Affine Restaking Risk Engine: Simulating the Distribution of Eigenlayer Restaking Yields

Affine Protocol Team*

May 2024

Eigenlayer is a novel protocol that extends Ethereum's proof of stake security for other services through *restaking*. Currently, the risks and the returns of restaking are in the speculative phase as Eigenlayer smart contracts are under active development and the risk and reward parameters are largely unknown. In our research, we modeled restaking after Ethereum (PoS) staking with higher risk parameters and simulated the returns for the restakers validating 1-50 Actively Validated Services (AVSs). We incorporated both correlated slashing risks among the AVSs and correlation penalties similar to that of Ethereum in our simulations. We show both theoretically and experimentally that the *expected return* and the *standard deviation* for the restakers grow approximately *linearly* with validating more AVSs, suggesting a convergent Sharpe ratio and a limited probability of loss. Within our simulation parameters, the probability of loss remained at 0.73% even when validating up to 50 AVSs. However, potential losses during black swan events scale with the number of AVSs validated. In addition, we built a model where slashing risk of each AVS scales with the number of AVSs an operator is validating, demonstrating a tangible risk of cascading, recursive slashing events. Based on these findings, we propose **Affine Restaking Risk Engine**–an AVS selection framework and an optimal set of restaking strategies: 1) Risk-averse restakers should limit their exposure to black swan events by carefully selecting a smaller set of AVSs (up to 10), 2) restakers with higher risk appetites should aim to maximize expected returns by validating as many AVSs with similar risk profiles as possible. 3) Care should be taken to keep the slashing of each AVSs isolated, to avoid a potential cascading N-slashing event that could be catastropic.

* Affine Protocol team brings together blockchain, quantitative finance, research, and technology experts from YCombinator, Meta, Goldman Sachs, BitGo, MIT, and Harvard.

CONTENTS

1 INTRODUCTION

Eigenlayer [1] is an innovative, novel protocol operating on top of the Ethereum blockchain. It introduced the concept of *restaking*, an extension of delegated proof-of-stake that allows Ethereum stakers to secure multiple blockchains or services simultaneously. This mechanism facilitates a security marketplace where smaller protocols can tap into Ethereum's robust validator set, bootstrapping their operations without incurring the overhead of independent validator networks.

We can broadly categorize the key components of Eigenlayer as follows:

1. **Restakers:** These are the individuals or entities who have staked ETH, either directly or through liquid staking derivatives. Restakers have the option to restake their ETH via Eigenlayer to provide additional security to services, earning rewards and points in exchange.
2. **Actively Validated Services (AVSs):** These are applications including, but not limited to, rollups, data availability layers, oracles, or other specialized services that require security guarantees via consensus. By opting into Eigenlayer, AVSs access Eigenlayer's pool of staked ETH.
3. **Node Operators:** These participants play a pivotal role by choosing which AVSs to support and actively validate. In return for their service, node operators earn yields from the supported AVSs. Crucially, node operators face the risk of slashing (loss of staked ETH) due to misbehavior or failure of validation duties.
4. **Liquid Restaking Tokens (LRTs):** While not directly a part of the Eigenlayer, LRTs streamline the process for users to participate in the Eigenlayer. They function as staking pools for restaking. By depositing ETH into LRTs (for example, eETH [2], ezETH [3], pufETH [4], etc.), users receive representative tokens. These tokens represent a share in a node operator network and earn Eigenlayer points, with potential additional rewards offered by the LRT providers themselves.

1.1 Research Context

It is important to note that restaking yields are currently in a speculative phase. Slashing mechanisms within the Eigenlayer are under development (aimed to be released by Q3 2024), and their implementation will substantially influence actual yields. Moreover, the correlated nature of the risk profiles of the AVSs can further impact potential returns.

This research report presents simulation-based analyses that project potential future scenarios based on various slashing conditions and degrees of interconnection within the restaking ecosystem. The aim is to offer insights into the risk profile associated with restaking as Eigenlayer's slashing implementation and participation evolve.

1.2 Outline

The subsequent sections of this report explore:

- **Risk Factors in Restaking:** A detailed analysis of the operational, technological, and market-related risks inherent to restaking participation.
- **Simulating Restaking Returns modeling after Ethereum:** Explanation of the simulation methodology, assumptions, and results projecting potential yield scenarios based on different slashing parameters.
- **Framework for AVS Selection** A summary of exisiting AVS selection research, and our proposal based on the quantitative methods proposed in this report.
- **Limitations of the Study:** A discussion of factors influencing the simulations and the potential deviations between simulated results and real-world outcomes.

2 RISK FACTORS IN RESTAKING

Currently, restaking rewards on Eigenlayer are effectively risk-free, as slashing conditions have not yet been implemented within its smart contracts. However, this scenario is expected to change, and understanding the potential risks associated with restaking is crucial for node operators and restakers alike. A strong starting point is to examine the established risks and penalties associated with Ethereum staking.

2.1 Ethereum Slashing Conditions

Ethereum validators can be slashed for actions that compromise network integrity, such as:

- Proposing and signing two different blocks for the same slot.
- Attesting to a block that 'surrounds' another one (effectively changing history).
- By 'double voting' by attesting to two candidates for the same block.

When a violation is detected, the validator is slashed. When a validator is slashed, they immediately lose 1/32 of their staked ETH (up to a maximum of 1 ETH), which is permanently removed from circulation. The validator also begins a 36-day removal period during which their remaining stake gradually decreases due to ongoing inactivity penalties. [5]

At the midpoint of the removal period (Day 18), the validator receives an additional penalty that scales with the total amount of staked ETH from all validators slashed within the 36 days surrounding the event. This "correlation penalty" ensures that mass slashing events are punished more harshly, potentially resulting in the loss of a validator's entire stake. [5] In contrast, isolated slashing incidents incur significantly less severe penalties.

