Evmos

Agent-Based Model Analysis

Examining the effects of the redesigned Token Economy on the Evmos Chain.

POL

pol.finance

Evmos Token Economics

(Agent Based Model)

POL Finance[1],
Federico Kunze Küllmer[2], and
Daniel Burckhardt[3]

[1] *dipa@pol.finance*, *agustin@pol.finance*
[2] *federico@evmos.org*
[3] *daniel@evmos.org*

June 13, 2024

Abstract

This report examines the implementation and impact of the redesigned token economy for the Evmos chain, utilizing a custom agent-based model to assess performance under varied conditions. By simulating interactions among key agents—SimEngine, the Evmos chain, and validators—across scenarios of fluctuating demand and price volatility, the effects of changes in parameters such as chain demand, price volatility, and emission structure are highlighted. Results demonstrate a significant reduction in inflation, from 30% to 4%, and improvements in decentralization, as evidenced for example by a Gini coefficient reduction from 0.7 to 0.1. These findings validate the effectiveness of the new design in promoting a more equitable and sustainable ecosystem, confirming that strategic adjustments to token issuance and reward distribution can lead to a more decentralized Evmos chain economy with controlled inflation.

Disclaimer

The following disclaimer is intended to clarify the purpose and limitations of the report produced by POL Finance and the Evmos Core Team for the study of the Evmos protocol. Please read this disclaimer carefully before reviewing the report.

No Financial Advice: The information provided in this report is for informational purposes only and should not be construed as financial advice, investment recommendations, or an invitation to engage in any buying, selling, or holding of any specific asset. POL Finance, Evmos Foundation and the authors do not provide financial advice, and the content of this report should not be considered as a sole or definitive source for making decisions or relying upon.

Limitation of Liability: POL Finance, Evmos Foundation and the authors assume no responsibility or liability for any financial losses, direct or indirect, incurred from using or interpreting the information provided in this report. Readers are responsible for their own investment decisions and should undertake their due diligence.

Information Obsolescence: It is crucial to acknowledge that the information provided in this report may become obsolete due to dynamic market conditions, regulatory shifts, technological advancements, or other factors. Readers should be aware that the report's content is subject to change, and it is advisable to conduct timely research and analysis before making any investment decisions.

Contents

Chapter 1

Introduction

Agent-Based Simulations (ABS) are computational models that simulate the behavior and interactions of individual agents to understand the collective behavior of a system. In the context of stress tests, ABS can be a valuable tool for assessing the resilience and performance of various systems. It is the primary tool we employ to conduct simulation-based stress tests on Evmos's new tokenomic design.

ABS has wide-ranging applications in stress testing. For example, in the cryptocurrency industry, ABS can be used to simulate extreme market conditions and evaluate the impact on Decentralized Finance (DeFi) protocols, such as analyzing the effects of a significant drop in token prices on liquidity pools and user funds. In the Decentralized Exchange (DEX) sector, ABS can be employed to assess the resilience of trading platforms against flash crashes or large trading volumes, aiding in risk management and protocol optimization. ABS is also utilized in simulating yield farming behavior and market dynamics, allowing for the evaluation of different farming strategies and market conditions.

Fine-tuning ABS models can be a complex task. It involves calibrating the behavior and characteristics of individual agents, as well as adjusting the system-level parameters to achieve desired outcomes. The challenge lies in striking a balance between accuracy and computational efficiency, as ABS models can become computationally intensive and time-consuming. However, with careful design and iterative testing, ABS models can be fine-tuned to provide valuable insights and replicate real-world scenarios effectively.

1.1 Environment

Having a modular, generic ABS framework offers several advantages such as flexibility and adaptability for easy customization and adaptation to different scenarios and systems; reusability for utilizing it across different domains and applications, enabling the reuse of code and models; and scalability for handling simulations of varying complexity and scale, making it suitable for simulating complex real-world scenarios.

The POL platform enables the execution of realistic and precise simulations by providing a modular interface. Moreover, agent behavior was modeled using a Python domain-specific language which promotes quick prototyping and iterative experimentation. Therefore, the POL platform allows analysts to focus on exploring different scenarios while abstracting implementation details and facilitating rapid prototyping.

1.2 Impact Analysis

This final report details the impact analysis of restructuring the native token economy on the Evmos chain, using an advanced custom agent-based model. Focusing on the practical implementation of our proposed design, the study employs a dynamic simulation approach to evaluate the robustness of the new tokenomic design under a variety of market scenarios and parameter configurations. Incorporating key agents such as SimEngine, the Evmos chain, and validators, the model simulates complex interactions and decision-making based on behavior profiles, market trends, and risk aversion under stochastic conditions and fluctuating demands.

Through simulations varying key parameters like chain demand, token price volatility, emission weights, and validator power distribution, the analysis reveals how different configurations affect inflation, decentralization, and other key performance indicators (KPIs). Scenarios of high demand, significant price trend factors, and adjustments in emission weights associated with the staked ratio and base fee, along with optimal initial staked ratios, emerged as critical factors influencing the success of the new design.

The outcomes indicate a trend towards inflation reduction, from 30% to 5% in nearly all tested scenarios, and even to below 3.5% in some cases, highlighting the proposed design's effectiveness in emission control. Furthermore, a notable improvement in decentralization is observed, with a Gini coefficient drop from 0.7 to approximately 0.1, and an increase in participation from small and medium agents, confirmed by a shift in the stake distribution towards lower values. These findings underscore the revised design's capacity to foster a more equitable and sustainable economy on the Evmos chain.

This report not only validates the viability of the proposed design through comprehensive analysis and simulations but also establishes a framework for dynamic adaptation and resilience of the Evmos ecosystem against market fluctuations and changes in agent behavior. The implementation of the new design promises to advance towards controlled inflation, increased decentralization, and healthy, sustainable growth of the Evmos ecosystem.

Chapter 2

ABM: Key Variables and Influences

2.1 Model Objectives

The model will assist us in testing the impact of implementing tokenomics v2 in relation to the objectives we aim to achieve. Based on the Evmos' token economic redesign report, the network should:

i. Reduce the annual emission rate. This rate should be such that:

   - It does not exceed an annual inflation limit of 5%.

ii. Dynamically adjust the emission rate. Issuance must:

   - Decrease when the staked ratio is above the optimal value $s > s^*$, in order to discourage the newly issued to be used for staking purposes.
   - Increase when the staked ratio is below the optimal value $s < s^*$, in order to encourage the newly issued to be used for staking purposes.
   - Be inversely proportional to the base fee factor, in order to behave contrary to network demand.

iii. Achieve higher decentralization measured in terms of the total staked of each validator.

2.2 Description of Agents

- **SimEngine**: This agent models the entire system, acting as the model itself. It is responsible for:
  - Implementing the agents.
  - Implementing the system's step (agents' and model's updates).
  - Adjusting exogenous parameters in each step, such as EVMOS price, staked ratio, base fee, black swans (e.g., extreme volatilities, extreme values in BaseFee or staked ratio, random inputs in stochastic processes, etc), and agent behavior profiles.
  - Collecting data on the KPIs of interest.

- **Evmos Chain**: This agent replicates the Evmos chain, mimicking the updated attributes associated with all agent's actions. It adjusts attributes associated with agent's actions like stake/unstake and delegate/re-delegate, and with self actions such as mint/burn, transfer rewards, transfer/burn base fee.
- **Validator**: This agent, which encompasses the role of both a validator and also a delegator, validates transactions, aiming to replicate a behavior profile based on:
  - The utility function of rewards received for staking minus the cost of staking.
  - A trend following strategy that represents the validator's appreciation of value towards the EVMOS token.
  - A risk aversion related to the volatility of EVMOS' price.

  Its actions include stake/unstake and delegate/re-delegate.

2.3 Hypothesis

2.3.1 Temporal Horizon (Step)

We assume that a **daily** step interval is the optimal temporal horizon in terms of computational efficiency and obtaining reliable and dynamic system data.

2.3.2 Parameter Estimation

We consider that the daily average of associated parameters provides a precise estimation of the general system behavior. This approach allows us to effectively capture the system dynamics by reflecting the mean value of updates made at shorter time intervals than a day, such as transaction updates.

2.3.3 EVMOS Token Price

The update of the EVMOS token price follows a stochastic process given by a Geometric Brownian Motion (GBM). The model parameters are defined by μ and σ, which are the mean and standard deviation of the underlying Wiener process, respectively. To estimate these parameters, we use historical Ethereum price data. We assume that the short price history of EVMOS is not sufficiently representative to use as data, and the Ethereum price is more appropriate.

2.3.4 Validators

Each validator in the simulation is assigned to a specific category that defines its profile and behavior in the ecosystem. These categories are defined by a risk profile associated with each validator at its initialization. We assume that each validator has a different risk aversion, which we will divide into 3 groups according to high, medium, and low risk. This risk aversion will determine:

- The target reward rate they aim to earn at each step (day).
- The percentage of capital they are willing to move in each action (stake, unstake, delegate).

In addition, to make the model more flexible and to relax the assumptions and behaviors assumed, each agent will be initialized with a probability that will govern:

- Apart from checking their target rewards, which action (check volatility or check price trend) they will take in the current step.
- The optimality of the choice of the delegating validator in delegations.

Risk Profile

As mentioned earlier, we assume risk profiles for validators `low`, `medium`, `high`:

- Low Risk Profile (`low`):
  - Daily Reward Target: 0.20 / 365
  - Percentages of funds that will move in each action (stake, unstake, delegate): 0.025
  - Volatility Tolerance: 0.15
- Medium Risk Profile (`medium`): Double the values of the low-risk profile.
- High Risk Profile (`high`): Triple the values of the low-risk profile.

These three profiles are assigned by dividing the total number of validators into 3 groups and assigning a different profile to each group. Since the initial stakes are randomly assigned to each validator at their initialization and since each step is not executed in initialization order but in random order, this risk assignment by group does not bias the model.

Wallet

Each validator has a wallet where they will store the EVMOS received as rewards for their actions, which will be used to pay the costs at the end of each month.

Initial Deposits (Stakes)

To establish the initial stakes of the validators, we have data on the deposits of 16 validators in one day. However, since this sample is relatively small, we have chosen to expand our data set using the bootstrapping technique. This methodology involves generating additional replicas of the original set, accurately maintaining the authentic distribution of deposits. Throughout the report we will refer interchangeably to stake or deposits.

Initial Holdings

At initialization, half of the available EVMOS tokens that are not being staked, are distributed equally among validators. This gives the model more flexibility to perform actions during the simulated year. The other half of the outstanding available supply will be assigned to a reserve in order to be used by the burning mechanism and the assignment of the priority tip at each step.

Initial Costs

We assume that each validator has an initial technological cost, which is calculated taking into account the information provided by validators through a google form and, as mentioned, expanded through data augmentation.

Actions

At each step, every validator executes the action associated to their daily target reward and then randomly decides whether to perform an action associated with volatility and trend, based on their initialization probability (self.action_prob). The actions are:

- Check the target reward.
- Check the trend.
- Check the target volatility.

In each of these cases, validators can only stake, unstake, delegate, or bring funds from a previous delegation (to their own wallets or re-delegate to another better validator).

- **Stake**
  - check target reward: If the reward they receive in each daily issuance exceeds their daily target, the validator assumes that this could continue to happen. Therefore, they stake a percentage of their available capital.
  - check trend: If the token price follows a designated upward trend metric:
    - The Simple Moving Average (SMA) of the token's closing prices over a period of 15 days for upward movement and 10 days for downward movements, is calculated.
    - The most recent price is compared to the corresponding SMA value.
    - If the current price is higher than the SMA, a percentage of the available capital is staked.
  - check volatility: If the token's volatility does not exceed a certain target volatility given by the risk profile of each validator, a percentage of the available capital is staked.

  The capital used for staking comes from the initial holdings and rewards accrued in the wallet as a result of ongoing actions.

- **Unstake**
  - At each step, if the reserved capital to pay the costs does not cover the monthly technological cost, after reviewing the redelegated capital and pending transfers (pending unstakes and pending redelegations), the validator will unstake the necessary amount to pay that cost.
  - check trend: If the token price adhere to a specific downward trend indicator:
    - If the current price is lower than the downward SMA, a portion of the total stake is unstaked.

  - check volatility: If the token's volatility exceeds the target volatility associated with the risk profile of each validator, a portion of the total stake is unstaked.

- **Delegate**
  - check target reward: If the reward they receive in each daily issuance from their own stake does not exceed the target, they will delegate to another validator. Although the validator could unstake, we assume that delegating is better than withdrawing funds from the total investment (unstake).
  - check target reward: For each delegation made, if the last rewards received by the delegation exceeds the daily target reward, an extra percentage of the total holdings will be assigned to the agent that generated these rewards.

- **Re-delegate**
  - check target reward: In each step, validators will go through all the delegations made on previous days and will make the following reasoning:
    * If the last rewards received from this delegation does not exceed the target reward, the owner of those funds will re-delegate to another (better) validator. In the case where the best validator are themselves, they will withdraw the delegated funds.
    * If the agent managing these delegated funds is deactivated, the owner of the funds will re-delegate the funds to another (better) agent.

In delegation and re-delegation actions, the logic combines an optimal and sub-optimal choice to provide greater flexibility to the model. Based on the probability associated with initialization, validators will randomly choose among the "top" (smallest) 5 validators, this being the optimal action, as the new reward distribution design benefits smaller players. Similarly, with the complement of this probability, they will randomly choose any validator and transfer the funds to him, representing a sub-optimal behavior. Both actions are performed on the set of **active** validators.

Delegations have an associated cost given by the commission that the delegator must pay to the validator who will manage these funds.

