Wealth Inequality in CasperLabs' Proof-of-Stake Blockchain

Master's thesis

University of Zurich

Department of Informatics

Blockchain and Distributed Ledger Technologies

Professor:	Prof. Dr. Claudio J. Tessone
Author:	Thomas Levrand
Matriculation number:	14-721-542
E-Mail:	thomas.levrand@uzh.ch
Address:	Im Hundsrücken 2
	8600 Dübendorf
Major:	Business Administration
Minor:	Data Science
Date of Submission:	15.08.2022

Abbreviations

APY	Annual percentage yield
BFT	Byzantine fault tolerance
CBC	Correct-by-Construction
PoW	Proof-of-Work
PoS	Proof-of-Stake
ROI	Return on Investment

List of Figures

List of Tables

Abstract

The mining of cryptocurrencies such as Bitcoin, or more specifically the Proof-of-Work (PoW) consensus algorithm behind it, has come under criticism for its high energy consumption. In the increasing spotlight is the alternative consensus algorithm Proof-of-Stake (PoS), which only consumes a fraction of the power during operation. Regardless of the consensus algorithm, decentralization and a certain degree of fairness in the distribution of incentives are considered important cornerstones of a blockchain. While a body of research has already theoretically and empirically studied large PoW blockchains such as Bitcoin and Ethereum in terms of decentralization and fairness, corresponding research for PoS blockchains is rather scarce. This paper builds on the existing research and derives a set of appropriate measures that can be used to characterize the properties of the Casper PoS blockchain in terms of wealth inequality, fairness of reward distribution, and decentralization. Subsequently, based on real-world data extracted from the Casper blockchain, a comprehensive exploratory analysis is conducted and the defined decentralization and fairness measures are computed. Furthermore, an agent-based model is constructed based on the characteristics of the Casper blockchain and other common PoS blockchains, which is then used to simulate the respective reward mechanisms. Through the data analysis and the applied measures, it was found that the Casper blockchain exhibits high values of wealth and income inequality. The decentralization measures also indicate a high level of centralization, especially at the beginning, but a constant increase in decentralization can be observed over time. Moreover, high expectational fairness values were measured, indicating that participants largely receive fair rewards according to their input (staking balance). The simulations carried out indicate that the reward mechanism of the Casper blockchain displays weaker wealth compounding effects compared to common PoS blockchains and thus is less conducive to the concentration of wealth.

1 Introduction

1.1 Motivation

Blockchain is considered a trustless technology because it eliminates the need for trust between two parties. However, in order for participants to be confident in the blockchain system and willing to participate to maintain the system, enabling fairness is a fundamental requirement. Regardless of which protocol a blockchain is based on, a large number of participants is essential both for security and decentralization.

The consensus mechanism, which makes it possible to validate blocks and thus transactions without a central authority, is crucial for a blockchain to function in a decentralized manner. The most widespread blockchains are those based on PoW, which achieves this consensus by requiring participants to successfully solve cryptographic puzzles as a condition to be allowed to validate a block. However, as this method consumes increasingly large amounts of computing power and energy (Digiconomist, 2022), the waste of resources of PoW is a major point of criticism of blockchain technology, especially in view of increasing environmental awareness.

The alternative consensus mechanism PoS is becoming increasingly popular precisely because of its higher energy efficiency. In its original form, the right in PoS to validate the next block and thus to receive the block reward is based on the tokens that are already held (King & Nadal, 2012). Thus, higher income is not achieved by the participant who spends more computing power, but by the one who already owns more tokens. While this largely solves the problem of power consumption, it creates a new problem: the so-called “rich get richer” effect describes the phenomenon that wealth becomes increasingly concentrated in PoS networks, as high income in one period leads to even higher income in the next period due to compounding effects.

One of these PoS networks that has been attracting more attention lately is Casper, a network that has been optimized for more sustainability due to its new consensus mechanism, allowing the network to outperform Ethereum (before the switch to PoS) and Bitcoin while consuming only a fraction of the energy. A recent power usage study conducted by developers at CasperLabs (CasperLabs, 2021) claims that the Casper network in its current state with 100 nodes is 191400% and 545400% more energy efficient when compared Ethereum to and Bitcoin, respectively. Even with 400 nodes in the future, the Casper network requires 47000% less energy than Ethereum and 136000% less than

Bitcoin. Moreover, in a direct comparison with other PoS networks such as Cosmos, the Casper network performs twice as effectively in terms of energy consumption (Cosmos, 2021).

Previous studies have addressed the implications of general PoS consensus protocols and different reward functions on income and wealth distribution (Dimitri, 2021; Fanti et al., 2019; Roşu & Saleh, 2021). In contrast, little research has examined how wealth distribution behaves within a real-world cryptocurrency with its specific PoS implementation. However, as the respective PoS system implementations differ, distinct monetary dynamics may be the result, making wealth and income concentrations stronger or weaker depending on system characteristics.