In general, Slashing events for Ethereum have been rare and less than 0.04% of the active ETH validators have been slashed. [6]

2.2 Implications for Restaking Risks

Eigenlayer's model has a different risk profile for restakers when compared to traditional Ethereum staking. The slashing conditions for individual AVSs will determined by themselves, resulting in the following implications:

- **Correlated Slashing Risk:** Restakers face amplified risk from unintentional errors, technical failures, or malicious behavior. Since a restaker's staked ETH backs multiple AVSs, a single misstep could trigger slashing penalties from numerous sources at once. This creates a cascading effect where the consequences of an incident can be more severe than in isolated staking scenarios in Ethereum.
- **AVS-Specific Risk:** The nature of Eigenlayer could attract AVSs that cannot secure a dedicated validator set by themselves. These AVSs might be incentivized to take on greater risks, operating under the assumption that they can rely on borrowed capital via restaking. This risk-taking behavior could increase the likelihood of security breaches or technical failures that may trigger slashing events for restakers supporting those AVSs.
- **AVS-Specific Correlation Penalties:** Similar to Ethereum's correlation penalty, Eigenlayer could see implementations where the severity of individual slashing events increases proportionally to the number of other restakers or node operators slashed within a specific timeframe. This interconnectedness means a widespread incident affecting multiple AVSs could result in disproportionately harsh penalties for all involved restakers, even those with no direct fault in the triggering event.

3 SIMULATING RESTAKING RETURNS

To explore the potential impact of slashing risks on restaking returns within Eigenlayer, we conducted simulations based on a simplified model that draws inspiration from Ethereum staking. It's essential to remember that this simulation is intended to illustrate broad trends and highlight the significance of correlated risks under various scenarios. Real-world outcomes may vary as Eigenlayer's slashing implementations and the AVS landscape evolve.

3.1 Simulation Setup and Assumptions

Our simulation operates under the following assumptions:

1. **10x Slashing Risk:** Each AVS has a 0.4% probability per year of slashing a validator, a rate that is 10 times the observed frequency on Ethereum to highlight the potential impact of amplified risk.
2. **Ethereum Identical Penalty:** The slashing penalty is 1 ETH for every 32 ETH staked, mirroring Ethereum's initial slashing penalty.

3. **High Correlation (50%) Risk:** The probability of simultaneous slashing events across AVSs is modeled with a pairwise correlation coefficient (ρ) of 0.5. This indicates a high degree of interconnectedness between AVS security incidents.

4. **Adjusted Correlation Penalty:** For each slashing event, the validator incurs a correlation penalty proportional to the total number of slashings within the surrounding 36-day period. This penalty is capped at 50% of the validator's stake and is triggered only when the network experiences more than 2 slashings during the surrounding 36-day period. This penalty is less harsh than Ethereum's *correlation penalty*, in which case the validator may lose all of their funds. In Eigenlayer, we expect the AVSs to slash more, so an overly harsh correlation penalty will disproportionately affect the validators. Figure 1 shows the curve used for imposing correlation penalty.

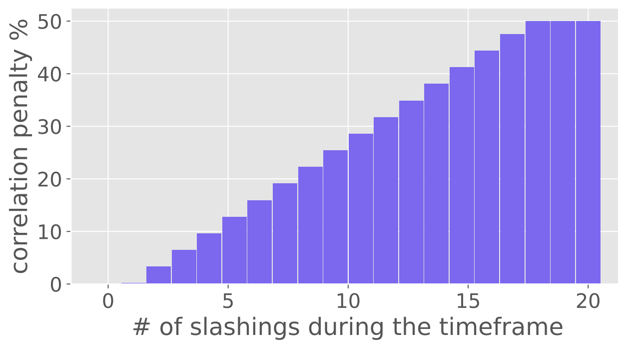

Figure 1: A validator may lose up to 50% of their staked funds due to *Correlation Penalty*.

We simulate scenarios where a single validator supports N different AVSs, with N being 1, 10, 30, or 50. Each AVS offers a fixed return of 0.5% per year. Simulations are run for one year to analyze the final annual yield distribution for the validator. A yield lower than zero indicates a loss.

Two sets of simulations are conducted:

- **No Correlation Penalty:** Validators are subject only to the initial slashing penalty for misbehavior.
- **Correlation Penalty Present:** In addition to the initial penalty, validators face increasing penalties based on the total number of validators slashed within a specific timeframe.

3.2 Technical Details

We simulated restaking returns 3000 times for one validator validating different numbers of AVSs either with or without correlation penalty. We ran the simulation one day at a time for one year (365 days). This is longer than Ethereum's epoch length of 6.4 minutes. We chose a longer epoch to save on computation and because we didn't observe a noticeable difference in the restaking returns with shorter epoch lengths. In every simulation, we assumed the network to have 200 validators in total, and the simulations were implemented in Pytorch [7] with a batch size of 75 simulations. The simulations took approximately 1 hour to run on an Nvidia T4 GPU instance on Google Colab.