Additional Assumptions: To simplify the model, we will allow actions that only move a minimum amount of tokens (1e4 EVMOS tokens). This solves problems with operations involving insignificant token amounts, which would increase the required computational capacity without a significant impact on the results.

Additionally, we will only allow 50 pending re-delegations open at the same time. These two are strong restrictions, but after running the simulations without these limitations, no significant differences were observed in the final results, but there was a considerable computational cost.

Finally, each action executed by a validator involves an amount of EVMOS rounded down to two decimal places. This avoids problems of moving insignificant amounts and moving negative amounts (if the rounding is up, in certain steps it is observed that a validator wants to move more capital than it has).

Deactivation

Every 30 days, the technological costs that allow validators to perform their work are charged. If these agents do not have sufficient and available capital to make these payments, they will be deactivated until they normalize their debt.

When they are deactivated, the funds that have been delegated to them will be released, allowing all other agents to re-delegate these funds following the logic of re-delegation.

2.3.5 Evmos Chain

Network Demand (Base Fee)

Since the focus of the model is on studying the most relevant interactions regarding the objectives we want to track and achieve, we will model network demand through a daily value of gas used.

Since we do not have such abundant historical information, this value will be generated through a stochastic process based on Ethereum network data. However, to adapt it to the magnitudes of gas used in the Evmos chain, we will normalize the output of this process to match the values observed in minstcan's on-chain data.

Furthermore, to model scenarios of high, medium, and low demand, this normalization process will be made by using a parameter that we can change to represent these configurations.

Priority Tip

We assume that the priority tip and the network demand are correlated, so that higher priority tip results in validators earning more rewards. To correlate both factors, each time we generate a daily base fee, we will also generate a daily priority tip that will be a multiple (between 1% and 5%) of the base fee. For this, we will use the half of the available tokens moved to the reserve at initialization, to ensure that this token amount does not exceed the actual total available tokens in the entire ecosystem.

This priority tip will be distributed among activated validators based on their stakes.

Issuance

Issuance occurs daily, generating an amount that is a proportion of the total outstanding supply, including the burning from the previous step. This proportion is dynamically adjusted so as not to exceed an annual emission of 5%.

This adjustment is a weighted sum of two system's variables: staked ratio, compared to the optimal stake considered 0.66; and base fee, modeled through a Brownian process. This calculation, as mentioned in previous chapters, takes into account the influence of weights w_s and w_b, which were discussed in the Evmos' token economic redesign report and that will be updated in Chapter 4.

Reward Transfer

The Evmos chain agent distributes a proportion of the issuance to each activated validator on a daily basis, which is determined by the square root distribution design minus the staking commission of the validator/delegator, if it is the case. This reward transfer process is complemented by the burning process, where the daily base fee is burned at the end of the step.

Additionally, as we mentioned previously, the daily priority tip generated is distributed among activated validators randomly.

Burning Process

The burning process involves burning the gas used generated daily.

Token generation for Burning Process and Priority tip

To ensure the model is consistent, the tokens burned and the priority tips come from two sources of available tokens:

- Tokens not staked at initialization: Validators are initialized with an initial stake, which adds up to the initial staked ratio. As we mentioned several times before, the difference between total supply and total stake is a surplus that is allocated as follows: 50% will be assigned to validators' wallets to have capital to pay their costs and to execute actions; 50% will be reserved for use in burning and priority tip.

- Capital paid by validators at the end of the month in the form of technological costs.

Chapter 3

Parameter Sensitivity and Scenario Analysis

In this chapter, we will focus on studying two things. Firstly, the impact that different configurable parameters of the model have on the behavior and evolution of the metrics of interest. Secondly, the various values that the parameters will assume represent stressful scenarios for the model. These scenarios are necessary to cover in order to have greater confidence in the achievability of the proposed objectives.

Analyzing the sensitivity of parameters in a simulation model involves understanding how variations in these parameters affect the output metrics and it is also helpful in order to revise the stability of the overall system. Additionally, it is crucial to comprehend the range within which each key factor operates, as this knowledge enables the simulation of black swan scenarios. Understanding these extreme scenarios is vital for assessing the resilience and robustness of the system in unpredictable conditions.

It is important for a consistent analysis to establish a baseline, that is, to determine the output metrics (KPIs) under normal or baseline conditions. This will allow us to have a point of comparison.

The baseline will be the results under the following standard configuration:

- Steps = 365 days: This parameter represents the duration of the simulation, set to a full year.

- Chain demand = Medium (M): This corresponds to the average daily gas used, based on data obtained from mintscan. It provides a moderate level of network demand for the simulation: $daily\_base\_fee = 1e9$ and $daily\_gas\_used = 150$.

- Maximum delegations = 50: This parameter corresponds to the maximum number of validators that can be delegated to.

- μ, σ token price's stochastic process: These parameters are derived from historical data, where μ represents the average daily return of the token price, set to 0.005, while σ is the volatility of the token price, set to 0.075.

- Issuance weight: These are the weights used in the tokenomics report to calculate the emission, W_S represents the weight for stake, set to 0.000233, and W_B represents the weight for the base fee, set to 0.000249.

- Initial staked ratio = 0.5: This is the starting proportion of total tokens staked by validators.

- Number of validators = 150: This parameter determines the number of validators participating in the network.

- Power of the rewards distribution ($p = 0.5$): This parameter corresponds to the power of the rewards distribution function. A value of 0.5 indicates a square root distribution, affecting how rewards are distributed among validators.
- Fixed demand = True: This parameter specifies whether the network demand is fixed or varies over time. In this case, it is set to True, meaning the demand won't vary.
- Risk and optimality profile distributions: These distributions determine the risk and optimality profiles of validators, and are represented by random Dirichlet distributions with $\alpha = [1, 1, 1]$ (equally distributed). The highest number in α means that the Dirichlet distribution will return a list of three probabilities that sum up 1, being the highest the one associated with the bigger value of the triple used in the input.

3.1 Parameters to vary

In each of the simulations we will conduct, we will fix all parameters to the standard configuration and only vary the parameter under study. We will analyze the following parameters within the specified ranges. The choice of these parameters is justified by the following reasoning:

Seed

To give the model greater flexibility and relax the assumptions made about the behavior and profile of each validator, many of the actions and order of activation of the agents are performed randomly. To ensure that the results obtained are not a product of the noise generated by this random input, we must run the model several times by changing the randomness process. Seed: [0, 1, 2, 3, 4, 5, 6, 7, 8, 9].

Chain demand

Varying this parameter shows its sensitivity but also serves as specific and extreme study scenarios. We can study situations where the demand for the network remains at extreme conditions for an entire year. Chain demand: ["Low", "Medium", "High"]

Maximum delegations

With the aim of limiting the size of the model, we have established a maximum number of possible delegations. In this way, each validator cannot exceed this number. To assess the impact that this restriction has on the results, we have proposed different maximum quantities. It is important to note that this limitation does not exist in reality; therefore, the higher the possible number of delegations, the model becomes more realistic. Maximum Delegations: [15, 30, 50, 100, 150].

Token price's stochastic process (μ, σ)

As the costs are in dollars but the income is in a non-stable token, it is interesting to see if the evolution and behavior of the agents change significantly when we subject the token to varied and extreme situations regarding its stochastic dynamics. μ, σ token price's stochastic process: μ: [0.005, 0.01, 0.05], σ: [0.075, 0.1, 0.2].

Issuance weights

These parameters are crucial in the entire issuance process, as the emission function depends on two fundamental terms: one associated with how far the staked ratio is from the optimal staked ratio, which is 0.66, and the other determines an inversely exponential relationship with the network demand, represented by the base fee in the formula. The issuance depends on two weights that controls the impact of each term in the emission. We can verify how changes in the staked ratio and in the base fee will impact the dynamics of inflation and decentralization of the network by varying these weightings. Issuance weights: w_s: [0, 2, 4] * W_S, w_b: [0, 2, 4] * W_B,

Staked ratio

We believe it is important to analyze the network's stability in terms of certain initial conditions, in this case the initial staked ratio. We consider a staked ratio of 0.66 to be a good target measure to achieve. Therefore, it is necessary to understand how the network behaves when we are far from or close to this target. As a result, we decided to take the following values: [0.1, 0.5, 0.9].

Number of Validators

The objective is to simulate a year with the maximum possible number of validators, also evaluating how the number of validators impacts the time needed to achieve the desired objectives. This will give us a better understanding of the sensitivity of this parameter and explain its impact on the model's evolution. Number of validators: [50, 100, 150].

Power of the rewards distribution (p)

Although we have decided to implement the economic system of the new design under the square root distribution approach, it is interesting to study the impact that the power p has on the behavior and general evolution of our system. It is expected that when p is smaller, the reward distribution tends to benefit validators with lower stakes more, and distribute less to validators with higher stakes. Power of the rewards distribution (p): [0.1, 0.5, 0.9].

Fixed demand

The variability associated with this parameter represent one important aspect. By allowing it to vary, we introduce flexibility to the model, enabling the simulation to represent a more realistic behaviour and covering a wider range of situations. Fixed demand: [True, False].

Risk and optimality profile

The purpose of varying this parameter is to ensure flexibility within the model and to encompass different configurations of risk and random choices in actions. This approach allows us to ensure the model's impartiality and that, regardless of the validator profiles, the overall behavior aligns with the desired objectives for the network. We will stress the model by allocating the highest proportion of agents to: low, medium and risk profile, and to low, medium and high optimality. Risk and optimality profile distributions: Random Dirichlet distributions with α in the list $[[10, 1, 1], [1, 10, 1], [1, 1, 10]]$.

3.2 KPIs charts

For every scenario analyzed, we will track the changes on inflation and decentralization of the system though three different charts

- Inflation Chart: The simulations will be examined in conjunction with three charts depicting the evolution of key performance indicators (KPIs). We will compare the trends of annualized inflation and daily emission against the staked ratio and the base fee. Since the emission function dynamically adjusts based on these two variables, which are closely linked to validators' behavior and network demand, it is intriguing to analyze how changes in these signals affect inflation and emission.
- Decentralization charts: We will use two charts to have a supported conclusion of the behavior of this metric
  - Decentralization chart: In this chart, we compare the Gini coefficient, which measures inequality associated with a population attribute, in our case, the total staked, against the progression of stakes of each validator. This evolution of stakes will be depicted through three metrics: the sum of the stake of those agents below the 90th quantile of the population (small and medium agents); the sum of the stake of agents above the 90th quantile (large agents); and a metric that measures the ratio of both amounts to know which magnitude is larger and by what proportion.
  - Histogram of stakes chart: To support the previous chart with an additional picture of the system's dynamic, we will show the distribution of stakes on day 0 (initialization) and on day 365 (completion).

3.3 Results Analysis

We have carried out all the sensitivity analyses of the previously described parameters and have concluded that it is sufficient to present the demand and the maximum number of delegations analysis, as almost all other cases are similar. This is because changes in other parameters produce consistent and comparable results in terms of inflation, decentralization, and other aspects of the model. Therefore, we will define an exclusive and comprehensive subsection to present the results of the sensitivity analysis for these two parameters.

3.3.1 Demand

As we mentioned before, we will focus first on analyzing in detail the impact that the demand parameter has on the evolution of the system and on the metrics of interest.

For our simulation, we have identified three levels of demand: high, medium, and low. Each of these demand levels will influence key indicators such as inflation, validator behavior, and decentralization. As you may notice, the case for medium demand is our baseline.

Inflation

The analysis of inflation in all cases is closely correlated with daily issuance. This is because inflation depends on the total supply of the token, which in turn is determined by what is minted and what is burned. In this model, the only thing burned is the base fee, so when it remains constant, daily issuance plays a crucial role in determining inflation.

In all scenarios, we notice that the staked ratio follows the same pattern: it starts at 0.5 and grows to approximately 0.66, as seen in Figure 3.1, even when demand is set at higher values.

We observe a drastic change in daily issuance when considering high demand compared to medium and low demand. This is predictable because in the issuance formula, when the base fee is high, the term associated with stake becomes more significant. Therefore, we see that issuance inversely correlates with the staked ratio. In contrast, for low and medium demand, the term associated with the base fee becomes more relevant in the issuance formula, resulting in a linearly increasing daily issuance, without many variations. Linearity occurs because for certain staked ratio values below 0.54, the daily issuance reaches the maximum allowed, and the *minimum* function involved in the theoretical design returns 5%. This critical staked value can be deduced from the minting function:

$$f(t) = (1 + R(t)) * Total\_supply_t$$

where

$$R(t) = \max\left\{0, \min\left\{r_s, \omega_s\left(1 - \frac{s(t)}{s^*}\right) + \omega_b e^{-BaseFee(t)1e^{-9}}\right\}\right\}. \tag{3.1}$$

We want $\overline{s}$ such that

$$\omega_s\left(1 - \frac{\overline{s}}{s^*}\right) + \omega_b e^{-\alpha_{\text{BaseFee}} BaseFee(t)} = r_s$$

$$\Rightarrow \overline{s} = s^*\left(1 - \frac{1}{\omega_s}\left(r_s - \omega_b e^{-\alpha_{\text{BaseFee}} BaseFee(t)}\right)\right) \sim 0.54.$$

The decrease in inflation that we observe for medium demand compared to low demand from day 25 onwards is caused by the staked ratio exceeding 0.54. This results in the issuance responding inversely proportional to the dynamics of the staked ratio. With high demand, we see how inflation decreases, reaching even the value of zero.

