc}, \sigma_{sc}^2\right) \tag{6}$$

As there are approximately 146 committee periods in a year, the total number of draws is:

$$n_{\text{total}} = 146 \times 512 \tag{7}$$

Expected Value:

$$\mu_{sc} = n_{\text{total}} \times p_{sc} = 146 \times 512 \times \frac{N}{S} \tag{8}$$

Variance:

$$\sigma_{sc}^2 = n_{\text{total}} \times p_{sc} \times (1 - p_{sc}) = 146 \times 152 \times \frac{N}{S} \times (1 - \frac{N}{S}) \tag{9}$$

To scale this back to one reward per slot, we multiply the distribution by the number of slots in a sync committee period. The new distribution is:

$$X'_{sc} \sim \mathcal{N}\left(\mu'_{sc}, (\sigma_{sc}^2)'\right) \tag{10}$$

with:

$$\mu'_{sc} = \mu_{sc} \times 8{,}192 \tag{11}$$

$$(\sigma_{sc}^2)' = \sigma_{sc}^2 \times (8{,}192)^2 \tag{12}$$

Attestation Rewards

Attestation rewards are deterministic and depend on the participation of the validators. Since these rewards do not have a stochastic component under normal operation, their distribution can be considered constant. The yearly attestation rewards for the group of N validators are given by a fixed amount and can be expressed as a normal distribution with zero variance:

$$X_{att} \sim \mathcal{N}\left(\mu_{\text{att}},\ 0\right) \tag{13}$$

where:

$$\mu_{\text{attest}} = 2,628,000 \times \frac{N}{S} \tag{14}$$

Total Yearly Rewards

To determine the total yearly reward for the group of N validators, we express the total yearly reward as the sum of rewards from the three categories:

$$X_{total} = X_{att} + X_{bp} + X'_{sc} \tag{15}$$

Since each reward category is independent, the sum follows a normal distribution with parameters:

$$\mu_{\text{total}} = \mu_{\text{att}} + \mu_{\text{bp}} + \mu_{\text{sc}} \tag{16}$$

$$\sigma_{\text{total}}^2 = \sigma_{\text{bp}}^2 + (\sigma_{sc}^2)' \tag{17}$$

Note that the attestation reward does not contribute to the variance since it is deterministic.

Staking Reward Rate

Our objective is to describe the return profile of a sub-group's reward rate. We transform the staking reward amount distribution to obtain the staking reward rate distribution:

$$X'_{total} = \frac{X_{total}}{N} \tag{18}$$

Thus:

$$X'_{total} \sim \mathcal{N}\left(\frac{\mu_{total}}{N}, \frac{\sigma^2_{total}}{N^2}\right) \tag{19}$$

Additionally, we scale the distribution to be centered (i.e., set the average to 1) to finally obtain the centered distribution of the Staking Reward Rate:

$$X'_{SRR} \sim \mathcal{N}\left(1, \sigma^2_{SRR}\right) \tag{20}$$

with:

$$\sigma^2_{SRR} = \frac{\sigma^2_{\text{total}}}{\mu^2_{total}} \tag{21}$$

Reusing all previously established equations, we can express the staking reward rate (SRR) variance as:

$$\sigma^2_{\text{SRR}} = \frac{(S-N)(1+TS)(2,628,000 \times TS^2 + 659,628 \times TS) + 4,898,988,463.657}{N(21,683.701 + 2,628,000)^2} \tag{22}$$