4 SIMULATED RESULTS

4.1 Simulations without Correlation Penalty

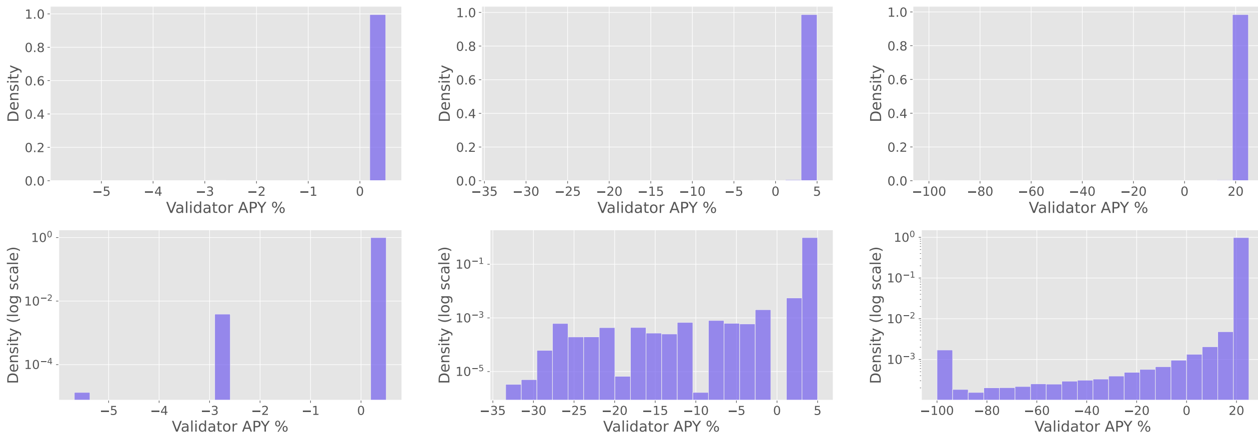

(a) 1 AVS (b) 10 AVSs (c) 50 AVSs

Figure 2: Simulated distribution of Restaking returns for a single validator validating 1, 10, and 50 AVSs (without Correlation Penalty). The top plot shows the results in the linear scale and the bottom plot shows the same results in the log scale (for better visibility). A small number of simulations ended in negative returns for the validator.

In section 8, we show that the expected return from validating N AVSs is N times the expected return from validating 1 AVS. In contrast, the standard deviation of the returns grows only $\sqrt{\rho N^2 + (1-\rho)N} \approx \sqrt{\rho} \cdot N$ times.[1] In Figure 3, we observe that a validator's expected return and the standard deviation of the returns in the simulations match with the theoretical values very closely, verifying the accuracy of our simulations. Table 1 summarizes our findings.

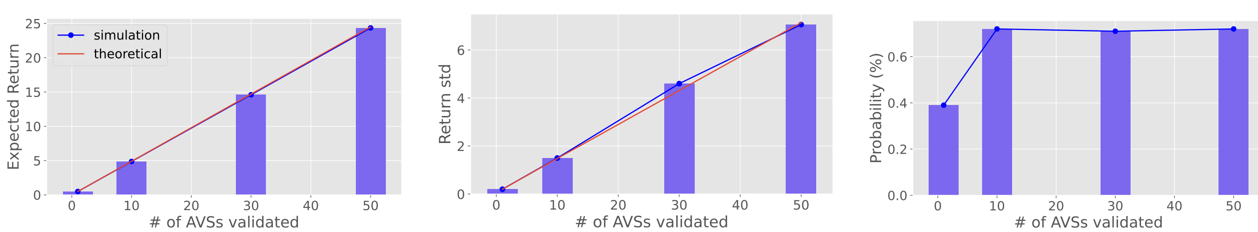

(a) Expected Return (b) STD of Return (c) Risk of loss

Figure 3: Simulated returns from restaking for a single validator validating 1-50 AVSs (without Correlation Penalty). Validating a large number of AVSs does not increase the probability of loss proportionally.

We also showed in section 8 that the risk of loss for validating any number of AVSs has an upper limit lower than 1. With our simulation parameters, this theoretical upper limit of loss is

$$\frac{\rho \cdot s^2 \mathrm{Var}\,(S_1)}{(-s\mathbb{E}\,[S_1] + \mathbb{E}\,[Y_1])^2} = \frac{0.5 \cdot (1/32)^2 \cdot 365 \cdot 0.004/365 \cdot (1 - 0.004/365)}{(-1/32 \cdot 365 \cdot 0.004/365 + 0.005)^2} \approx 8.22\%$$

1 ρ is the pairwise correlation between the slashing events of the AVSs. In our simulations, ρ was 0.5.

In other words, regardless of how many AVSs a validator validates, the risk of making a negative return is never more than 8.22%. The theoretical upper limit of loss is not tight and we observe in Figure 3c that the actual probability of a loss from the simulations is merely 0.72%.

Intuitively, this upper limit can be understood by noticing that both the standard deviation of returns and the expected return grow approximately linearly with N. Consequently, the Sharpe ratio approaches a constant, and thus validating many AVSs is not proportionally riskier in expectation. Mathematically, due to the constant Sharpe ratio, Chebyshev's Inequality gives that the area under 0 must be small and bounded.

Nonetheless, it's important to note in Figure 6 that during the black swan events, a validator's maximum possible loss increases significantly by validating a larger number of AVSs. **When many AVSs slash the staked funds simultaneously, the validator's losses can be substantial.**

Number of AVSs	Expected APY	STD of APY	Risk of loss (%)
1	0.49	0.20	0.39
10	4.87	1.50	0.72
30	14.60	4.60	0.71
50	24.35	7.06	0.72

Table 1: Simulated returns from restaking over one year without correlation penalty.

4.2 Simulations with Correlation Penalty

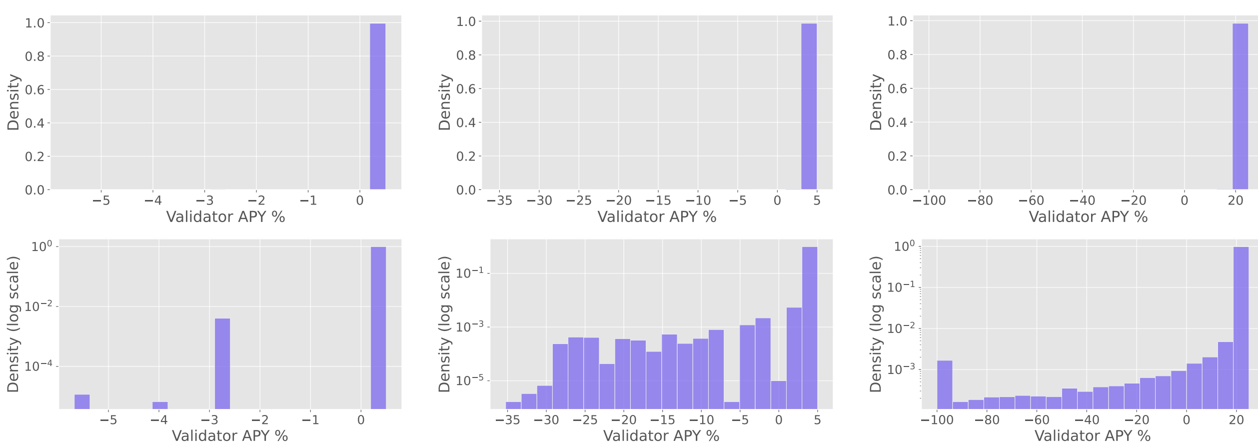

(a) 1 AVS **(b)** 10 AVSs **(c)** 50 AVSs

Figure 4: Simulated distribution of Restaking returns for a single validator validating 1-50 AVSs (with Correlation Penalty). The top plot shows the results in the linear scale and the bottom plot shows the same results in the log scale (for better visibility). A small number of simulations ended in negative returns for the validator.