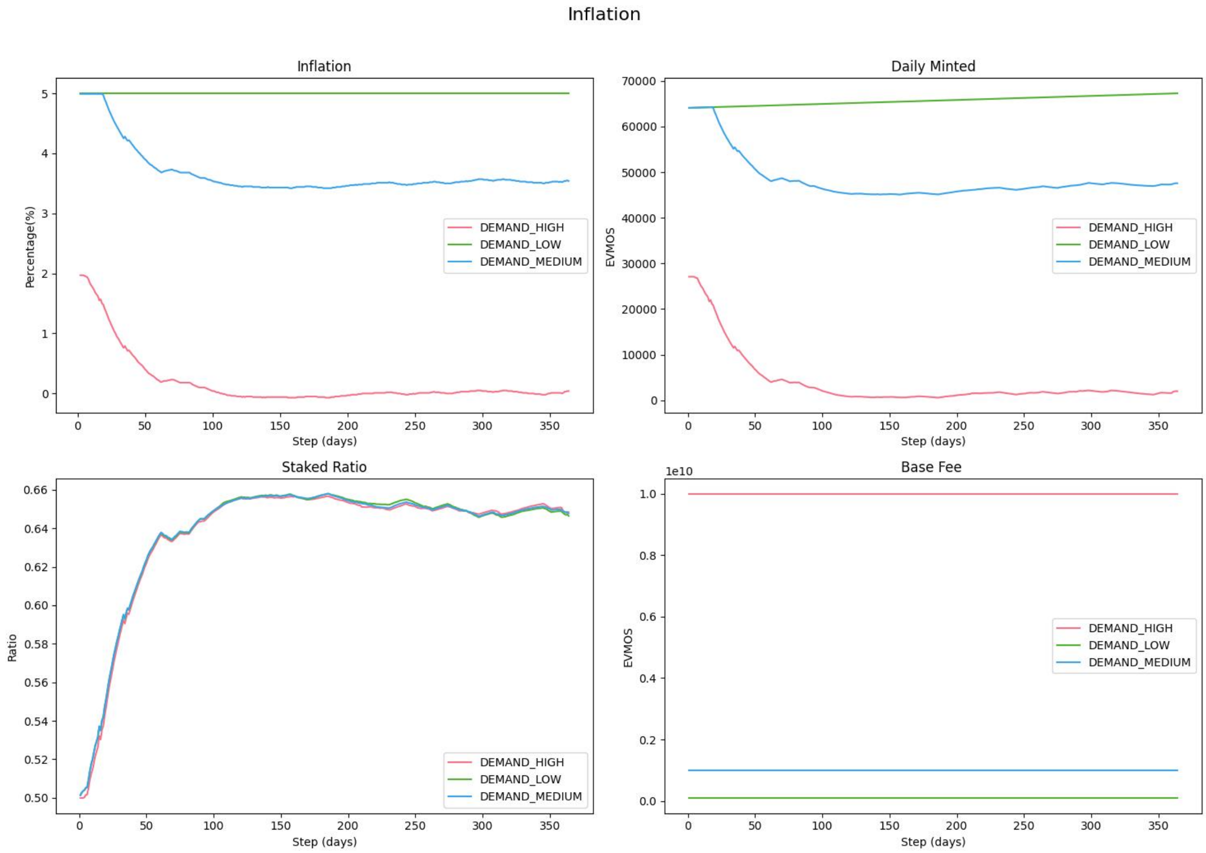

Figure 3.1: Inflation - Demand Scenarios

Decentralization

In all demand scenarios, we observe significant network decentralization in Figure 3.2, where the Gini coefficient decreases from values close to 0.65 to a range between 0.08 and 0.12. It's noteworthy that in these simulations, the charts exhibit quite similar behavior, though we may observe slightly lower decentralization achieved for low and medium demand.

In addition to this overall trend towards decentralization, a closer examination of the charts in Figure 3.2 reveals other relevant aspects:

- The large validators, which represent the 10% with the largest participations, reduce their share, going from an initial total of a little more than 150M EVMOS to values between 40M and 60M EVMOS, during the analyzed period.
- Conversely, medium and small validators (the remaining 90%) exhibit growth in their share, increasing their stake from a total of 80M EVMOS to values between 250M and 275M EVMOS.
- Additionally, we observe a significant increase in the ratio between the deposits of medium and small validators compared to the large ones. This ratio, initially at 0.6, grows to values between 6 and 7 times in the first 100 days before oscillating 5.5 and 6.5. This chart describes a high final decentralization, in the same sense than with the Gini coefficient.

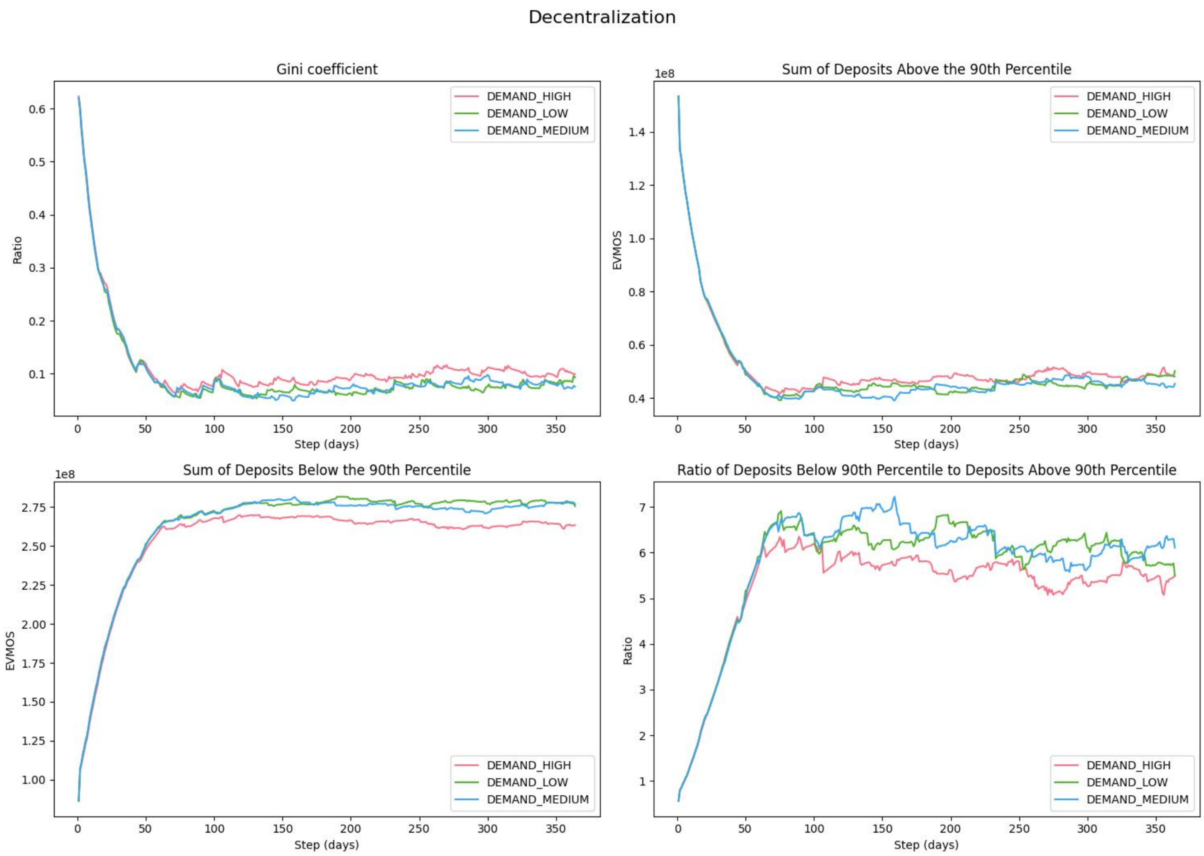

Figure 3.2: Decentralization - Demand Scenarios

This last conclusion is also supported by the progression of stakes observed in the histograms shown in Figure 3.3. Regardless of the network demand, after one year, there is a higher frequency of validators with lower stakes and a reduction in the frequency of validators with higher stakes. This suggests a clear trend towards a more equitable distribution of EVMOS among the participants of the Evmos network, thereby supporting the decentralization process observed in the previous results.

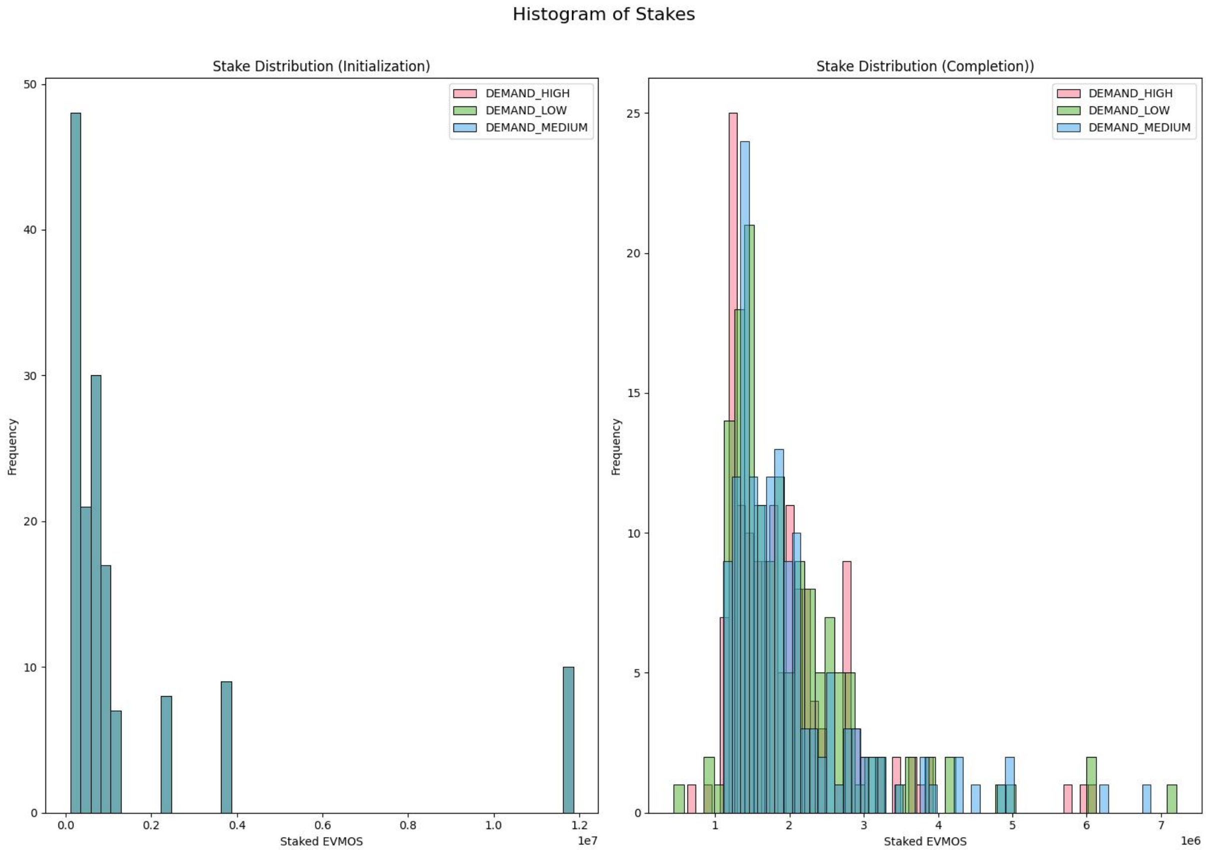

Figure 3.3: Histogram of stakes - Demand Scenarios

3.3.2 Maximum Delegations

In this section, we will analyze how the model changes when varying the maximum number of delegations. Specifically, we will study its impact on inflation and the decentralization of the network.

Inflation

In Figure 3.4, we observe that as we allow a higher number of delegations, the staked ratio increases and exhibits less volatile movements. It is noteworthy that allowing 150 delegations (the maximum considered in this analysis) over the course of a year, the value hovers around 0.62, approaching the optimal staked ratio of 0.66. Considering what was mentioned in Section 3.3.1, inflation reacts solely and inversely to the staked ratio, as the base fee is fixed, decreasing further as we allow a greater number of delegations.

It is worth mentioning that in the case of 15 maximum delegations, it is observed a rather chaotic behavior due to the higher impact that the actions associated to volatility and trend have over the action associated to the target reward. For lower values of maximum delegations, the possible outstanding open delegations is limited and, therefore, the staked ratio is capped.