1.2 Purpose

The research presented in this thesis aims to investigate the wealth inequality and related fairness and decentralization measures in the Casper blockchain. Previous research focused on differences in reward functions and their effects on wealth distribution, but often did not distinguish between the implemented reward distribution mechanism. Therefore, this paper analyzes the impact of the Casper reward distribution mechanism on wealth inequality compared to this of a basic PoS protocol. Furthermore, this paper addresses the research question of how wealth and income distribution evolve over time in the Casper network, which is measured by the inequality coefficient (Gini). The decentralization of a blockchain is seen as its core concept, which often raises the question of whether this is actually achieved. In addition, especially in a PoS system, the power in the system is closely linked to the concentration of wealth, which closely connects wealth inequality and decentralization. For this reason, the degree of decentralization and fairness in Casper is measured by three additional metrics (Nakamoto coefficient, expectational fairness, and information entropy).

The thesis thus aims to contribute to the state of research to date by offering insights into a blockchain that has not yet been studied for wealth inequality, decentralization and fairness, and by examining which individual characteristics of the blockchain influence these values.

1.3 Research Design

This research begins with a literature review that establishes a theoretical foundation for developing the agent-based models and conducting analysis of real-world data of the Casper blockchain. The selected literature consists of empirical research that has investigated fairness in different consensus mechanisms. For the study of reward distribution mechanism, several simulations are conducted that compare the impact of the Casper reward mechanism on inequality in comparison with a basic PoS reward mechanism. Subsequently, the data obtained from the Casper blockchain is analyzed using a set of measures that have been identified in the literature review.

This thesis is structured as follows: First, section 2 provides the reader with a theoretical basis and gives an overview of relevant theories and research findings on inequality and fairness in blockchains. In section 3, the existing literature is put into context in more detail and the research questions addressed in this paper are elaborated. Section 4 discusses the simulation and presents its results. The next section contains information on the data acquisition, descriptive and exploratory analysis and the results of the fairness measures with the data obtained from the Casper blockchain. In a next step, section 6 discusses and critically comments on the results with suggestions for future research. Section 7 consists of concluding remarks.

2 Theoretical Foundation

2.1 Proof-of-Work vs. Proof-of-Stake

The blockchain is a decentralized ledger in which new blocks of information are continuously created and then backed up by a network of decentralized nodes. Because the blockchain does not involve a central party that decides whether the information is correct and which block should be added to the database, alternative methods are used to validate blocks. The two most commonly used consensus mechanisms are the PoW and PoS mechanisms.

In PoW blockchains, the participants, usually called miners, use their computing power to solve a cryptographic puzzle and thus obtain the right to add new blocks to the blockchain. In order to find a valid block, a hash must be determined, which must meet a requirement specified in the network protocol. The correct hash value is calculated using a combination of the hash values of previous blocks, the recent transactions to be added to the new block, and a “nonce”. The level of difficulty is defined in the protocol and specifies the number of zeros the hash must start with. The miner must find the nonce (short for "number used once"), which can only be done by brute force, therefore the probability of finding the solution is, a priori, proportional to the computational power each miner contributes per time unit. The individual miners are in competition with each other and only the first to solve the cryptographic puzzle will receive his reward in the respective blockchain currency. The only way to raise the probability of solving the puzzle is to increase the computing power, which requires the use of very expensive and specialized hardware (ASICs) as well as miners joining large mining pools.

In the battle for more computing power, a competition has broken out between miners to find the fastest hardware and the cheapest power source. As a result, the search for the cheapest source of power has led to increased centralization of mining pools in some countries and regions. How centralized the Bitcoin network really was became apparent after a government-led crackdown in China effectively banished all crypto miners - with more than 50% of the hashrate dropping out of the network (MacKenzie, 2021). The massive consumption of computing power that must be expended to solve the cryptographic puzzle cannot be used for other purposes, such as disease research. The wasteful use of energy that PoW requires to ensure the security of the system is becoming the focus of political discussion in light of impending global warming.

Alternative consensus mechanisms have been developed that are far less energy intensive. This has resulted in an ever-increasing number of cryptocurrencies relying on PoS, or, as in the case of Ethereum, plans to switch to it. The PoS consensus mechanism was first introduced by King and Nadal (2012). In its pure form, the right to validate the next block depends on the number of currency units of the respective cryptocurrency held and the duration they have been held. The participants in a PoS blockchain are called validators, bakers or block proposers, depending on the blockchain. To be eligible as a validator, an amount of native cryptocurrency must be staked on the network. In order to receive the reward, a validator must participate in the block finalization process, and his tokens are locked in the network as a sort of security deposit. If a validator would be caught in attacks like malicious forking or double spending, his coins could be confiscated as a punishment.

A heavily debated topic in the blockchain community is the concern that PoS protocols could potentially contribute towards making wealthier nodes even wealthier. In addition, this rich-get-richer effect (also known as the Matthew effect) not only is unfair to the less wealthier validators involved, but could also pose a threat to the decentralization of the network. Unlike PoW blockchains, the rewards from maintaining the blockchain can be added to the validator's staking balance account almost instantaneously upon receipt. Moreover, the probability of finalizing a block depends on the proportional share of total assets staked, hence for the next blocks the validator then has a higher probability of finalizing those blocks and earning the corresponding rewards.