Under the correlation penalty model, each slashing event incurs an additional penalty based on the total number of validators slashed within a specific timeframe. As illustrated in Figure 1, a small number of isolated slashings do not trigger the correlation penalty. However, in scenarios with many coordinated slashing events, a validator could lose up to 50% of their staked assets.

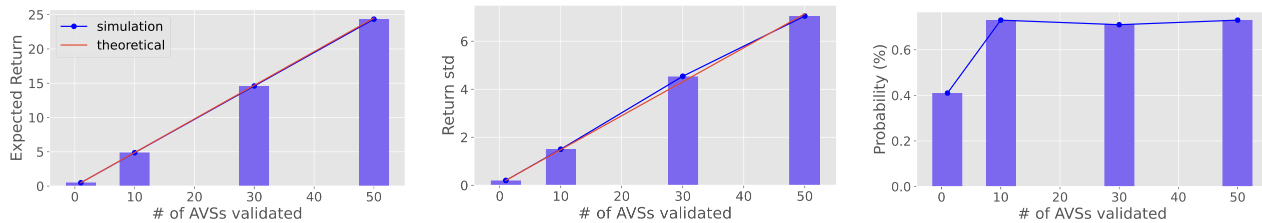

(a) Expected Return (b) STD of Return (c) Risk of loss

Figure 5: Simulated returns from restaking for a single validator validating 1-50 AVSs (with Correlation Penalty). The correlation penalty slightly increases the Risk of loss.

In the simulations, we observe that even with a correlation penalty, a validator's expected return and the standard deviation of the return increase almost linearly with validating more AVSs.

While the risk of loss increased slightly under this model, we did not observe significantly more loss due to correlation penalties. We suspect several reasons behind this. Firstly, the slashings are assumed to be rare events with 0.004 slashings per AVS per year. Secondly, isolated slashings from one validator do not incur correlation penalties. Since the correlation coefficient is 0.5, no significant large slashing event occurred in a short time frame in the simulations. Hence, the contribution of the loss from the correlation penalty was minimal. Table 2 summarizes our findings.

Number of AVSs	Expected Return	STD of Return	Risk of loss (%)
1	0.49	0.20	0.41
10	4.87	1.50	0.73
30	14.60	4.54	0.71
50	24.34	7.05	0.73

Table 2: Simulated returns from restaking over one year with correlation penalty.

4.3 Returns during Black Swans

In traditional finance, black swan events denote extremely rare occurrences that can affect the markets significantly. In this analysis, we define *black swans* as the tails of return distributions with a probability of 0.27% or less (analogous to 3-sigma events in traditional finance). In the Eigenlayer ecosystem, black swans will occur if the validator experiences substantially more slashings than expected. These can happen either due to the validator's mistakes or due to some incident affecting many AVSs within the Eigenlayer ecosystem.

We have executed 3000 simulations with 200 validators. Thus, the resulting dataset contains $200 \times 3000 = 600,000$ validator return values per scenario – a sufficient sample size for black swan loss analysis.

In Table 3, we observe that the black swan losses for validators increase significantly with validating a larger number of AVSs. Furthermore, the correlation penalty doesn't seem to affect the losses as much.

These observations suggest a critical trade-off. While validating a larger number of AVSs does not impact the overall likelihood of loss, it does increase the potential losses during a black swan.

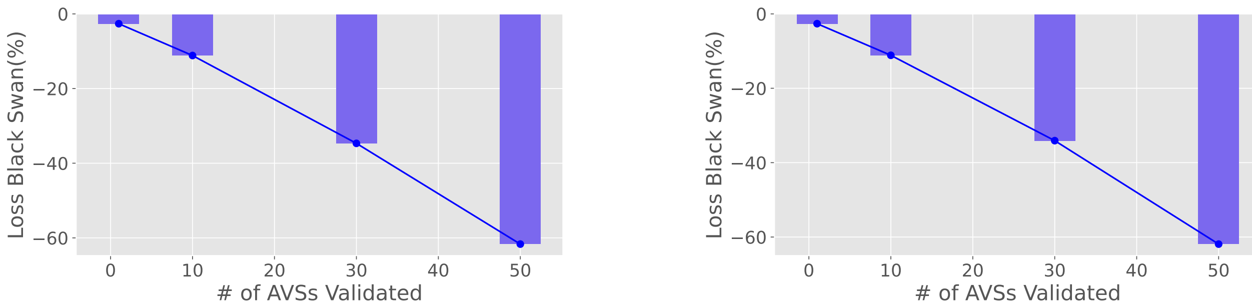

(a) Without Correlation Penalty **(b)** With Correlation Penalty

Figure 6: Simulated loss (%) of staked capital for a validator during a black swan (0.27% probability or less).

Validators with a higher risk tolerance may choose to validate a larger set of AVSs to maximize their potential returns. However, validators seeking to mitigate black swan losses should carefully consider the number of AVSs they validate.

Number of AVSs	Loss (%) without Corr. Penalty	Loss (%) with Corr. Penalty
1	-2.63	-2.63
10	-11.15	-11.13
30	-34.68	-34.08
50	-61.67	-61.90

Table 3: Simulated loss (%) of staked capital for a validator during a black swan (0.27% probability or less).