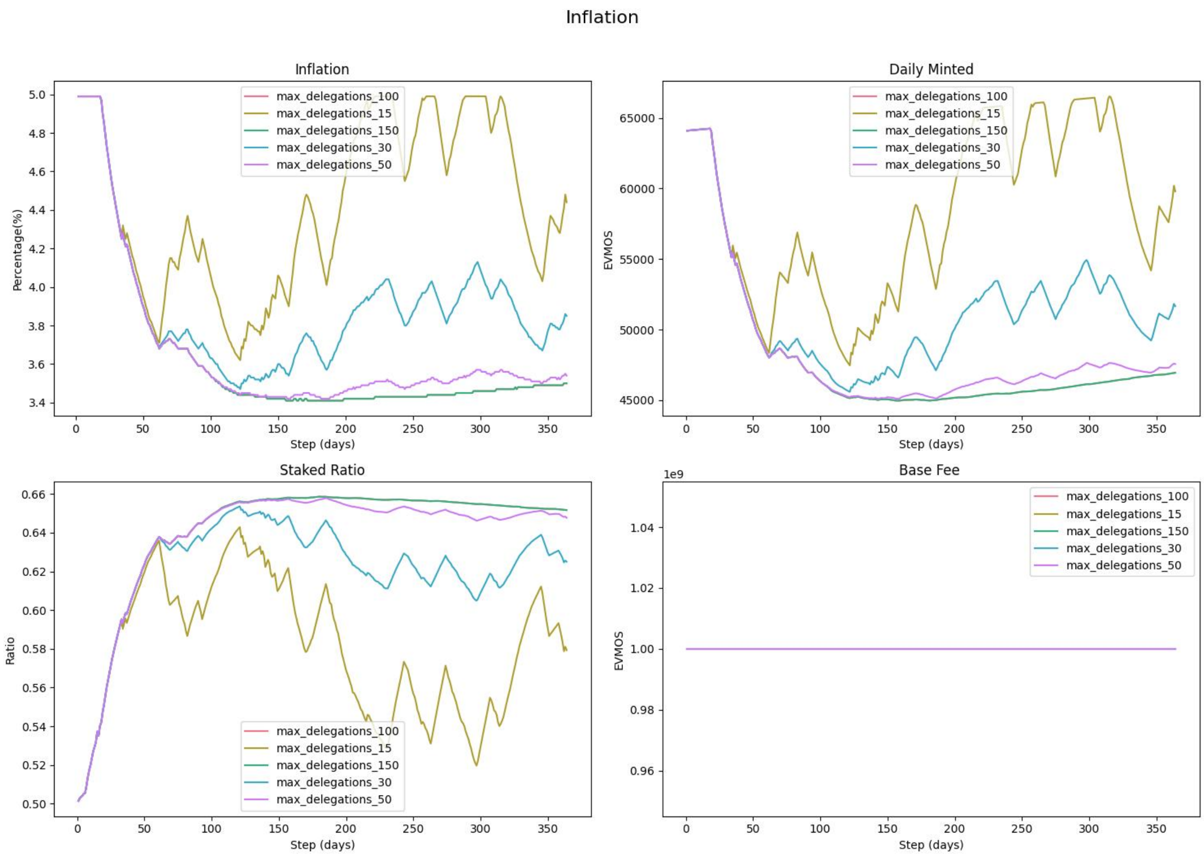

Figure 3.4: Inflation - Max Delegations Scenarios

Decentralization

The number of delegations is a crucial factor in the model, as it has a significant impact on the decentralization of the network. As shown in Figure 3.5, the Gini coefficient is inversely proportional to this parameter, meaning that as we allow more delegations, the network tends to decentralize over time. It is interesting to note that this change is observed as time progresses; for example, with 50, 100 and 150 allowed delegations, we see that until day 25 the Gini coefficient remains the same as when we have 15 delegations, but then it begins to decrease even further to levels below 0.1, while with 15 delegations it reached levels of almost 0.3.

In the other graphs of Figure 3.5, more details about decentralization can be observed. It is noted that the proportion of the total stake of small and medium validators over the total stake of large validators, considering 50, 100 and 150 as the maximum number of delegations, reaches a peak of approximately 6 times around day 75. If we consider 15 delegations, this proportion remains between 2.5 and 3.

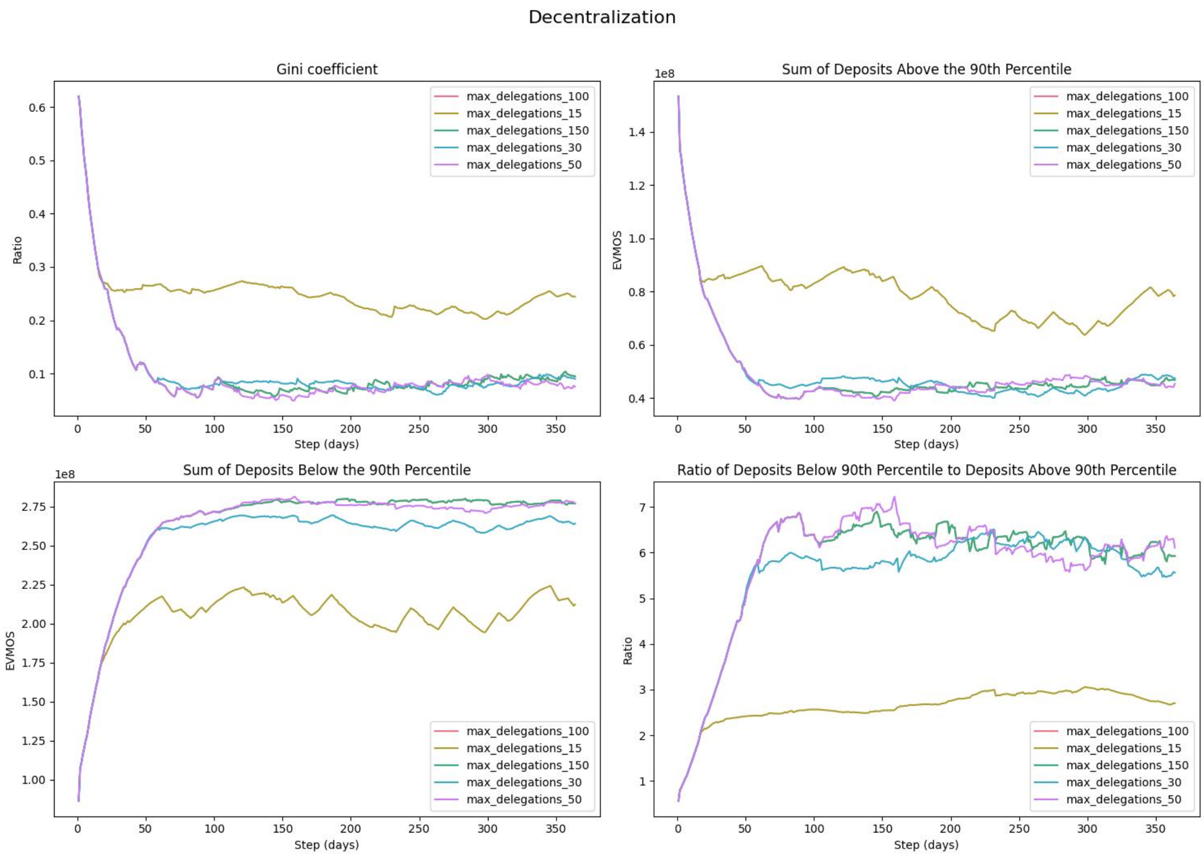

Figure 3.5: Decentralization - Max Delegations Scenarios

Another interesting feature to notice is that, as we can see in all the charts, considering 150 possible delegations give the same result as considering 100. This suggests that an equilibrium is reached beyond which any increase in the number of delegations becomes sub-optimal.

Finally, while in practice there's no upper cap on the number of delegations, our findings illustrate that adjusting this threshold impacts decentralization. Concurrently, altering this parameter could lead to substantial computational savings. There's a balance to be struck between the computational resources allocated for processing delegations and achieving desired levels of decentralization.

3.3.3 Other Parameters Impact

In this subsection, we will comment on details about the parameters that, when varied, do not essentially change the results of the model. In order to do this, we will briefly explain what changes in the results and show the evolution of inflation and Gini coefficient.

Seed

By varying the seeds, everything configured as random in the model (validator actions, EVMOS price dynamics, and the base fee) undergoes changes. For all seed changes, we observe an annualized inflation with a similar pattern than in the medium demand scenario, which does not exceed 5% annually and declines to a range of 2.8% to 4.5%.

Despite the seed variations, the Gini coefficient tend to decrease also, indicating a trend towards greater decentralization, regardless of the model's randomness.

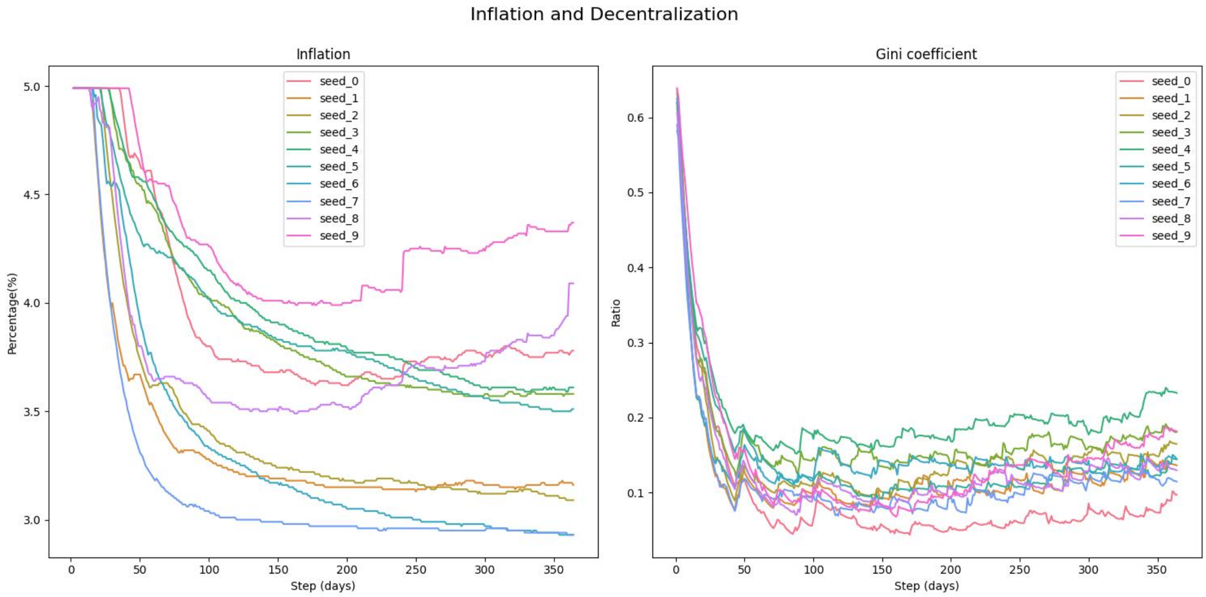

Figure 3.6: Inflation and Gini Coeff. - Seed Scenarios

Mint Weights

These scenarios illustrate the impact and sensitivity that the emission function has concerning the weights assigned to variables associated with the percentage of tokens staked and the network demand.

The dynamics of inflation are alike to those in the demand scnearios. For example, when we take the weight associated with the absent staked ratio, $w_s = 0$, the emission would remain at 3.4% annually as it would only react to the dynamics of the base fee. When we take four times the optimal weight associated with the staked ratio, we see how the emission increases as a consequence of the weighted sum reaching the annual limit of 5%. When we consider the weight associated with the base fee as absent, $w_b = 0$, inflation would only react to the staked ratio, leading to a value less than 0.5%.

In all cases, we observe a trend towards decentralization, experiencing a decrease in the Gini coefficient to below 0.1

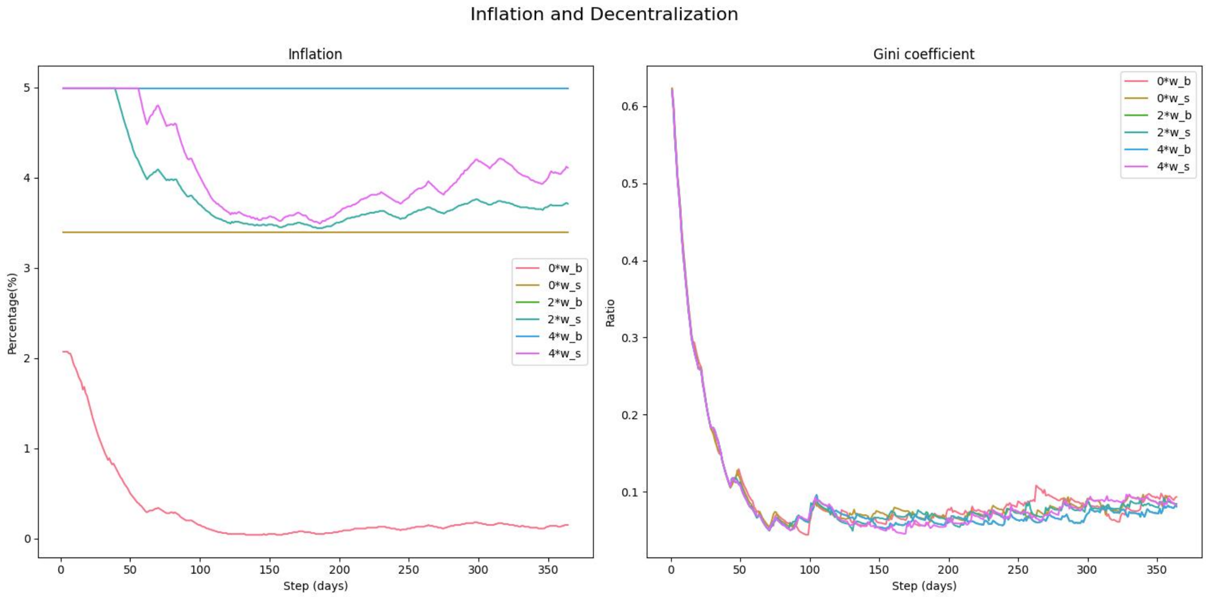

Figure 3.7: Inflation and Gini Coeff. - Mint Weights Scenarios

Number of Validators

It is noteworthy that inflation tends to decrease more sharply when we have fewer validators. For instance, with 50 validators, we obtain around 3.5% annualized daily emission, and with 150, just slightly exceeding 3.5.%. This is explained by the staked ratio reaching higher values. However, take into account that it does not necessarily mean that a lower number of validators will lead to a higher staked ratio, It simply relates to the number of simulated steps. Lower the number of validators, given the same total supply and the same initial staked ratio, the magnitude of each movement of funds is higher than considering more validators.

As the number of validators decreases, we observe that the Gini coefficient tends to take higher values, dropping from 0.2 to below 0.. However, it is important to note that overall the model tends towards decentralization.