A growing body of literature examines the economic dynamics of PoS platforms. In particular, Roşu and Saleh (2021) and Saleh (2021) question which role rewards play for the concentration of wealth in PoS blockchains. Besides, Fanti et al. (2019) and Wang et al. (2020) investigate the long-term behavior that reward functions have on users' cryptocurrency distribution. According to these studies, there appears to be a relationship between the degree of asset compounding and the reward function implemented in the consensus protocol.

2.2 Selfish Behavior

In 2014, Eyal and Sirer identified a mining attack strategy in PoW blockchains. In so-called selfish mining, the miner (or mining pool) does not publish the mined block immediately, but only selectively. The goal of the selfish mining pool is to waste the computing power of honest miners. Li et al. (2020a) showed in their study that mining pools do indeed use selfish mining strategies. Furthermore, they found the existence of mining cartels in Bitcoin. Due to the fact that cryptocurrencies are maintained by distributed consensus, while there is no trust between miners, ensuring fairness is essential. Consequently, Li et al. (2020b) investigated mining fairness in popular PoW blockchains, using inequality indices such as the Gini index. All PoW blockchains studied showed a very uneven distribution of revenue, from which the authors concluded that the sometimes persistently abnormally high success rates of some mining pools are due to the fact that they pursue a selfish mining strategy.

Similarly, in PoS blockchains, there are circumstances in which some validators may benefit from violating the protocol by behaving dishonestly. In their work, Neuder et al. (2020) use the Tezos PoS protocol as an example to investigate which incentives encourage rational participants to engage in selfish behavior. In a PoS blockchain, if several validators agree to form a validator cartel, they can decide which members to blacklist. Censoring the blacklisted members results in their messages being ignored, which in the eyes of the consensus protocol looks like they never sent them. The censored members cannot prove to the protocol that they have done their work and, depending on the project, may not receive a reward or even be punished for their absence.

2.3 Proof-of-Stake Block Reward Functions

In order to ensure the security of a PoS blockchain, incentives for validators must be set in such a way that both large and small participants are encouraged to take part. This can only be achieved if the process of reward distribution is fair and even the smallest participants can expect their reward on average. Important components of a blockchain that can also significantly affect the fairness of reward distribution is the design of the block reward function and the reward distribution mechanism that decides on distribution of the corresponding rewards to participants. For a fair compensation, not only the adequate average compensation is of importance, but also the variance, in other words, the uncertainty with which the payout takes place. A major concern in PoS systems is the rich get richer effect, which can result in a dramatic concentration of wealth. This phenomenon is also known as compounding. Whenever a node receives a reward, it is immediately added to the node's staking account, increasing the likelihood that the node will be selected as a block proposer in the next round and receives the associated reward. The compounding of wealth can lead to a very unfair distribution in which smaller participants are disadvantaged. Both the block reward function and the reward distribution mechanism therefore play a crucial role in the degree of unequal distribution.

In their paper, Fanti et al. (2019) studied the concentration of wealth in PoS systems resulting from different block reward functions. In an experiment they constructed, the effects of constant, geometric, and decreasing reward functions on long-term reward distribution were observed. In the PoS systems studied, a validator is chosen as the block proposer with a probability proportional to his fractional stake. Once elected, the validator receives the total reward for that time period. Thus, several block reward functions with one type of a reward distribution mechanism were investigated. They designed a closed blockchain for their experiment, meaning that no participant could exit or enter the system. They thus ensured that all participants received their average expected value of rewards, and the final variation in the reward distribution could thus be explained by the differences in the variance that arose from the three reward functions.

For the majority of cryptocurrencies today, a constant block reward function is used, whereby the reward remains the same over a certain period of time (e.g. years). The block reward is intended to compensate the validators for the cost of running the node and proposing new blocks. Many cryptocurrencies follow the maxim that in the short term each block should generate a similar reward, but it is not specified whether this value is measured in fiat or tokens. However, most cryptocurrencies measure their reward value in their native tokens.

When a blockchain is put into operation for the first time, there are usually only a few participants in the network, with a constant reward function, these few receive a relatively large reward. The geometric function tries to address the issue of wealth compounding by initially distributing smaller rewards. With the geometric reward function, only a small reward is paid out at the beginning, which rises over time and with the increasing number of participants. This distribution is intended to help ensure that the rewards are distributed more evenly.

Some cryptocurrencies, such as Monero, use a third type of reward function to dispose of their block rewards, known as decreasing reward function. In this case, fewer and fewer tokens are paid out over time.

Three different scenarios were examined: first, all participants behave honestly; second, participants violate the protocol; and finally, the existence of mining pools. Equitability is measured through the variance of each of the functions. The smaller the variance, the lower the uncertainty and the greater the equality of the final reward distribution. The authors conclude that the geometric function is notably superior to the other two. It performs better in equitability and could even avoid selfish mining attacks.

The study by Roşu and Saleh (2021) replicates the simulation of the geometric reward function by Fanti et al. (2019) and arrives at the same result. However, the authors go further and examine the wealth distributions of the reward functions over a longer period of time. In the long run, the simulations paint a different picture, with the geometric reward function leading to a more unequal distribution than the constant function .