4.4 Simulation with Ultra High (100x) Slashing Risk AVSs

While we have done this analysis with 10x (0.4% yearly) risk compared to the current Ethereum ecosystem, given the novelty and maturity level of newer AVSs, it is fair to assume that they will have varying levels of slashing risk and in some cases, it can be multiple order or magnitude of risk of Ethererеum slashing risk. So, we present an analysis where we assume 0% (current baseline), 25%, 50%, and 100% of the AVSs being 100x (4.0% yearly) risk compared to the 10x baseline slashing risk.

Both the low-risk group and the high-risk group of AVSs were assumed to have a pairwise group correlation of 50%, and an in-between group correlation of 10%. In other words, multiple high-risk AVSs are likely to slash the validator simultaneously, and similarly, multiple low-risk AVSs are likely to slash the validator simultaneously. However, two groups of AVSs are unlikely to influence slashings of each other.

Running similar simulations for the mix, we obtain the following results.

As expected, as the percentage of 100x, high-risk AVSs goes up in this mix, the expected return goes down–more importantly, the standard deviation (STD) of this return goes up significantly, indicating volatile return periods and higher frequency of black swan events of significant capital loss.

Percentage of High Risk AVSs	10 Total AVSs		30 Total AVSs		50 Total AVSs	
	Expected Return	STD	Expected Return	STD	Expected Return	STD
0%	4.87%	1.50	14.60%	4.54	24.34%	7.05
25%	4.61%	1.87	13.67%	5.67	22.71%	9.24
50%	4.30%	2.70	12.80%	8.02	21.24%	13.34
75%	4.07%	3.46	11.99%	10.67	19.78%	18.04
100%	3.62%	5.12	10.83%	14.55	18.36%	21.66

Table 4: Distribution of simulated return for different percentages of high-risk AVSs

The risk of loss also goes up considerably – from 0.72% to 7.61% when validating 100% high-risk AVSs as shown in Table 5. However, the risk of loss still does not grow with growing number of AVSs.

Percentage of High Risk AVSs	10 Total AVSs Risk of loss (%)	30 Total AVSs Risk of loss (%)	50 Total AVSs Risk of loss (%)
0%	0.72	0.71	0.72
25%	2.76	2.90	3.00
50%	5.12	4.83	4.95
75%	6.12	6.02	6.25
100%	7.43	7.44	7.61

Table 5: Risk of loss % while validating an increasing mix of high-risk AVSs.

This phenomenon can be explained by noticing that the riskiest condition, that is returns for 100% high-risk AVSs, is identically distributed as the returns for 100% low-risk AVSs but with a higher risk parameter. Consequently, due to the same mathematical reasons, the probability of loss will be higher but still bounded and will not grow with validating more AVSs. Because the other rows are less risky than this condition, their probability of loss will also be bounded.

Nonetheless, the risk of loss is still concerning, and if we consider a generally good Sharpe ratio guideline of 1, often followed by traditional finance practitioners, we would likely want to limit the high-risk AVSs in this to 0-50% of the total AVSs.

5 CASCADING AND VARIABLE SLASHING RISK ANALYSIS

One of the key assumptions during simulations and also during the theoretical analysis was that the slashing risk of an AVS does not change when it is validated with other AVSs. A 50% correlation ensures that AVSs are likely to slash in conjunction, however, correlation alone does not make each AVS slash more frequently than it does by itself.

In reality, it is possible that when two AVSs are validated together, they slash more often than they do when validated individually. An example of this could be an optimistic rollup AVS, which also depends on a fraud proof AVS–it is possible that when a fraud proof AVS faces a majority slashing, the rollup AVS could also slash, causing a cascading slashing event. Under this risk model, the

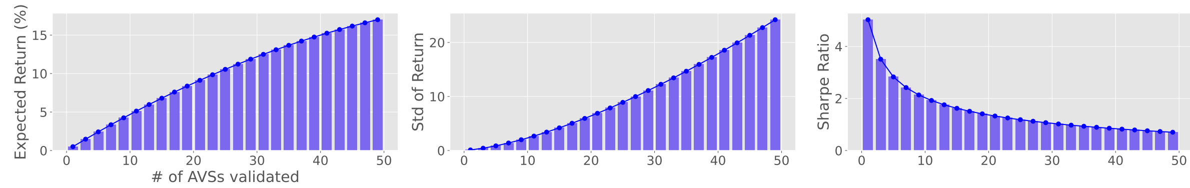

(a) $p_s(N) = 0.25 \cdot N \cdot p_s(1)$

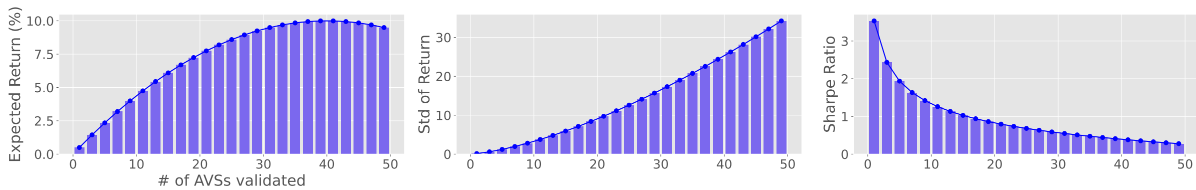

(b) $p_s(N) = 0.50 \cdot N \cdot p_s(1)$

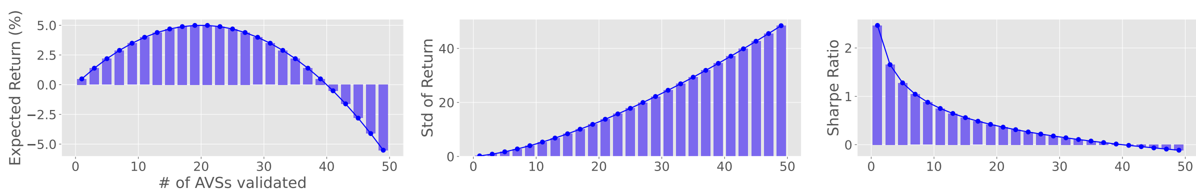

(c) $p_s(N) = 1.0 \cdot N \cdot p_s(1)$

Figure 7: Theoretical expected return, standard deviation of return, and Sharpe ratio for validating N AVSs when the slashing risk of each AVS grows linearly when validated with other AVSs. All other parameters were assumed to be the same as previous simulations.