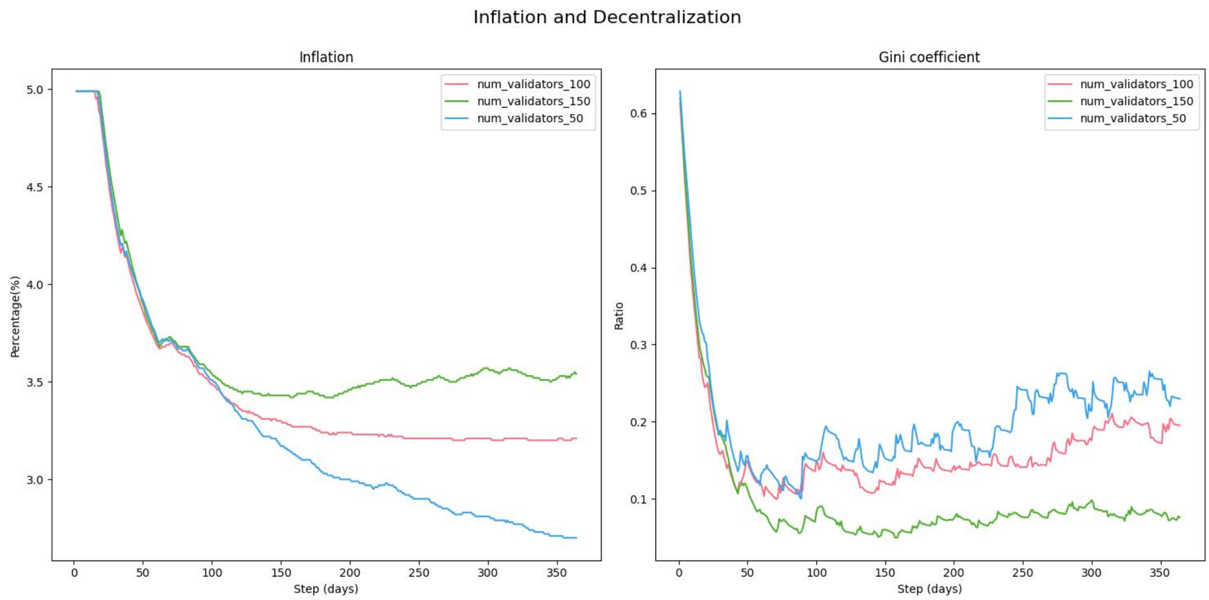

Figure 3.8: Inflation and Gini Coeff. - Number of Validators Scenarios

Power of the Rewards Distribution (p)

We know that with a lower p, the design change has a stronger impact on the rewards received by the agents, benefiting smaller agents and disadvantaging larger ones. With $p = 0.9$, we see a Gini coefficient stabilize at 0.09, and with $p = 0.1$, it stabilizes at 0.08 meaning that for higher players is preferable to delegate more as p increases. The dynamics of inflation is similar for changes in p, with lower emission occurring when p is smaller.

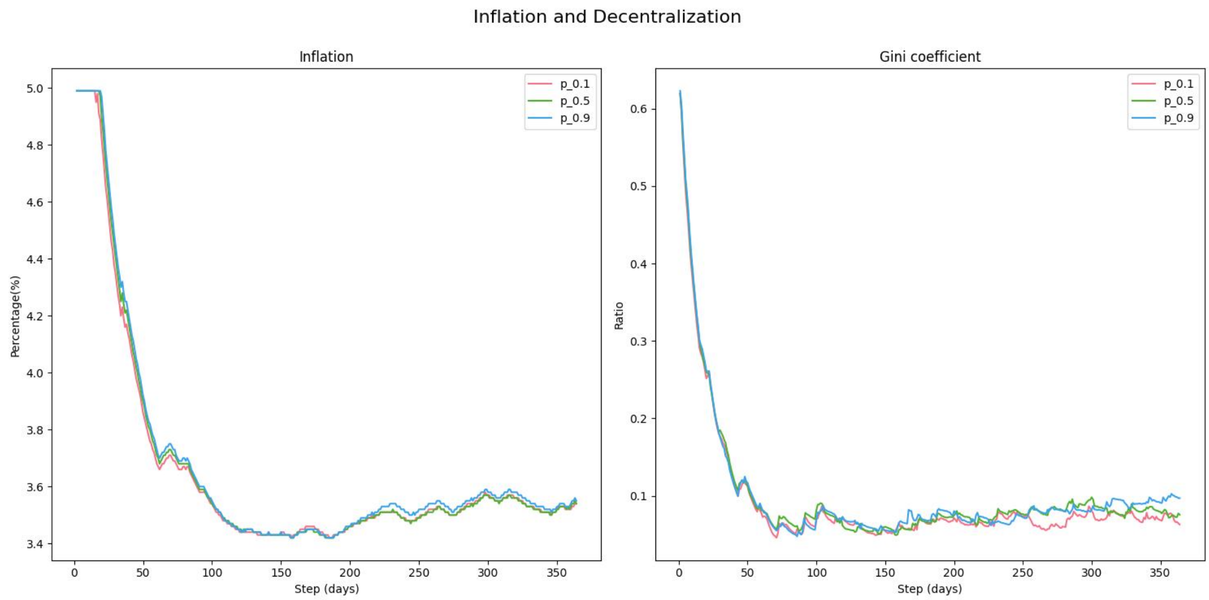

Figure 3.9: Inflation and Gini Coeff. - Power of Rewards Dist Scenarios

Fixed demand

To relax the assumption of a fixed demand throughout the year, which implies a constant number of transactions during the simulated year, we allow this attribute to vary following a stochastic process based on historical data.

In the variable case, inflation has a more volatile dynamic as expected due to the different magnitudes that the base fee takes. The base fee experiences abrupt changes, going from 500M to 2.0 billion EVMOS on day 250, which impacts the daily emission inversely proportional, reaching values below 1.5%.

The decentralization of the network tends to increase in all scenarios, as seen in the Gini coefficient graph (Figure 3.10).

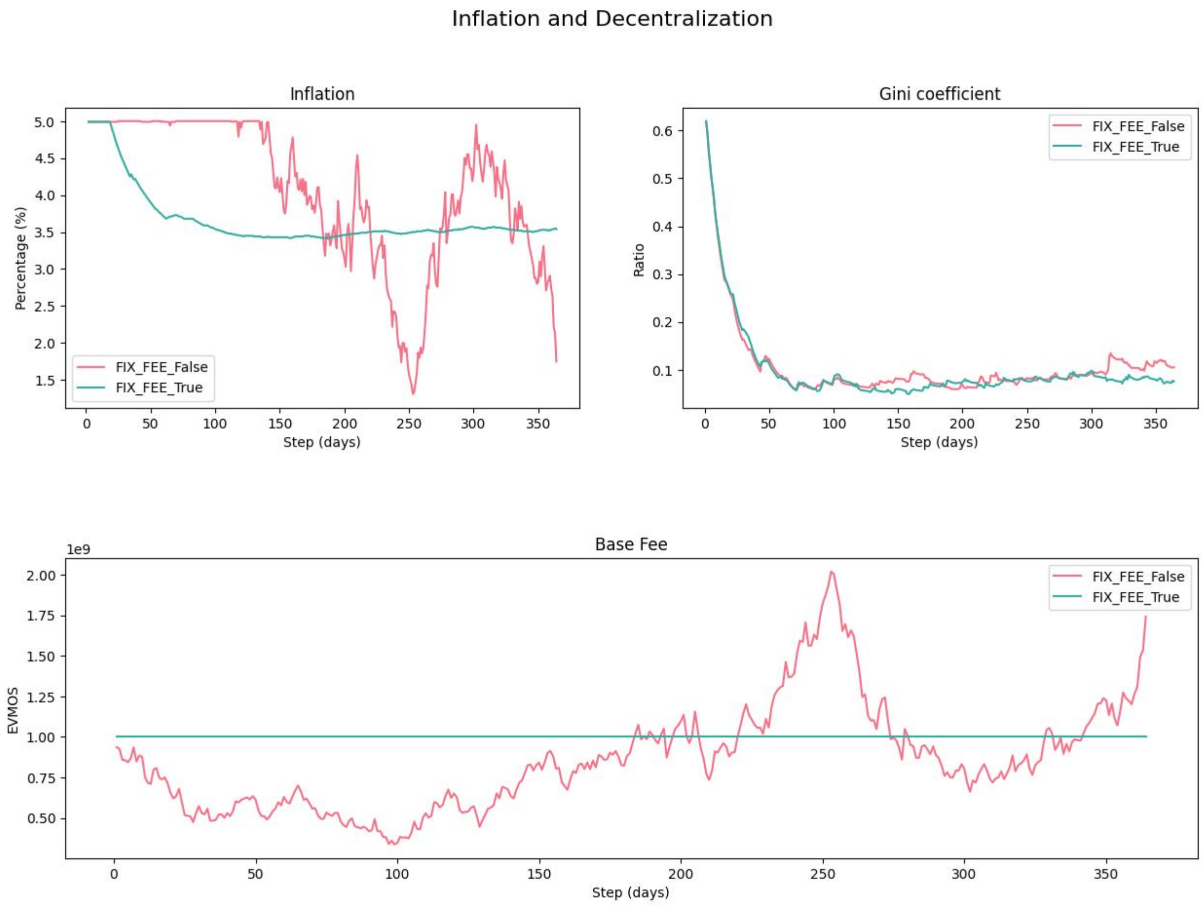

Figure 3.10: Inflation and Gini Coeff. - Fixed Demand Scenarios

Risk and optimality profile

In order to vary the behavior of the validators, different risk and optimality profiles are considered.

An initial observation reveals that as the network features a greater number of high-risk validators, the inflation rate tends to decrease. This is attributed to the increase in stake, as high-risk validators typically delegate more, and inflation is inversely related to stake.

On the other hand, it is observed that in all cases, the Gini coefficient decreases from 0.6 to values between 0.1 and 0.3. In the scenario with the highest number of high-risk validators, an upward correction is observed from step 80; however, the coefficient remains below 0.3 at all times.

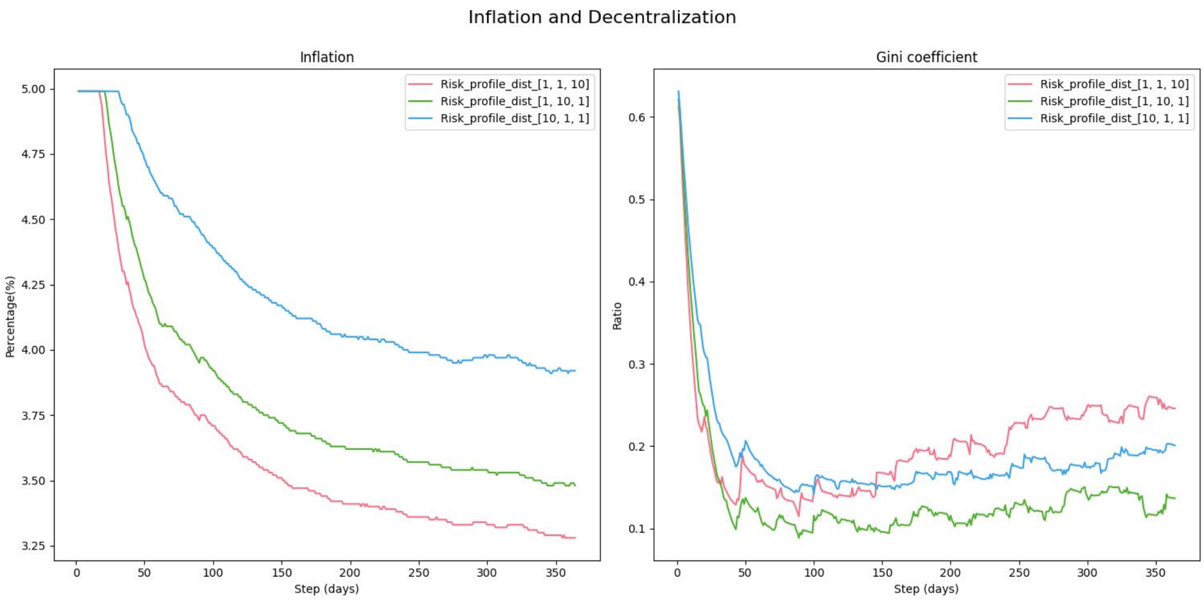

Figure 3.11: Inflation and Gini Coeff. - Risk Profile Scenarios

With regards to the optimality profile of validators, we can observe in Figure 3.12 that when we consider a larger number of validators with a more optimal behavior (meaning they are more likely to choose to delegate to the top 5 validators with the lowest stake), inflation tends to decrease and decentralization tends to increase in the system. Similarly, in all the cases considered, we see a similar trend.