Furthermore, Wang et al. (2020) show that the reward function alone is not sufficient to attract enough participants through a fair distribution. Another important ingredient is the incentives that must be offered so that the parties are willing to maintain the consensus mechanism. However, the determination of fairness in a blockchain is not only dependent on the design of the reward function, but other factors such as the specific reward distribution mechanism or the penalization of violations of the protocol could have a significant influence (Fanti et al., 2019; Wang et al., 2020).

2.4 Wealth Inequality Measures

In order to evaluate fairness in blockchains and to investigate the income inequality of rewards, other studies have often relied on measures from the field of economics. While most research has focused on fairness in PoW consensus mechanisms, research on PoS consensus mechanisms is scarce. In the following, the measures used in this thesis will be explained in more detail.

2.4.1 Gini Index

The Gini coefficient is one of the most commonly used statistical constructs to numerically determine the spread of wealth. It is defined as the ratio of the area under the line of equality to the area above the Lorenz curve. The Gini value ranges from 0 to 1. A Gini coefficient close to 0 would imply that the wealth distribution of a country is perfect, thus the Lorenz curve would be exactly the line of equality. On the other hand, a Gini coefficient of 1 means that all wealth is controlled by a single individual. The Gini coefficient can be calculated as follows:

$$A = \frac{n \sum_{i=1}^{n} x_i}{2}$$

$$B = \sum_{i=1}^{n} i\, x_i$$

$$Gini = \frac{A - B}{A}$$

To calculate the Gini of a country, it is assumed that the population within this geographic community spends most of its time and thus also uses its resources locally. This results in a Gini coefficient for a country or geographic community that reflects the inequality in total resources available to the population. Unlike a country, the inequality in the cryptocurrency community can originate from two sources. First, the amount invested in cryptocurrencies depends on the total resources available to the participants. A second important factor is the interest a participant has in the project and the participation in the community, which determines how much he is willing to invest (Buterin, 2021). A participant who has only a low staking balance in a blockchain and could therefore be considered poor does not necessarily suffer from limited resources, but may simply not have a particularly high level of interest. Consequently, the Gini coefficient of a cryptocurrency is not comparable to that of a country.

In recent years, the Gini coefficient has been used in a growing number of studies, especially in the context of blockchain, and is a relatively simple and easy-to-understand statistic. However, the Gini index alone cannot answer questions about the distribution of participants in the network and whether the concentration of power is high or not. As a result, researchers in the field have incorporated additional measures and developed alternative metrics explicitly designed to study inequality in blockchain networks. In order to provide a holistic picture of the Casper network, further measures are included alongside the Gini coefficient.

2.4.2 Nakamoto Coefficient

The Nakamoto coefficient is a measure that quantifies the degree of decentralization in a blockchain and is thus directly associated with its security. The coefficient is defined as the minimum number of participants that are needed to collectively obtain more than 51% of the total mining power to control the blockchain (Srinivasan & Lee, 2017). While the operational threshold for Bitcoin, as in most other PoW blockchains, is 51%, it may differ for other systems. In PoS blockchains that are Byzantine fault tolerant (BFT) and have a binary threshold, such as Algorand or Tendermint, more than 33.33% of the total stakes are needed to prevent consensus (Gilad et al., 2017; Tendermint, n.d.). Even though the threshold in Casper is not binary but instead represented by a number, it still requires a majority of 66.67% of all stakes to reach consensus. The Nakamoto coefficient can be modified according to the required threshold of a blockchain in order to determine the minimum number of participants whose combined propositions are sufficient to control the system. A higher Nakamoto coefficient translates into more decentralization and the need for more validators or staking pools to work together in order to gain control. The Nakamoto coefficient can be calculated as follows:

$$N = \min\{\, k \in [1, \cdots, K] : \sum_{i=1}^{k} p_i \geq X \,\}$$

Where X indicates the threshold of the respective network, in the Casper network it follows that:

$$p_i \geq 0.3333$$

2.4.3 Expectational Fairness

Given that fairness plays a very important role in the design of incentive mechanisms within a blockchain, the measure of expectational fairness was designed to examine fairness in different blockchain systems (Huang et al., 2021). In a blockchain system that would be perfectly fair, every participant could expect the same Return on Investment (ROI). An incentive mechanism preserves expectational fairness if the expected amount of the reward λ_a of a miner A is equal to the proportional amount of its resources a (for example, in PoS the amount of tokens placed in the network, in PoW the provided computing power) used and therefore $E[\lambda_a] = a$ holds.