slashing probability for *each AVS* grows when it is validated with more AVSs. **If this is the case, then neither the expected return nor the standard deviation of the return will grow linearly.** We show this below with a theoretical proof:

Suppose a validator validates N AVSs together. For each AVS, suppose the slashing probability is $p_s(N)$ per year, which is a monotonic function in N, and the pairwise correlation between the AVSs is ρ. Assuming the same notations for $(Y_n)_{n=1}^N, (a_n)_{n=1}^N, (S_n)_{n=1}^N$ as in the section 8, we can observe that the slashings, S_n-s, are binomially distributed. Therefore, the expected return from restaking N AVSs is

$$\mathbb{E}\left[\sum_{n=1}^{N} -s \cdot S_n + Y_n\right] = N \cdot (-s\mathbb{E}[S_1] + \mathbb{E}[Y_1]) \tag{1}$$

Similarly, the variance of restaking returns is

$$\mathrm{Var}\left(\sum_{n=1}^{N} -s \cdot S_n + Y_n\right) \approx s^2 \mathrm{Var}(S_1) \cdot \left[\rho N^2 + (1-\rho)N\right] \tag{2}$$

However, note that

$$\mathbb{E}[S_1] = M \cdot p_s(N)$$
$$\mathrm{Var}(S_1) = M \cdot p_s(N) \cdot (1 - p_s(N))$$

Both of these terms now depend on N and hence neither the expected return nor the standard deviation would grow linearly under these assumptions. In Figure 7, we show some plots showing the growth of these terms for different $p_s(N)$ functions.

6 AFFINE RESTAKING RISK ENGINE: AVS SELECTION FRAMEWORK

Given restaking is a nascent primitive and we have virtually no slashing data beyond Ethereum slashing, there has not been many studies around restaking. However, based on our simulation and a few existing studies, we present **Affine Restaking Risk Engine**, a framework for managing risks for restaking.

6.1 Existing Studies: Gauntlet's Framework

Recently, Gauntlet [8] published a framework for selecting AVSs. Under this framework, AVS selection is seen as an optimization problem at each rebalancing step t.

Suppose $\mathcal{AVS}$ is the list of live AVSs and $\mathcal{O}$ corresponds to the validators associated with an LRT. According to Gauntlet's suggestion, the LRT should optimize the following cost function, which depends on both the returns from the AVSs and the slashing conditions:

$$\max_{\mathcal{AVS}} g\left(\sum_{o\in\mathcal{O}}\sum_{a\in\mathcal{AVS}} w_f f(c_{a,o},t) - w_s s(c_{a,o},t)\right)$$

subject to some boundary conditions. While analyzing the formula for the full extent is tangential for our research, we will explain some of the key terms:

- $f(\cdot)$ is a function of the yield from the AVS a and validator o. Similarly $s(\cdot)$ is a function of the slashing penalties.
- w_f and w_s are importance given to the yield and slashing loss. From a mathematical perspective, these weights could be part of the definition of f and s, and explicit weight terms were unnecessary.
- $c_{a,o}$ represents the weight of the portfolio of validator o used to stake AVS a.
- $g(\cdot)$ is a function of the overall return.

One important point to note is that we ran our simulations for one validator and not from the perspective of an LRT. Consequently, in our optimization problem, the double summation above becomes a single summation over the $\mathcal{AVS}$.

6.2 Affine's Framework to Optimize Restaking Returns

As of writing this document, slashing penalties have not been implemented in the Eigenlayer smart contract. Consequently, the Eigenlayer points are risk-free at this moment. However, once the slashing conditions are implemented, the restakers need to be careful about AVS selection and monitoring to optimize their returns. Based on our simulations and Gauntlet's framework, we suggest the following optimization functions for investors with different risk appetites.

6.2.1 Investors with High Risk Tolerance

We observed from the simulations that the risk of the loss does not grow with validating more AVSs while the expected return does. Investors who think the risk of loss is acceptable should choose to validate as many AVSs as possible. In reality, not all AVSs will have the exact same slashing probability or the same pairwise correlation. So, the investor should also have a threshold for maximum slashing penalty (S) from one AVS, minimum yield (Y), and maximum allowed pairwise correlation (ρ_{max}). Mathematically, this strategy can be written as:

$$\max_{\mathcal{AVS}} \sum_{a \in \mathcal{AVS}} \mathbb{E}\left[w_f f(c_a, t) - w_s s(c_a, t)\right]$$

subject to

$$\begin{aligned} w_s s(c_a, t) &\leqslant S \\ w_f f(c_a, t) &\geqslant Y \\ \mathrm{Corr}(a_i, a_j) \leqslant \rho_{max} \quad \forall a_i &\neq a_j \in \mathcal{AVS} \end{aligned}$$

6.2.2 Investors Who Want Principal Protection

In our simulations, we also found that the risk of maximum loss during black swan events grows with validating more AVSs. The investors who want principal protection and minimum loss during black swan events should only validate a few, carefully selected AVSs. In particular, they should *minimize* the probability of a maximum loss $-1 < -L < 0$. Like before they should still have thresholds for maximum slashing penalty (S) from one AVS, minimum yield (Y), and maximum allowed pairwise correlation (ρ_{max}).