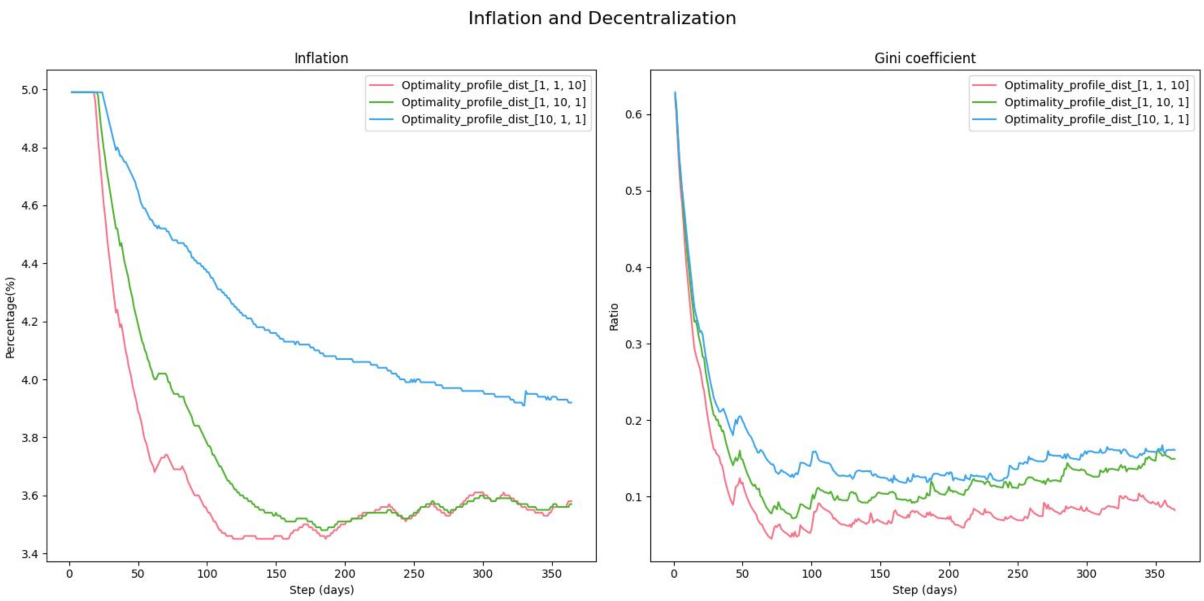

Figure 3.12: Inflation and Gini Coeff. - Optimality Profile Scenarios

3.4 Special Cases Analysis

In this section, we will analyze cases that exhibit a slightly different behavior from validators than what was observed in the previous section. It is intriguing to study these to understand which parameters affect their behavior. Nonetheless, in all these cases, inflation and decentralization exhibit similar behaviors to those described in the previous section, so we will focus only on the differences.

3.4.1 EVMOS Price

We have designed several scenarios for the EVMOS price, considering that this price is generated through a stochastic process that takes into account two important parameters: the price trend of the token, represented by μ, and the volatility, indicated by σ. We did not find any substancial difference when varying the trend factor μ so we will focus on the strange findings when moving the volatility factor σ.

Inflation

In Figure 3.13, we observe the effect of varying the token's volatility, represented by different values of σ:

For the lowest volatility $\sigma = 0.075$, the staked ratio increases from 0.5 to 0.64, resulting in a decrease in inflation from day 40, reaching a value of 3.8% at the end of the year. A similar behavior is observed with $\sigma = 0.1$.

For higher volatilities, for example $\sigma = 0.2$, we observe a significant reduction in the stake ratio triggered by a collapse of the price. What is happening in these extreme cases is that as the price reaches 0, only large players can afford their expenses, while the rest of the players must exit the system. This results in inflation tending towards 5%, the maximum allowed limit.

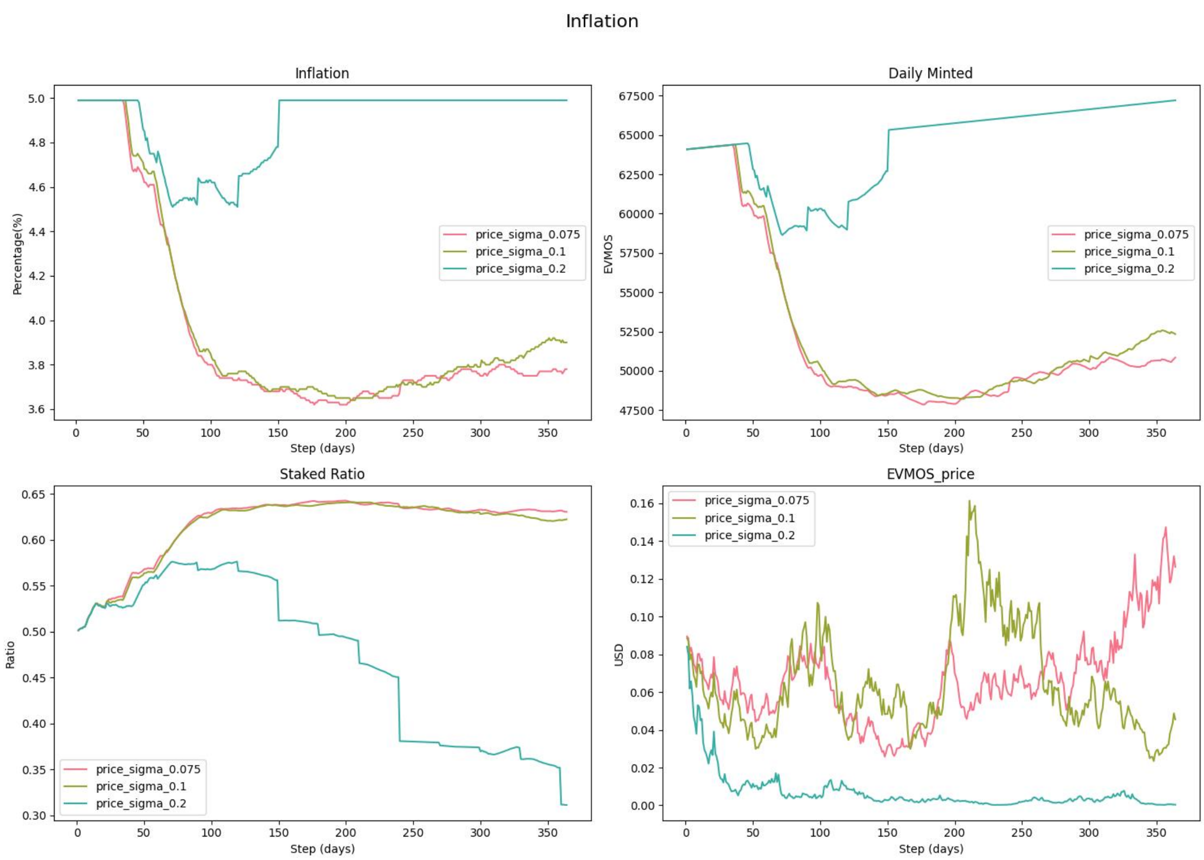

Figure 3.13: Inflation - EVMOS Prices Special Case (sigma)

Decentralization

Now let's analyze how changes in the volatility of the EVMOS token impact the decentralization of the system.

In Figure 3.14, we observe that the Gini coefficient begins to decrease until some step where it increases or remains constant. Higher volatility means a lower stake ratio due to more and more agents deactivating, leading to a higher Gini. For cases of low volatilities, decentralization is similar to that in the base case.

We observe that in the case of $\sigma = 0.2$, the Gini coefficient decreases until day 100 and then increases. Additionally, we can see how the proportion of the total stake of small and medium agents to the top 90% decreases from a ratio of 5 to around 1.5, contrasting with the low volatilities that remain in a range of 6 to 7 times.

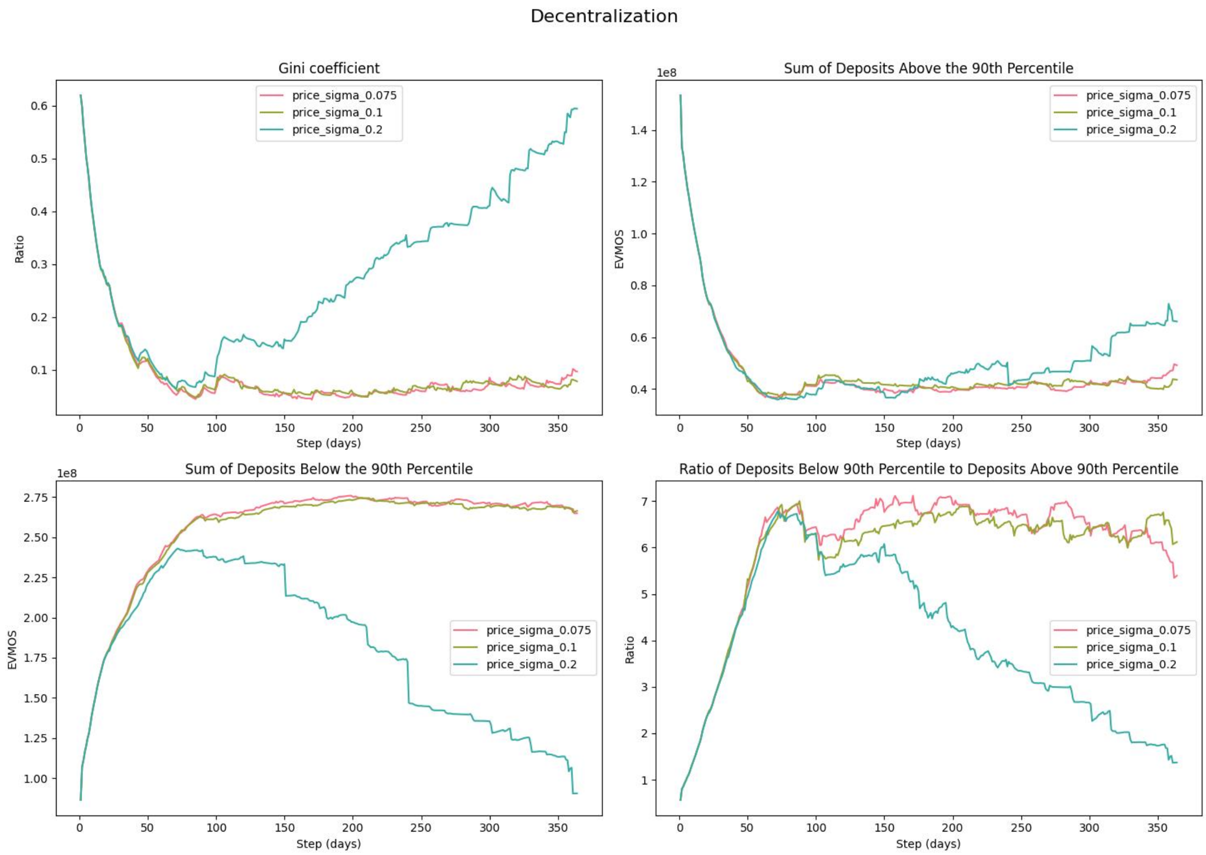

Figure 3.14: Decentralization - EVMOS price scenarios (sigma)

3.4.2 Initial staked ratio

Let's now move on to the analysis of scenarios that start with extreme staked ratios to stress the emission function and the general behavior of validators.

Inflation

In Figure 3.15, we observe that when we start with a low *staked ratio* of 0.1, it tends to increase over time. This results in a linearly increasing daily issuance, reaching almost 70,000 EVMOS, which translates in a final issuance of 5%. On the other hand, if we consider a high *staked ratio* of 0.9, we see that over time it tends to decrease until reaching 0.85. This results in an issuance of around 0.5% annually. Finally, if we start with a medium initial staked ratio of 0.5, we see that it tends to increase until reaching the optimal value of 0.66 which gives an issuance of 3.5%. In every case, staked ratio tends to move towards the optimal staked ratio of 0.66 and inflation never exceeds an annual 5%.

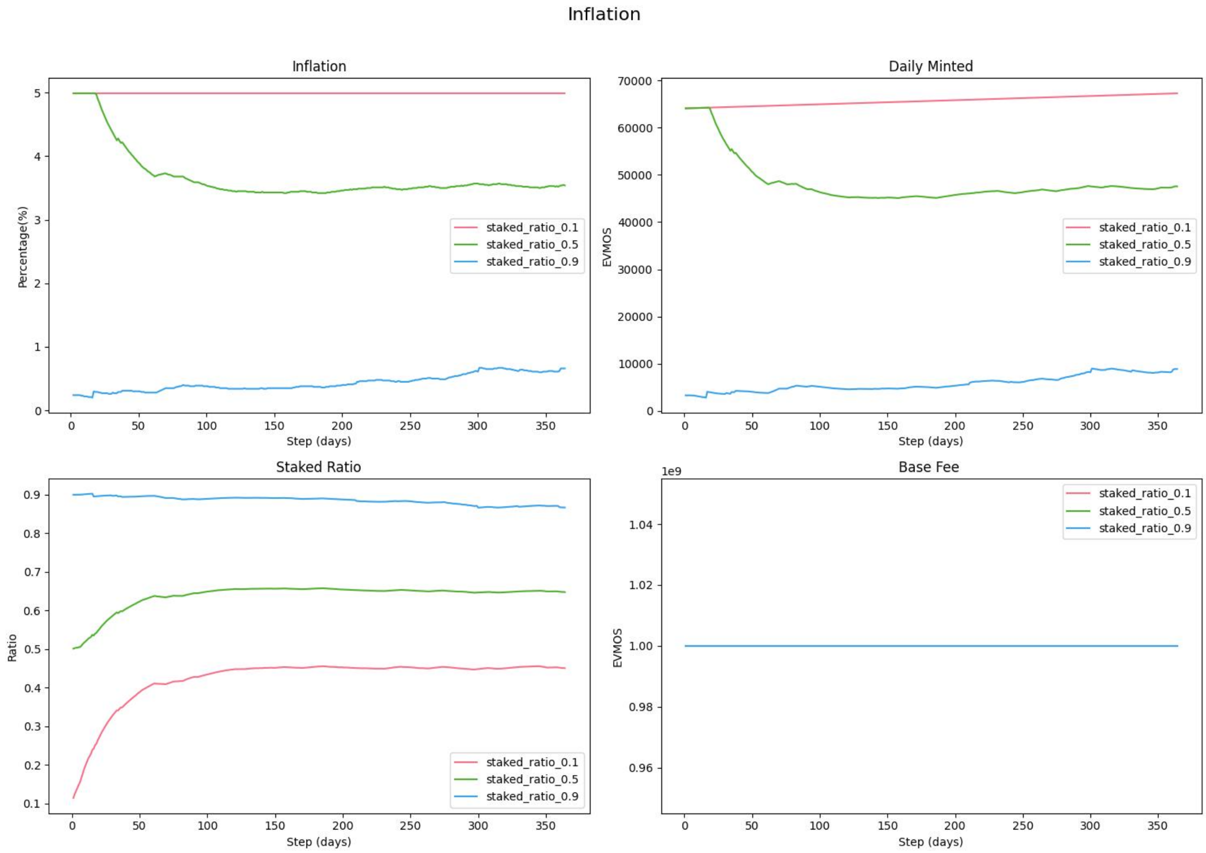

Figure 3.15: Inflation - Initial staked ratio Scenarios

Decentralization

We observe a significant reduction in the Gini coefficient. In every case, we notice how in the first 100 days it decreases rapidly, going from 0.65 to values between 0.05 and 0.1. From step 100 onwards it increases until stabilizing between 0.08 and 0.2.