2.4.4 Information Entropy

Entropy is a physical concept from thermodynamics that measures the degree of chaos or disorder within a system. Jia et al. (2022) introduced this measure to determine decentralization in different PoS systems. Derived from the thermodynamic concept of entropy, the information entropy describes how much event information is transmitted on average. A smaller entropy implies that the system is more orderly and contains more information. In contrast, the larger the entropy, the more disorganized and chaotic the system, and the less information it contains. Extended to the blockchain scenario, the information entropy can thus capture the randomness and chaos of the distribution of tokens. The information entropy is calculated as follows:

$$p_i = \frac{x_i}{\sum_{i=1}^{n} x_i}$$

$$Entropy = -\sum_{i=1}^{n} p_i \log_2 p_i$$

2.5 Casper Network

The Casper network is a layer one PoS blockchain, that claims to be fully decentralized and designed to help companies develop their blockchain-enabled products and services effectively and quickly. Casper is the first live implementation of correct-by-construction (CBC) Casper consensus protocol, and, according to its own statements, is trying to achieve higher performance without sacrificing too much in terms of decentralization, scalability, or security. The native token to the Casper network is called CSPR. As a PoS blockchain, Casper uses its CSPR to pay out rewards to its validators for carrying out on-chain transactions via Casper's PoS consensus mechanism. Furthermore, network participants use their CSPR to interact with Casper's services and features and to pay fees. Additionally, it plays a vital role in the community governance of the network.

The current consensus protocol of the Casper network is called Highway and extends Vlad Zamfir's research on the original CBC Casper family protocol (2017). The Casper Highway protocol is a secure and live consensus model in the sense of a classic BFT consensus protocol, but offers two significant advantages for the network. Distinguishing characteristics of this protocol are its flexibility and finality (Kane et al., 2021). First of all, Casper is flexible, each validator can choose its own finality threshold, this allows them to optimize for different roles in the network. Unlike other protocols, even if for some reason many nodes crash and, for example, only sixty percent of the nodes are still online, the protocol can continue to run for a while with the remaining finality threshold. Second, the Highway protocol allows the network to achieve higher finality thresholds. Finality guarantees the immutability and irreversibility of a transaction once it has been propagated, increasing the certainty and trust in the network. Compared to classical BFT PoS models, the block finality is no longer binary but rather expressed by a number. Additionally, the Highway protocol allows more than 67% of the nodes to be considered honest, thus higher finality thresholds of, for example, 80% can be achieved.

As a server operator, a validator must operate and maintain a Casper node. In order to stake tokens, participants do not need to set up a Casper node themselves, instead they can delegate their tokens to a validator, they are therefore referred to as delegators. Usually, delegators are charged a fee by the validator for staking, the amount is determined by each validator individually and typically equals a few percent of the reward. Both validators and delegators will receive a reward that corresponds to their proportional share of total staked tokens. In addition, the validator who was elected as block proposer will receive the transaction fees of the respective block. However, since transaction fees on the Casper network are extremely low, these additional rewards are negligible. The rewards that validators and delegators receive for participating in the consensus and finalizing blocks are termed seigniorage. To pay out the rewards, new tokens are minted and transferred to the validators and associated delegators. The base annual reward rate is 8% of the total supply. Due to the fact that only a part of the total supply is currently staked, the current annual percentage yield (APY) ranges from 11 to 20 percent. Currently there is no slashing enabled in the Casper mainnet. Slashing means that, in the event of a validator's misconduct, both his own staked tokens and all staked tokens delegated to him can be taken away as a penalty.

In the Casper network, the rewards are not paid out per block but once per era. An era is currently set to 2 hours and originally a block was set to 64 seconds, which means that an average of 112 blocks per era could be expected. Due to an update in April this year, the block time was halved, resulting in an average of 225 blocks per era. The number of validators in the Casper network is limited to 100, and the 100 biggest nodes - measured by the size of their total staking balances - are selected. The reward function is constant and linearly increasing total rewards are expected per era.

As long as the violation of the protocol is not severely punished, and if there is the chance to increase profit, it is only a matter of time until the first validators decide to form a cartel. In other projects, if validators do not send a message in the time they are supposed to, they will not be paid any rewards or may even be penalized for their absence. This could be a false incentive for reward distribution, as the validator is not necessarily the one to blame, but rather a cartel may have a hand in it. The cartel can make the validator appear offline for the protocol by censoring that validator. The cartel ignores the messages sent and the validator has no way to defend himself against the ongoing censorship. As a consequence of such incentive mechanisms, small validators may decide to not to join the network. The Highway protocol (CasperLabs, 2020) takes a different approach in the design of its reward distribution compared to other projects, in order not to create incentives that contribute to cartelization in the first place. Therefore, the seigniorage paid in the Casper network is independent of who proposed the block. As a result, in the Casper network, the reward-per-staked-token is the same for each validator. In game theory, this concept is known as payoff symmetry and prevents smaller validators from receiving a smaller reward or even no reward at all compared to larger validators due to censorship.

In addition to payoff symmetry, a reward correction has been implemented. For this purpose, the total rewards paid out are adjusted to the proportion of participation in the network. In the case that only 90 percent of all validators send a message to finalize a block, the actual block reward *R* is corrected downwards and amounts to only 0.9 * *R*. Since all participants are in the same boat and everyone in this case is harmed by a lower payout, the only rational decision for a majority is not to censor. Without an increasing payout resulting from censorship, there is theoretically no reason to form a cartel.

2.5.1 Main Updates

In the following the most important updates of the Casper network will briefly be introduced, a table with further updates and references can be found in the Appendix A.