Disclaimer & Disclosures

- This information is provided by CF Benchmarks Ltd.
- CF Benchmarks Ltd (‘CF Benchmarks”) is a limited company registered in England and Wales under registered number 11654816 with its registered office at 6th Floor One London Wall, London, United Kingdom, EC2Y 5EB.
- CF Benchmarks is authorised and regulated by the Financial Conduct Authority (FCA) as a Registered Benchmark Administrator (FRN 847100) under the UK Benchmarks Regulation.
- CF Benchmarks is authorised to undertake the following regulated activity ‘Administering a Benchmark.” Administering a Benchmark” is a regulated activity under article 63S of the Financial Services and Markets Act 2000 (Regulated Activities Order) 2001 (SI 2001/544) (RAO), which, in summary, means acting as the administrator of a benchmark as defined in article 3.1(3) of the benchmark regulation.
- CF Benchmarks is **NOT** a registered investment advisor and does **NOT** provide investment, tax, legal, or accounting advice in any geographical locations. You should consult your own financial, tax, legal, and accounting advisors or professionals before engaging in any transaction or making an investment decision.
- All information contained within, including recordings, is for educational and informational purposes **ONLY**. None of the Information constitutes an offer to sell (or a solicitation of an offer to buy) any cryptoassets, security, financial product, or other investment vehicle or any trading strategy. No member of CF Benchmarks nor their respective directors, officers, employees, partners, or licensors provide investment advice, and nothing contained herein or accessible through CF Benchmarks products, including statistical data and industry reports, should be taken as constituting financial or investment advice or a financial promotion.
- Information containing any historical information, data, or analysis should not be taken as an indication or guarantee of any future performance, analysis, forecast, or prediction. Past performance does not guarantee future results. The Information should not be relied on and is not a substitute for the skill, judgment, and experience of the user, its management, employees, advisors, and/or clients when making investment and other business decisions. All Information is impersonal and not tailored to the needs of any person, entity, or group of persons.

- Charts and graphs are provided for illustrative purposes only. Index returns shown may not represent the results of the actual trading of investable assets/securities.
- The Information may contain back-tested data. Back-tested performance is not actual performance but is hypothetical. There are frequently material differences between back-tested performance results and actual results subsequently achieved by any investment strategy. The back-test calculations are based on the same methodology that was in effect when the index was officially launched. However, back-tested data may reflect the application of the index methodology with the benefit of hindsight, and the historic calculations of an index may change based on revisions to the underlying economic data used in the calculation of the index.
- All information and data contained in this publication is obtained by CF Benchmarks from sources believed by it to be accurate and reliable. Because of the possibility of human and mechanical error as well as other factors, however, such information and data is provided 'as is" without warranty of any kind.
- No member of CF Benchmarks nor their respective directors, officers, employees, partners, or licensors make any claim, prediction, warranty, or representation whatsoever, expressly or impliedly, either as to the accuracy, timeliness, completeness, merchantability of any information or of results to be obtained from the use of any CF Benchmarks products.
- No responsibility or liability can be accepted by any member of CF Benchmarks nor their respective directors, officers, employees, partners, or licensors for
    1. any loss or damage in whole or in part caused by, resulting from, or relating to any error (negligent or otherwise) or other circumstance involved in procuring, collecting, compiling, interpreting, analyzing, editing, transcribing, transmitting, communicating, or delivering any such information or data or from use of this document or links to this document; or
    2. any direct, indirect, special, consequential, or incidental damages whatsoever, even if any member of CF Benchmarks is advised in advance of the possibility of such damages, resulting from the use of, or inability to use, such information.
- The user of the Information assumes the entire risk of any use it may make or permit to be made of the Information. **CF BENCHMARKS DOES NOT MAKE ANY EXPRESS OR IMPLIED WARRANTIES OR REPRESENTATIONS WITH RESPECT TO THE INFORMATION (OR THE RESULTS TO BE OBTAINED BY THE USE THEREOF), AND TO THE MAXIMUM EXTENT PERMITTED BY APPLICABLE LAW, IT EXPRESSLY DISCLAIMS ALL IMPLIED WARRANTIES (INCLUDING, WITHOUT LIMITATION, ANY IMPLIED WARRANTIES OF ORIGINALITY, ACCURACY, TIMELINESS, NON-INFRINGEMENT, COMPLETENESS, MERCHANTABILITY, AND FITNESS FOR A PARTICULAR PURPOSE) WITH RESPECT TO ANY OF THE INFORMATION.**

- No part of this information may be reproduced, stored in a retrieval system, or transmitted in any form or by any means—electronic, mechanical, photocopying, recording, or otherwise—without prior written permission of CF Benchmarks Ltd. Any use of or access to products, services, or information of CF Benchmarks Ltd requires a licence from CF Benchmarks Ltd or its authorized licensing agents.
- CF Benchmarks is a member of the Crypto Research group of companies, which is in turn a member of the Payward group of companies. Payward, Inc. is the owner and operator of Staked, a provider of block production and validation nodes for decentralized PoS protocols on behalf of institutional investors. `Staked.us` is a source of input data for the CF Staking Series.