For the highest *staked ratio* considered, we observe that the top 10% of validators with the largest stake delegate their deposits, meaning a decrease from 270M EVMOS to 70M EVMOS, with the stakes of the medium and small validators increasing initially from 160M EVMOS to a maximum of 370M and then stabilizing around 350M EVMOS.

In the case of an initial staked ratio of 0.5, all the trends are similar to the previous case, with differences in the magnitudes of the stakes involved. The highest players reduce their stakes from 150M EVMOS to 50M EVMOS, while the other players increase their total sum of deposits from 80M EVMOS to almost 260M EVMOS.

Additionally, it is interesting to note that for the lowest initial *staked ratio*, we see how the top 10% of validators with the largest stake decides to delegate a little portion of their stakes in the first 25 steps, increasing the stakes of the remaining 90% of validators. This is explained by the fact that the total stake of these big players is much less than in the other two cases, therefore the optimal amount of tokens to be delegated is lower.

The final relation between the size of each group of agents reaches a maximum value of 8 times for the initial staked ratio of 0.1, and of around 7 times for the other two cases. It tends to decrease after step 100 and to stabilize between 5 and 6.2.

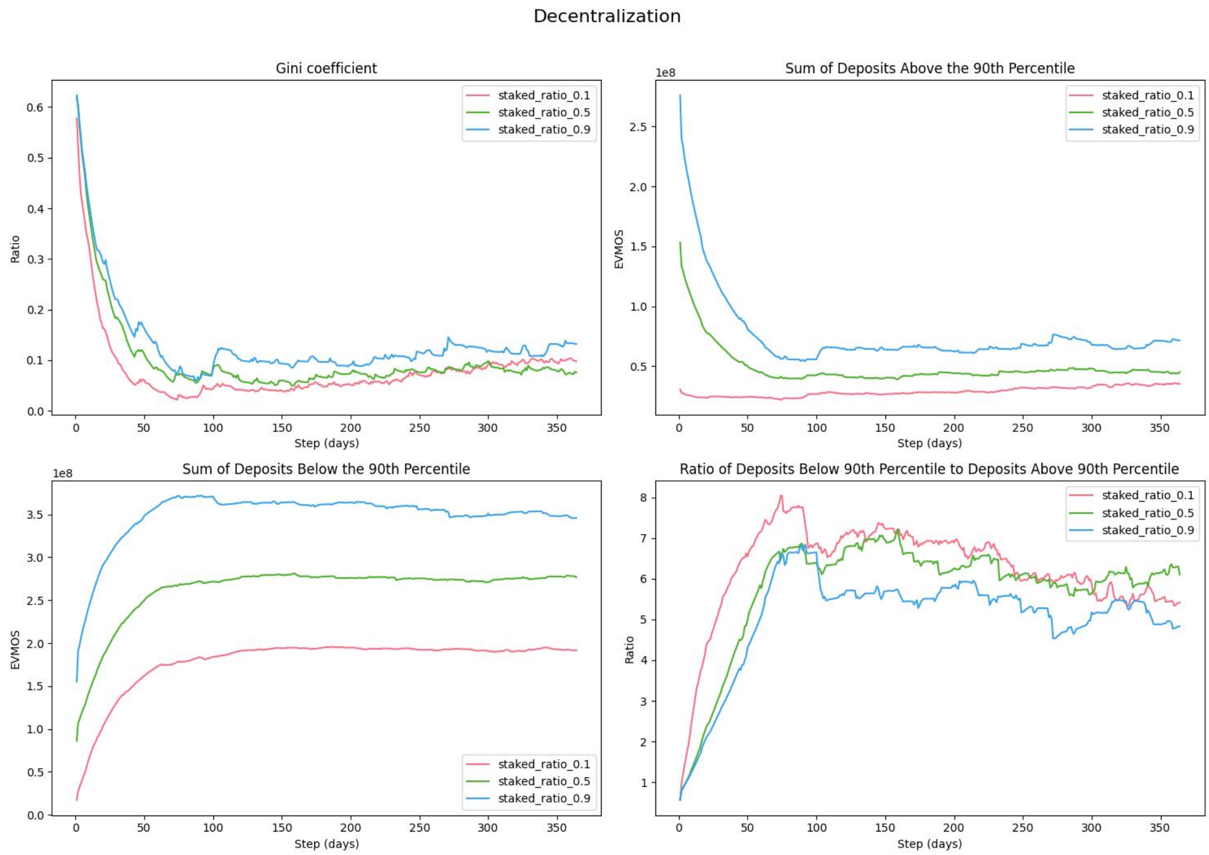

Figure 3.16: Decentralization - Initial staked ratio scenarios

3.5 Pearson / Spearman Correlation

In this section, we will study the correlation between inflation and the key factors that alter it: staked ratio and base fee. In order to metrize the impact that each variable has towards inflation, we will use two coefficients of correlation, *Pearson correlation* and *Spearman correlation*. We will focus on the case with the fix demand set as False, as we need the base fee to change in order to execute this analysis.

The *Pearson correlation* measures the linear relationship between two variables. It is assumed that the variables are normally distributed and that there is a linear relationship between them. A value of $+1$ indicates a perfect positive linear correlation, -1 indicates a perfect negative linear correlation, and 0 indicates no linear correlation. This correlation is shown in Figure 3.17.

The *Spearman correlation* measures the monotonic relationship between two variables, it means that whether as one variable increases, the other tends to increase or decrease. It does not assume that the variables are normally distributed or that they have a linear relationship. This correlation is illustrated in Figure 3.18.

The Pearson correlation between base fee and inflation is approximately -0.96 and of -0.98 for Spearman, indicating a strong negative correlation. This suggests an almost perfect inverse relationship between base fee and inflation: as base fee increases, inflation

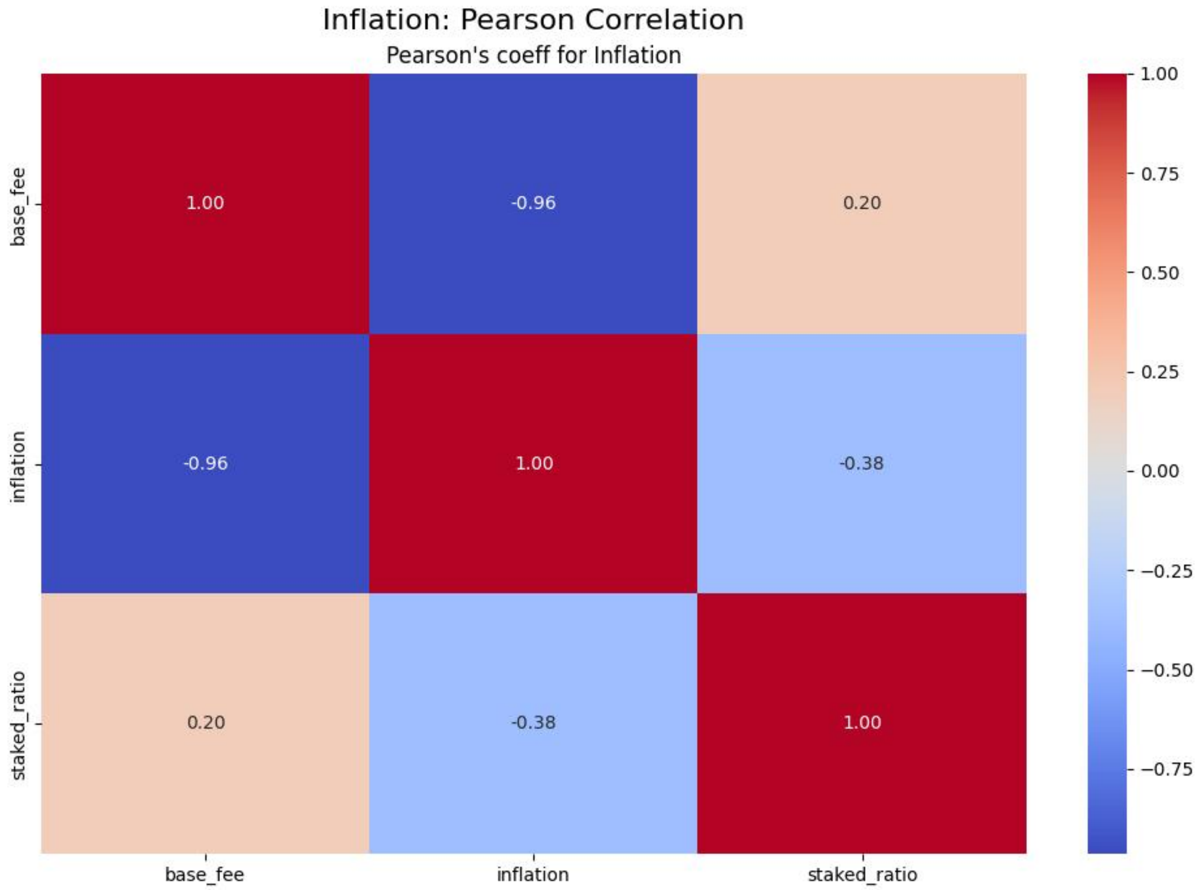

Figure 3.17: Inflation: Pearson Correlation

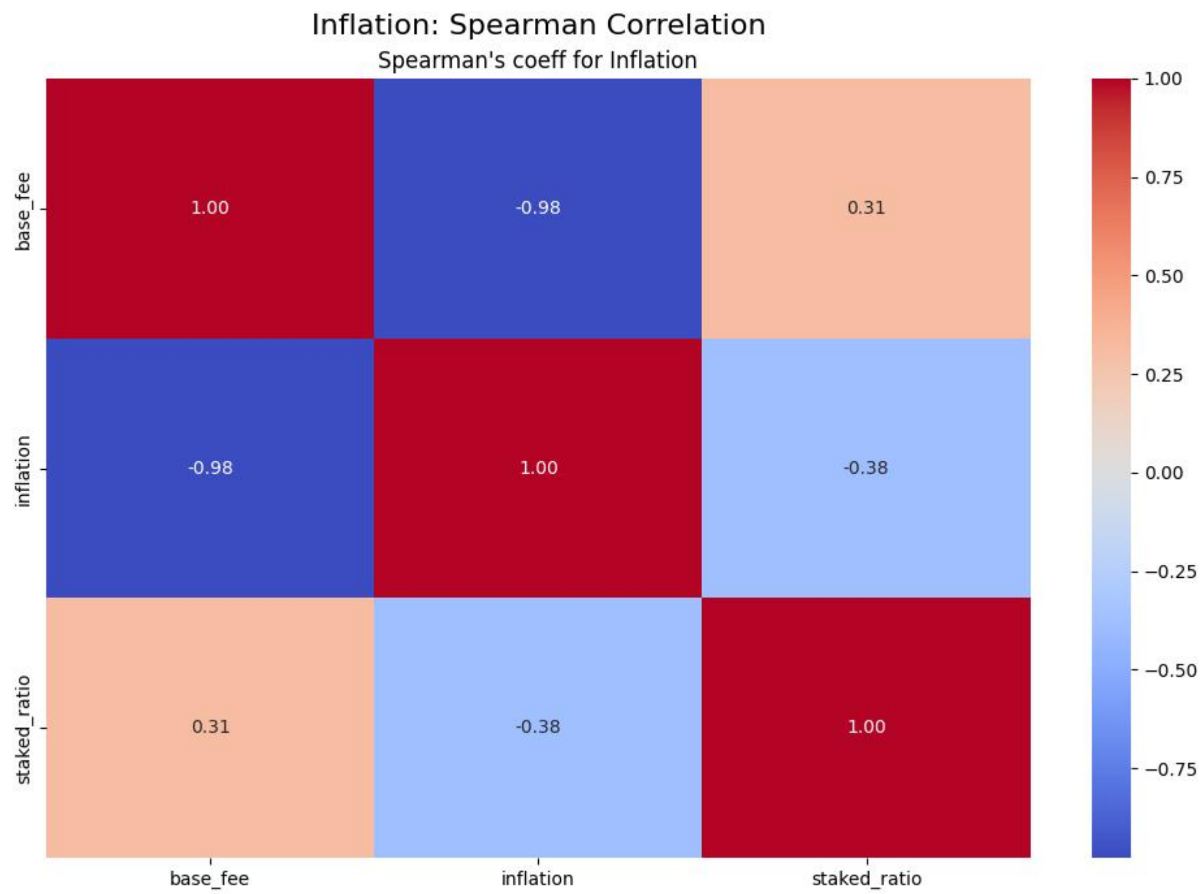

Figure 3.18: Inflation: Spearman Correlation

decreases, and vice versa.

Both correlations between staked ratio and inflation is approximately -0.38, indicating a moderate negative correlation. This suggests that as the staked ratio increases, inflation tends to decrease, and vice versa, although the relationship is not very strong.

The case analyzed starts with an initial staked ratio of 0.5 which tends to stabilize at the optimal 0.66. Moreover, only values above the critical value of 0.54 have an impact on inflation. On the other hand, the base fee starts at 1e9 EVMOS and reaches a value of almost 2e9 EVMOS, which means that duplicates its value. The difference in the correlation of each factor against inflation, is explained by this distincts magnitudes of changes.