Table 1 Caper's main updates and outages

Update	Time	Era	Notes
1.2.0	28/05/2021	694	Additional consensus and block proposer log messages were added Block validators now distinguish between transfers and deploys
1.2.1	13/07/2021	1281	Security updates
1.3.2	12/08/2021	1605	Using validator keys the network verifies the node as validator and prioritizes them over others Limiting outgoing traffic per node to a fixed amount Validators that were inactive in an era are prevented from proposing blocks in the next, to improve liveness by reducing missed proposals
Outage	03/11/2021	2595	Total outage of 4 hours and 21 minutes
1.4.1	03/11/2021	2600	Faster network block time: improved expected block times from 64 seconds to around 32 sec. Equivocator enhancement: equivocators are now added to equivocators list in the block, they are banned from participating in upcoming eras, this adds to resilience
Outage	22/11/2021	2819	Total outage of 9 hours and 45minutes
1.4.3	17/12/2021	3111	Missed block improvements Latency improvements
1.4.5	04/04/2022	4417	Support DeFi products including several bridges, NFT marketplaces, decentralized exchanges, automated market makers, yield farming and collateralized lending application

			Limit number of delegators per validator to 953 temporarily
1.4.6	20/05/2022	4968	Significant drop in disk usage growth and disk bandwidth

3 Research Goals

With their work, Huang et al. (2021) laid a first foundation and investigated the fairness of incentive models for PoW and PoS protocols with the objective of getting to the bottom of the question if the rich get richer. To characterize the relationship between the resources used by a validator and his reward, the fairness measure of expectational fairness was introduced. The authors studied four different incentive models: PoW, Single-Lottery (SL) PoS, Multiple-Lottery (ML) PoS and Compound (C) PoS.

This thesis extends the research by examining the fairness of another incentive model, the Casper CBC PoS protocol. Unlike other protocols, the rewards are distributed according to the fractional stake of the participants, regardless of who proposes the block. In addition to calculating the expectational fairness measure, an agent-based simulation is performed for this purpose, which analyzes the effects of the reward distribution mechanism in isolation.

Moreover, this work contributes to empirically support the academic works on theoretical analysis of fairness and decentralization in PoS blockchains (Huang et al., 2021; Nguyen et al., 2019; Roşu & Saleh, 2021) with real data obtained from a blockchain.

Measures such as the Gini coefficient (Campajola et al., 2022; Li et al., 2020b; Jia et al., 2022; Sai et al., 2021; Lin et al., 2021; Wang & Ge, 2022), Nakamoto coefficient (Campajola et al., 2022; Sai et al., 2021; Lin et al., 2021), Expectational fairness (Huang et al., 2021) and information entropy (Jia et al., 2022; Wu et al., 2019; Lin et al., 2021) have been widely used to determine the fairness of a blockchain in terms of wealth and reward distribution and in terms of its decentralization. Although there is a growing body of empirical research on PoW blockchains, the research on PoS blockchains remains very limited. This paper aims to contribute to an enhanced understanding of fairness in PoS consensus protocols through the empirical evaluation of the CBC Casper PoS blockchain.

4 Simulation

In the following part, two agent-based models will be designed, one to represent the basic principles of the staking reward mechanism of a PoS in a pure form (hereinafter referred to as basic PoS reward distribution) and the other according to the specifics of the Casper blockchain. Subsequently, both models will be tested with various parameters in order to gain insight into how the distribution of wealth, measured with the Gini coefficient, is influenced by the characteristics of the respective reward mechanisms. An identical linear block reward function is assumed in both cases.