Mathematically, their optimization function can be written as:

$$1 - \min_{\mathcal{AVS}} \sum_{a \in \mathcal{AVS}} \int_{-\infty}^{-L} w_f f(c_a, t) - w_s s(c_a, t) \; d\mu$$

Where μ is the probability measure over the return from restaking at time step t. The boundary conditions would be

$$\begin{aligned} w_s s(c_a, t) &\leqslant S \\ w_f f(c_a, t) &\geqslant Y \\ \mathrm{Corr}(a_i, a_j) \leqslant \rho_{max} \quad \forall a_i &\neq a_j \in \mathcal{AVS} \\ -1 < -L &< 0 \end{aligned}$$

6.3 General Guidelines to Optimize Restaking Returns

Here are some steps we recommend to all restakers to optimize their returns:

- Conduct in-depth analysis of an AVS's operating principles, code quality, audit history, and track record of security incidents before committing to validate them.
- Understand that the slashing conditions are determined by the AVSs. Scrutinizing the specific slashing rules and penalties of each AVS before validating them.
- Select validators that demonstrate consistent availability, active participation, and a minimal history of being slashed.
- For a chosen level of risk tolerance, identify AVS combinations that maximize potential returns while minimizing the probability of loss.
- Be aware that the probability of loss initially increases with the number of validated AVSs, then stabilizes to be approximately 0.73%.
- There could be heuristics-based slashing conditions, to limit the amount of possible loss in specific time windows and circuit breakers to potentially even halt the network and use social consensus to avoid catastrophic losses, while being mindful of the general liveliness of the network.
- Restakers could also look into insurance products that provide principal protection during the black swan events in exchange for some of the Eigenlayer yields.
- Recognize that restaking, in extreme circumstances, can lead to complete loss of staked assets. Implement risk mitigation strategies to limit exposure and protect against the impact of black swan events. In particular, the novel $bEIGEN token design [9] to use the potential Intersubjectivity of such black-swan event, could protect the stakers from massive losses.

7 LIMITATIONS OF THE SIMULATIONS

- In the simulation, each AVS is assumed to provide a return of 0.5% per year. In reality, the AVS returns will have variability, and the returns could be either more or less for many AVSs. As a result, the mean APY for a validator will also change.
- The slashing risk for each AVS was assumed to be the same, but in reality, the slashing risk will be different for each AVS.
- Currently the slashing conditions are not implemented in Eigenlayer. We have run the simulation assuming the slashing penalties will be similar to Ethereum. However, if that is not the case, these simulations cannot provide any insights.
- When a non-malicious Ethereum validator gets slashed, it's often due to a software bug, network error, unsynchronized clock, etc. In those cases, the validator re-joins the network after some inactivity period. For simplicity, the simulation does not take that into account and the validator is assumed to re-join the next day.
- The pairwise correlation factor was fixed to be 0.5. However, in practice, different AVSs will have different levels of correlation. Also, the correlation will vary over time.
- In the simulation, the epoch length was one day instead of 32 blocks. This reduced the amount of computation by 225x but made the simulation slightly inaccurate.

8 APPENDIX: THEORETICAL ANALYSIS OF RESTAKING RETURNS

In this section, we do a theoretical estimation of the number of slashings a validator may experience by validating a growing number of AVSs. We assume there is no correlation penalty to simplify the analysis.

Suppose a validator is validating N AVSs $(a_n)_{n=1}^N$ for M epoch with one unit staked capital. In every epoch, each AVS has an i.i.d. slashing probability of p_s, gives a yield $y\%$ on unit staked capital, and every two AVSs a_i and a_j have pairwise Pearson correlation of ρ. Lastly, assume that for AVS n, the validator gets slashed S_n times in total during the M epochs and gets a total yield Y_n on the unit capital staked.

8.1 Expected Return from Restaking

First, note that the expected yield of each AVS is identical, hence

$$\mathbb{E}\left[\sum_{n=1}^{N} Y_n\right] = N \cdot \mathbb{E}\left[Y_1\right]$$

On the other hand, the total number of slashings the validator can expect is the following:

$$\mathbb{E}\left[\sum_{n=1}^{N} S_n\right] = \sum_{n=1}^{N} \mathbb{E}\left[S_n\right] = N \cdot \mathbb{E}\left[S_1\right]$$

Suppose the maximum immediate slashing penalty is s. Then, the restaking returns can be written as

$$-N \cdot s\mathbb{E}\left[S_1\right] + N \cdot \mathbb{E}\left[Y_1\right] = N \cdot \left(-s\mathbb{E}\left[S_1\right] + \mathbb{E}\left[Y_1\right]\right)$$

Thus, the *expected return* from validating N AVSs is N times the return of validating one AVS.

8.2 Variance of the Restaking Returns

The variance of Y_1 is negligible compared to the variance of S_1 since the expected number of slashings by AVS a_1 over one year is very small. Thus, the variance of the returns will be dictated by the variance of the S_1.

Next, we shall compute the variance of the total number of slashings. First, note that S_n-s are binomially distributed. Therefore,

$$\begin{aligned}
\mathrm{Var}\left(\sum_{n=1}^{N} S_n\right) &= \sum_{n=1}^{N} \mathrm{Var}\left(S_n\right) + 2\sum_{n=1}^{N}\sum_{k=n+1}^{N} \mathrm{Cov}\left(S_n, S_k\right) \\
&= MNp_s(1-p_s) + 2\sum_{n=1}^{N}\sum_{k=n+1}^{N} \rho \cdot Mp_s(1-p_s) \\
&= MNp_s(1-p_s) + N(N-1)\rho \cdot Mp_s(1-p_s) \\
&= Mp_s(1-p_s)\left[\rho N^2 + (1-\rho)N\right] \\
&= \mathrm{Var}\left(S_1\right) \cdot \left[\rho N^2 + (1-\rho)N\right]
\end{aligned}$$

This is the variance of the *total number of slashings*. With the slashing penalty factor s, the variance of the *total slashing penalty* will have an s^2 multiplier to the above formula.

Also, for $\rho > 0$, this value grows quadratically in N. Therefore, for N large enough and $\rho > 0$, the variance of the return while validating N AVSs is approximately ρN^2 times the variance of the return for validating one AVS. Therefore, the standard deviation should be $\sqrt{\rho}N$ times the standard deviation of validating one AVS.

8.3 Upper Bound on the Risk of Loss

Chebyshev's inequality: For a random variable X with finite expectation μ and finite variance σ^2, and for any real number $k > 0$, the Chebyshev's inequality gives

$$P((X-\mu)^2 \geqslant k^2\sigma^2) \leqslant \frac{1}{k^2}$$

We know that for the total restaking returns, we have $\mu = N \cdot (-s\mathbb{E}\left[S_1\right] + \mathbb{E}\left[Y_1\right])$, and $\sigma^2 = s^2\mathrm{Var}\left(S_1\right)(\rho N^2 + (1-\rho)N)$.