Both coefficients shows the high correlation between inflation and the variables base fee and staked ratio. This makes sense, as the issuance design takes into account the demand of the network expressed in the base fee variable and the portion of the total supply that is being used for validations, staked ratio.

Chapter 4

New Weights Selection

Initial Theoretical Weights

Initially, we justified the selection of the weights $\omega = (\omega_s, \omega_b)$ through mathematical and statistical analysis. The idea was to minimize the average annual error of being away from the boundaries r_i and r_s for the staked ratio s and the base fee. This was done by simulating 50 market scenarios, each with 30 daily realizations of s and the base fee, searching for weights that would keep the daily issuance within the desired range for each scenario.

Improvement through Experimentation

When experimenting with the agent-based model and stressing the weights by exposing them to a staked ratio that evolves over an entire year, we observed that these theoretical weights were not optimal for keeping the issuance within the desired range. We found that dividing the theoretical weights by 2 produced a more consistent and dynamic issuance in the model. This suggests that the theoretical weights can be improved to better fit the system dynamics in the context of a year-long simulation.

Conclusion

In summary, the initial theoretical study allowed us to propose a consistent model based on minimizing the annual error at the issuance boundaries. However, through experimentation and parameter variation, we have found that adjusting the theoretical weights, by dividing them by 2, significantly improves the dynamics of the issuance function in the agent-based model. This suggests that while theory gave us a solid foundation, practice has led us to optimize the weights for a more dynamic and consistent issuance over a longer time horizon.

Chapter 5

Machine Learning

Up to now, we have been analyzing the behavior of validators, whose long-term effectiveness remains unclear. However, the proposed action profile for validators, mentioned in 2.3.4, appears to be appropriate if we aim to model an agent seeking immediate rewards. In order to avoid biasing the behavior of agents, we have developed an innovative approach to explore new strategies using Machine Learning.

Our solution involves implementing a machine learning algorithm, specifically the Q-Learning algorithm. Unlike our approach, where agents take predefined actions at each step of the model, the Q-Learning algorithm grants them the freedom to select any action at any time, without restrictions.

Not biasing the Q-Learning algorithm with specific actions allows agents to discover optimal strategies on their own, which may not have been considered previously. This approach not only increases the flexibility and adaptability of our agents but could also lead to improvements in the overall performance of the model.

In this chapter, we will focus on determining whether the Q-Learning algorithm achieves improved performance of the agents compared to the previously established approach. We will analyze the results obtained through the implementation of the Q-Learning algorithm and evaluate its effectiveness in enhancing the performance of our agents in the model.

Reinforcement Learning

The Q-Learning algorithm, which falls under the category of model-free methods, is particularly useful in our context. Model-free algorithms, like Q-Learning, learn the consequences of their actions through experience without the need for explicit transition and reward functions. This flexibility is advantageous because it allows our agents to adapt and discover optimal strategies on their own.

The value-based method, employed by Q-Learning, trains a value function to learn the value of each state and choose the best action. This approach is beneficial because it enables our agents to understand which actions are more valuable in terms of achieving their not necessarily immediate objectives.

Additionally, Q-Learning is an off-policy algorithm. This means it evaluates and updates a policy that differs from the policy used to take actions. This feature allows our agents to explore and learn from various strategies, even if they are not currently the most optimal. This can be advantageous in our model, as it enables our agents to consider a wider range of actions and potentially discover better solutions.

A key point to highlight is that during the validators' training phase, which consists of 30 episodes, it is carried out with 25% of Q-Learning (QL) validators distributed equally among those with low, medium, and high initial stakes. The remaining validators follow the aforementioned logic. Another significant aspect is that these training data are updated simultaneously by all QL agents, in what is known as multi-agent training. This means that each QL agent not only learns from its own experience but also benefits from collective learning, which enhances the efficiency and adaptability of the system.

In summary, we are using the Q-Learning algorithm because it is a model-free, value-based, and off-policy method. This allows our agents to learn and adapt to the dynamics of the environment, discover optimal strategies, and explore a wide range of actions to potentially improve the overall performance of our model.

Performance Comparison

In this section, we will compare the performance of conventional validators against validators trained using Q-Learning (referred to hereafter as "QL validators"). To conduct this comparison, we introduce 25% of QL validators into the model, evenly distributed among small, medium, and large validators. This approach allows us to observe the behavior of QL validators based on their size.

Upon comparing the performance of Q-optimal validators with the risk-profile reward model, we have observed a notable similarity in their overall performance. However, the risk-profile reward model consistently outperforms the Q-optimal validators in terms of reward across nearly all days of the year, as shown in Figure 5.1. This observed superiority could potentially be attributed to the necessity for training to occur over extended time horizons for non-QL validators to exhibit distinct trends, thereby enabling QL validators to glean valuable insights from these behaviors. Conversely, the consistent dynamics we observe indicate that QL validators tend to mimic the actions of those seeking immediate rewards, all while considering their respective risk profiles. This is expected as the minting function, which is the main source of rewards for both QL and non-QL agents, has not a long-term evolutionary behaviour but instead it dynamically reacts to the last state of the system.

It is worth emphasizing that our analysis also reveals strikingly similar trends in inflation and decentralization across both models similar to those observed in the demand scnario analysis, Figure 5.2. This consistency in results bolsters the reliability and robustness of the immediate reward model.

Conclusion

In conclusion, the examination of a scenario where validators learn from others and are not solely dependent on predefined immediate target rewards presents intriguing insights. Our findings suggest that over the course of a one-year period, the risk-profile reward model consistently yields higher target rewards compared to the Q-optimal approach. This underscores the importance of incorporating risk profiles into reward models to achieve more optimal outcomes.

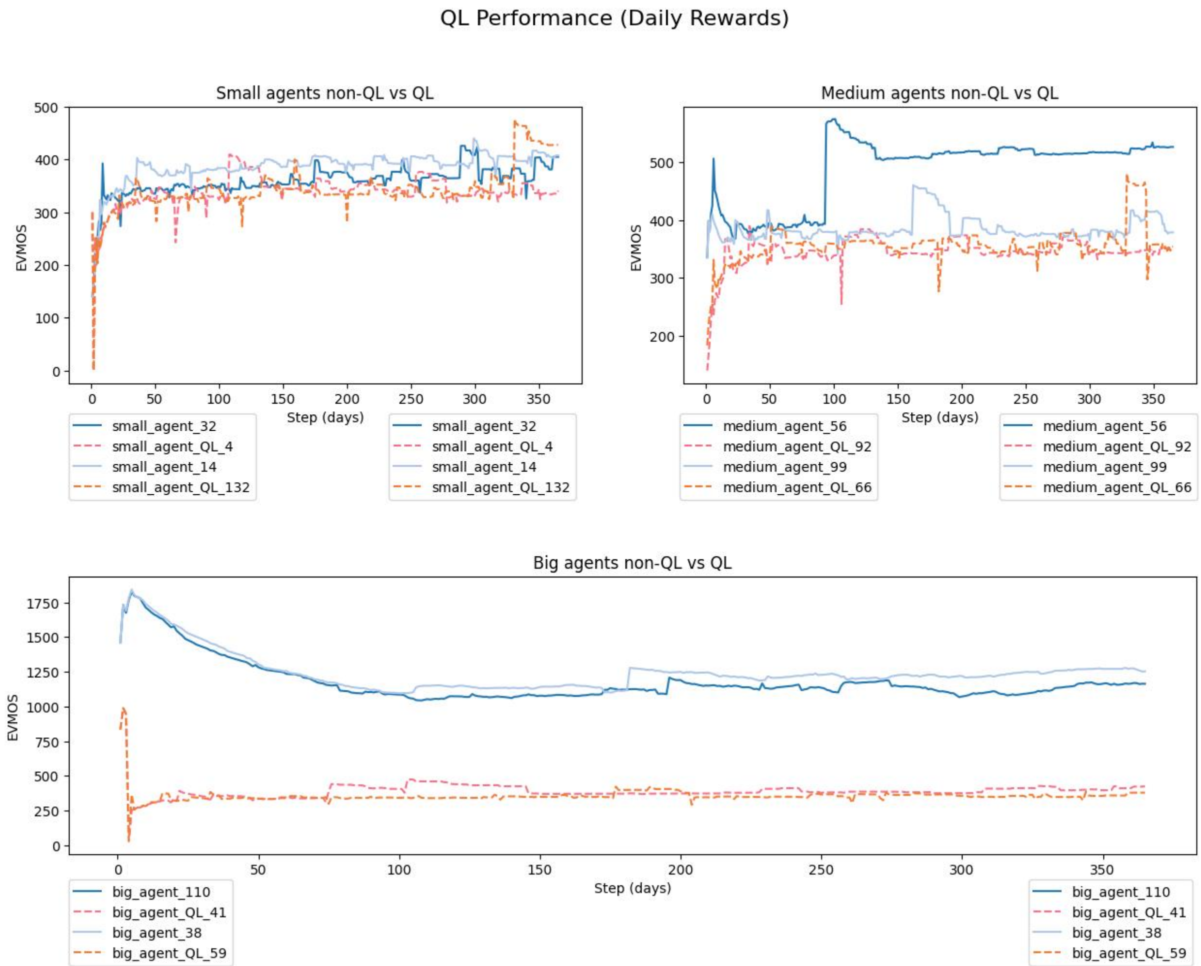

Figure 5.1: Daily Rewards: Non-QL Validators vs Validators QL

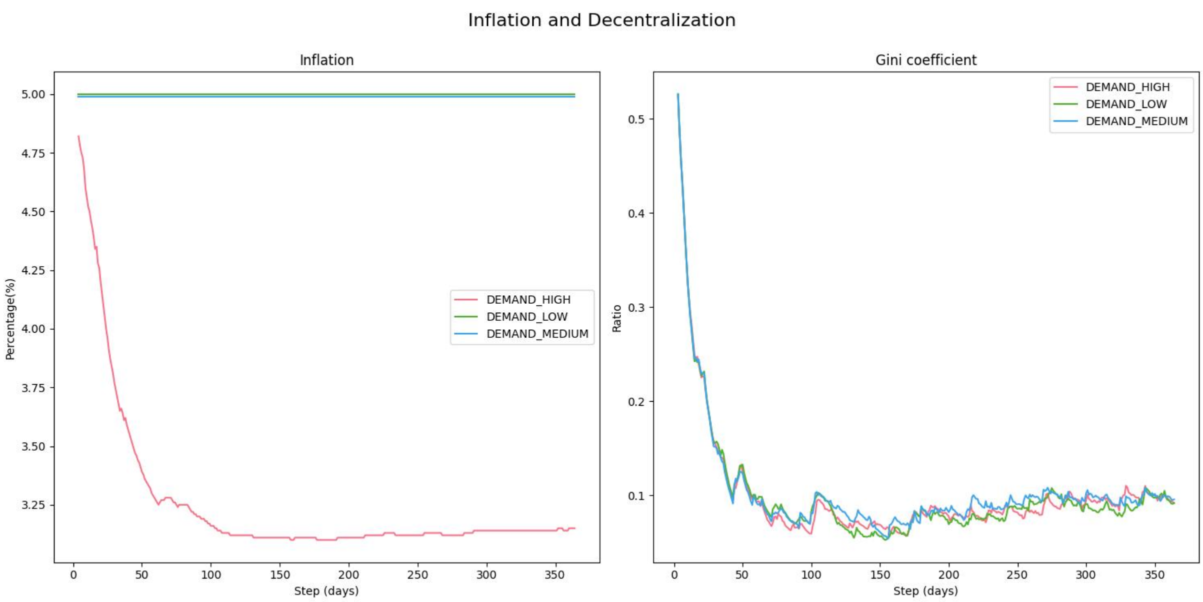

Figure 5.2: Inflation and Gini coeff with QL agents

Chapter 6

Conclusion

Concluding our comprehensive analysis and simulation-based evaluation of the redesigned token economy for the Evmos chain, the results underscore a significant move towards achieving a more sustainable, equitable, and decentralized ecosystem. Through meticulous adjustments in token issuance, reward distribution, and strategic modifications to validator dynamics, we've observed a tangible reduction in inflation rates, from approximately 30% in the first version of the tokenomics to around 4% in the current design, illustrating the effectiveness of our tokenomic interventions.

Furthermore, the decentralization metrics, particularly the drop in the Gini coefficient from 0.7 to values oscillating between 0.1 and 0.2, highlight the advancement towards a more balanced and inclusive network. This shift is further evidenced by the increased stake contributions from smaller and medium-sized agents, surpassing those of larger agents, which indicates a successful dispersion of power and influence within the network.

The behavior of validators, influenced by variations in risk and optimality profiles, has shown that even under conditions of increased volatility and price sensitivity, the system trends towards our desired outcomes of reduced inflation and enhanced decentralization. This is a testament to the resilience and adaptability of the proposed tokenomic design.

Additionally, the examination of a scenario where validators learn from others and are not solely dependent on predefined immediate target rewards presents intriguing insights. Our findings suggest that over the course of a one-year period, the risk-profile reward model consistently yields higher target rewards compared to the Q-optimal approach. This underscores the importance of incorporating risk profiles into reward models to achieve more optimal outcomes.

In conclusion, the annual simulation has confirmed that strategic adjustments in the token economy can effectively steer the Evmos chain towards its goals of reduced inflation and increased decentralization. These findings not only validate the proposed design changes but also provide a robust framework for future adjustments and optimizations. As the Evmos chain continues to evolve, the insights gained from this analysis will serve as a cornerstone for fostering a thriving, decentralized ecosystem, paving the way for sustainable growth and long-term viability of the network.