4.1 Properties of Reward Distribution

4.1.1 Casper Reward Distribution

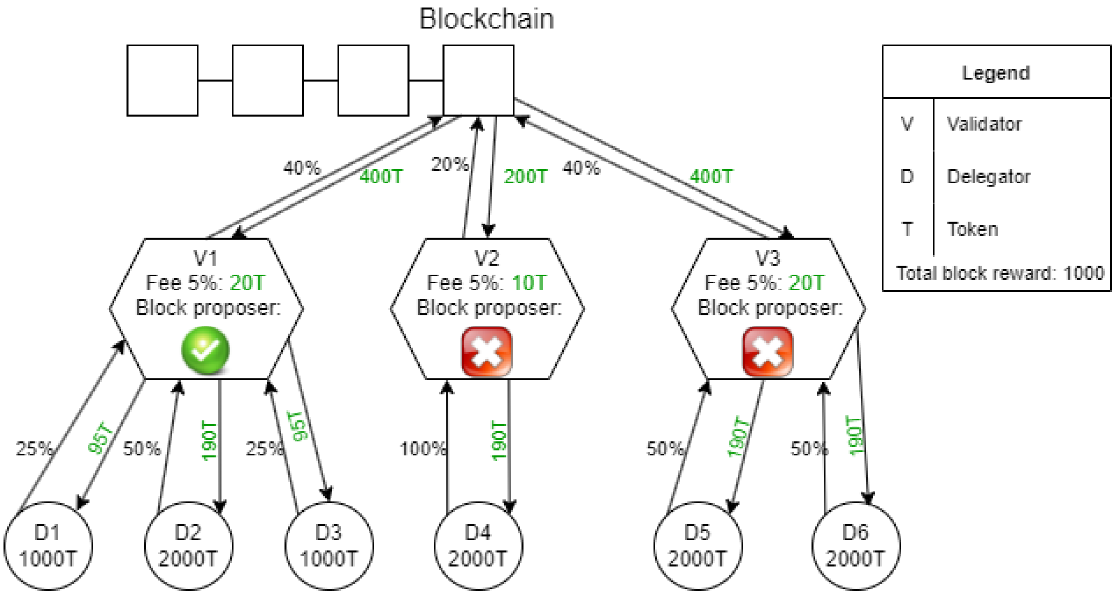

Figure 1 Reward distribution in the Casper blockchain

In the Casper blockchain, delegation describes the allocation of CSPR tokens from a token holder to a selected validator in the network who participates in the consensus protocol of the Casper network. Following the principle of a PoS network, a validator stakes CSPR tokens in order to be authorized to propose new blocks and approve blocks proposed by other validators - this requires the validator to ensure a constant uptime of the server or node by ensuring a reliable internet connection, providing sufficiently strong hardware and performing necessary updates, fixes and other maintenance in time to participate in the consensus mechanism without interruptions.

As a reward for participating in the consensus protocol, staking rewards in the form of CSPR tokens are paid out at regular intervals (at the end of each era, which lasts approximately 2 hours). These tokens are newly minted for this purpose at the time of the payout and thus increase the total amount of CSPR in circulation. The Casper network targets an inflation of 8% annually through newly issued tokens, which will be distributed among the respective eras. Accordingly, an APY for staking balances of 8% would result if the entire total supply would always be staked. However, significantly less than the total supply is staked, resulting in a higher APY than 8%, since the same amount of rewards per era is distributed to a lower amount of stakes.

In the Casper network, the rewards per era are distributed to all participating validators based on the validator weighting (calculated as the staking balance of all delegators and the validator's own stakes). The validator can set a percentage fee, which is deducted before the proportional distribution of the rewards to the delegators. The share of a delegator in the total rewards of an era thus corresponds to the weighting of its staking balance against all active stakes in the network minus the fee set by the validator. The income of a validator is, besides the commissions he earns through the validator fee, also the weighted share of the total staking rewards through his own stake. The block-proposing validator and its delegators also receive all transaction fees incurred in the era. Only one block-proposing validator is chosen per era, this choice corresponds to a weighted random selection from all active validators, where the weight of a validator is composed of all staking balances that have been staked to the corresponding validator by delegators and the validator himself. However, since transaction fees are extremely low in the CSPR network, these additional rewards for block-proposing validators are negligible. Figure 1 illustrates the reward mechanism and distribution of rewards to delegators and validators in the Casper network with exemplary values.

4.1.2 Basic PoS Reward Distribution

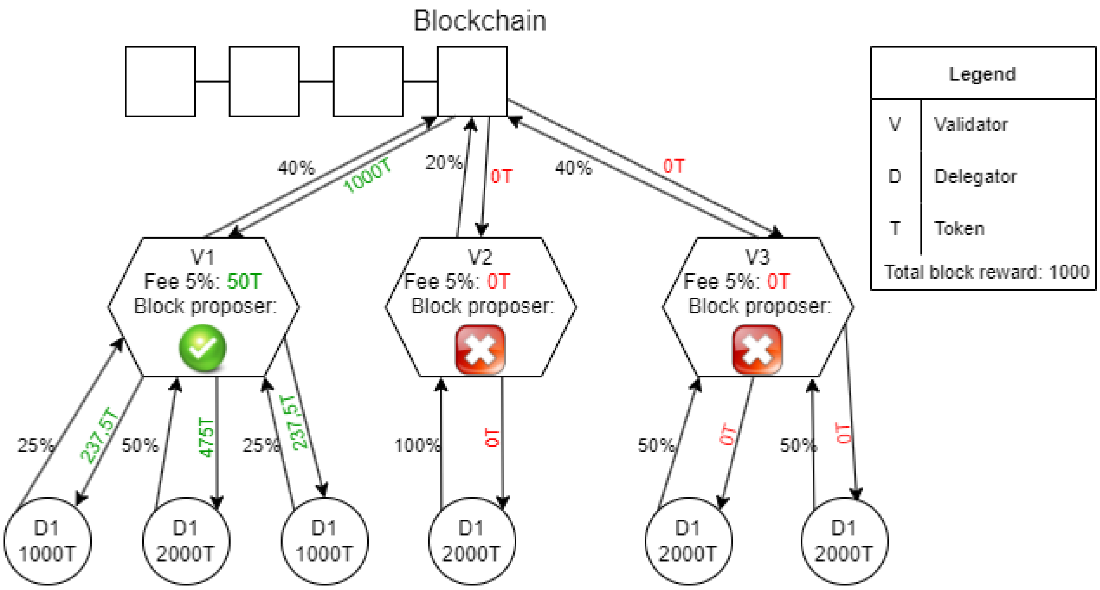

Figure 2 Reward mechanism in basic PoS blockchain

In the following, the reward mechanism of a PoS network of a more original form is described, hereafter it is referred to as basic PoS reward distribution. As with the Casper network, we specify that token holders can act as delegators by staking their tokens with a validator and that the rewards assigned to a specific validator are paid out to its delegators proportionally according to the weighting in the validator pool (based on the staking balances) minus a percentage validator fee. We also make identical assumptions regarding the origin of rewards (linear reward function based on total supply inflation), the negligibility of transaction fees (not to be confused with validator fees), and the payoff timing at the end of each era.