Assuming $\mu > 0$ and applying Chebyshev's inequality for $k = \frac{\mu}{\sigma}$ gives us,

$$\begin{aligned}
P((X-\mu)^2 \geqslant \mu^2) &\leqslant \frac{\sigma^2}{\mu^2} \\
&= \frac{s^2\mathrm{Var}\left(S_1\right)(\rho N^2 + (1-\rho)N)}{N^2 \cdot (-s\mathbb{E}\left[S_1\right] + \mathbb{E}\left[Y_1\right])^2} \\
&= \frac{s^2\mathrm{Var}\left(S_1\right)}{(-s\mathbb{E}\left[S_1\right] + \mathbb{E}\left[Y_1\right])^2} \cdot \left[\rho + \frac{1-\rho}{N}\right]
\end{aligned}$$

Using the last inequality, we can obtain an upper bound on the probability of loss:

$$\begin{aligned}
P(X \leqslant 0) &= P(X-\mu \leqslant -\mu) \\
&\leqslant P(X-\mu \leqslant -\mu) + P(X-\mu \geqslant \mu) \\
&= P(|X-\mu| \geqslant \mu) \\
&\leqslant \frac{s^2\mathrm{Var}\left(S_1\right)}{(-s\mathbb{E}\left[S_1\right] + \mathbb{E}\left[Y_1\right])^2} \cdot \left[\rho + \frac{1-\rho}{N}\right]
\end{aligned}$$

Now note that as $N \to \infty$, the right-hand side converges to a finite value:

$$\lim_{N\to\infty} \frac{s^2 \mathrm{Var}(S_1)}{(-s\mathbb{E}[S_1] + \mathbb{E}[Y_1])^2} \cdot \left[\rho + \frac{1-\rho}{N}\right] = \frac{\rho \cdot s^2 \mathrm{Var}(S_1)}{(-s\mathbb{E}[S_1] + \mathbb{E}[Y_1])^2} \tag{3}$$

Now it suffices to show that the right-hand side is less than 1.

Suppose $L = \frac{\mathbb{E}[Y_1]}{s\sqrt{\mathbb{E}[S_1]}}$. We will show that if $L > 2$, then the right-hand side is less than $\rho/(L-1)^2 < 1$.

First note that from the properties of the binomial distribution, we have $\mathrm{Var}(S_1) = \mathbb{E}[S_1](1-p_s)$. Additionally, $0 < \mathbb{E}[S_1] \ll 1$ by assumption.

Therefore,

$$\begin{aligned} \mathbb{E}[Y_1] &= s\sqrt{\mathbb{E}[S_1]} + (L-1)s\sqrt{\mathbb{E}[S_1]} \\ &> s\mathbb{E}[S_1] + (L-1)s\sqrt{\mathbb{E}[S_1]} \\ &\geqslant s\mathbb{E}[S_1] + (L-1)s\sqrt{1-p_s}\sqrt{\mathbb{E}[S_1]} \end{aligned}$$

This implies

$$\begin{aligned} (\mathbb{E}[Y_1] - s\mathbb{E}[S_1])^2 &> (L-1)^2 s^2 (1-p_s)\mathbb{E}[S_1] \\ &= (L-1)^2 s^2 \mathrm{Var}(S_1) \end{aligned}$$

Subsequently,

$$\frac{\rho \cdot s^2 \mathrm{Var}(S_1)}{(-s\mathbb{E}[S_1] + \mathbb{E}[Y_1])^2} = \rho \cdot \frac{s^2 \mathrm{Var}(S_1)}{(-s\mathbb{E}[S_1] + \mathbb{E}[Y_1])^2} < \rho \cdot \frac{1}{(L-1)^2}$$

For sufficiently large $\mathbb{E}[Y_1]$, the condition $L > 2$ is true and thus this upper bound applies.

Note that the values of $\mathbb{E}[Y_1]$ that make $L > 2$ are quite realistic and are satisfied by our simulation parameters. Furthermore, the upper bound from Chebyshev is also very loose. In particular, for our simulation parameters, the upper bound in Equation 3 is approximately 8.22%, while the actual probability of loss from the simulation was around 0.72%. This theoretical proof shows that an upper bound on the risk of the loss exists, it does not depend on N and is less than 1.

REFERENCES

[1] EigenLayer Team. Eigenlayer: The restaking collective. https://docs.eigenlayer.xyz/overview/whitepaper, 2024.

[2] ETH Re-staking — ether.fi — etherfi.gitbook.io. https://etherfi.gitbook.io/etherfi/ether.fi-whitepaper/eth-re-staking, 2024.

[3] Lucas Kozinski, James Poole, and Kratik Lodha. Renzo protocol. https://www.renzoprotocol.com/, 2024.

[4] Amir Forouzani. Puffer Finance. https://www.puffer.fi/, 2024.

[5] The Ethereum Foundation. Proof-of-stake rewards and penalties. https://ethereum.org/en/developers/docs/consensus-mechanisms/pos/rewards-and-penalties/, 2024.

[6] Matthieu Saint Olive and Simran Jagdev. Understanding Slashing in Ethereum Staking: Its Importance & Consequences — Consensys.io. https://consensys.io/blog/understanding-slashing-in-ethereum-staking-its-importance-and-consequences, 2024.

[7] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Köpf, Edward Z. Yang, Zach DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner, Lu Fang, Junjie Bai, and Soumith Chintala. Pytorch: An imperative style, high-performance deep learning library. *CoRR*, abs/1912.01703, 2019. URL http://arxiv.org/abs/1912.01703.

[8] Gauntlet. AVS Selection Framework - Gauntlet. https://www.gauntlet.xyz/resources/avs-selection-framework, 2024.

[9] EigenLayer Team. Eigen: The universal intersubjective work token, towards the open verifiable digital commons. https://docs.eigenlayer.xyz/overview/whitepaper, 2024.