As in the Casper network, a block-proposing validator is determined for each era by a weighted random selection whose weight is given by the sum of all corresponding delegator-stakes and the validator self-stake. Unlike the Casper network, in the basic PoS network a block-proposing validator is not only allocated the negligible transaction fees in addition to the proportional era rewards, but the entire staking reward for the era, while all other non-block-proposing validators and their delegators do not receive any reward for the era despite an active stake. This leads to the fact that within an era delegators of a block-proposing validator receive significantly more than their weight would account for in the total staking balance in the network. This distribution within an era is illustrated in figure 2 with exemplary values.

4.2 Reward Distribution Models

In the following, based on the elicited properties of the reward mechanisms, parameters are described, model assumptions are made and procedures used are explained with which the agent-based models used can sufficiently simulate the properties of both a Casper reward mechanism and a basic PoS reward mechanism.

The defined model parameters for both models are the initial total supply, the initial total staking balance, the inflation rate of the total supply per time step, the number of participants (N), the initial distribution of staking balances and the number of time steps for which the simulation should be performed. Furthermore, it can be specified from how many independent runs an average value should be formed, in order to generate a more reliable value without overrepresentation of outliers. In each time step, the current Gini coefficient of the wealth distribution of the staking balances is calculated. At the end of the simulation, a plot is created which visualizes the development of the Gini coefficient over the selected period for the selected parameters. The individual steps that are performed are described below. Note that only the reward distribution (step 4A for Casper, step 4B for basic PoS) differs between the two models.

1. Initial distribution of staking balances: At the beginning of the simulation, an initial distribution of staking balances is performed. Depending on the parameter set, the defined initial total staking balance is either distributed evenly to all N participants (fair distribution), which leads to an initial Gini of 0, or distributed according to a random uniform distribution, which leads to a Gini of approximately 0.33.

2. Designate validators and corresponding delegators: 100 participants are randomly selected to act as validators for the entire simulation period. Furthermore, all participants (including the defined validators) are randomly assigned a validator to which their staking balance is defined as delegated. Accordingly, the staking balances of the defined validators can be assigned to other validators as well as to themselves (self-stake).

3. Reward calculation: The total delegator rewards and validator fees for the current time step are calculated. The total rewards are calculated by multiplying the current total supply by the inflation rate per time step. Then, the rewards of the delegators and analogously the validator fees are calculated as follows:

$$Total\ delegator\ rewards = (1 - validator\ fee\ rate) * current\ total\ rewards$$

$$Total\ validator\ fee = validator\ fee\ rate * current\ total\ rewards$$

4A. Reward distribution for Casper: The calculated total rewards of the current time step are divided proportionally according to the weighting of the validators and delegators. Since a uniform validator fee rate has been defined for all validators, the individual delegator rewards can be calculated according to the proportional share of the total staking balance, while the validator fees can be calculated based on the validator weighting (sum of all stakes delegated to a validator) in the current time step.

4B. Reward distribution for basic PoS: In a first step, a probability-weighted random draw is performed to determine which validator acts as block-proposing in the time step. The probability of being selected as block-proposing validator is calculated by the current validator weighting. Subsequently, all rewards of the time step are allocated exclusively to the block-proposing validator and his assigned delegators. Consequently, the delegators receive their share of the total reward of the time step proportional to their share of the validator staking pool, while the validator receives the total calculated validator fees of the time step. All other 99 non-block-proposing validators and the respective delegators do not receive any rewards for the time step.

5. Updating staking balances and total supply: The staking balance of each participant is increased by the individually calculated rewards. The total supply is increased by the sum of the rewards paid out.

6. Gini coefficient calculation: The Gini coefficient of the current time step is calculated on the basis of the staking balances.

7. Repetition for time steps and averaging for runs: Steps 3-6 are repeated for the number of defined time steps. Then all values except the defined initial distribution of the staking balances are reset and the process starts again from step two. After all defined runs have been executed, the average value of the Gini over all executed runs is calculated for each time step.

8. *Staking balance Gini coefficient plot:* Output of a plot showing the development of the Gini coefficient of the staking balance as a function of time steps.

4.3 Simulation Results

For the simulation, the following model parameters were chosen for both the Casper reward mechanism and the basic PoS reward mechanism in order to obtain comparable values:

Initial total supply: 100’000

Initial total staking balance: 60’000

Inflation rate per time step: 0.1%, 0.5%, 1%

Validator fee rate: 10%, 20%, 30%

Time steps: 500

Averaged over: 20 runs

4.3.1 Casper Reward Distribution

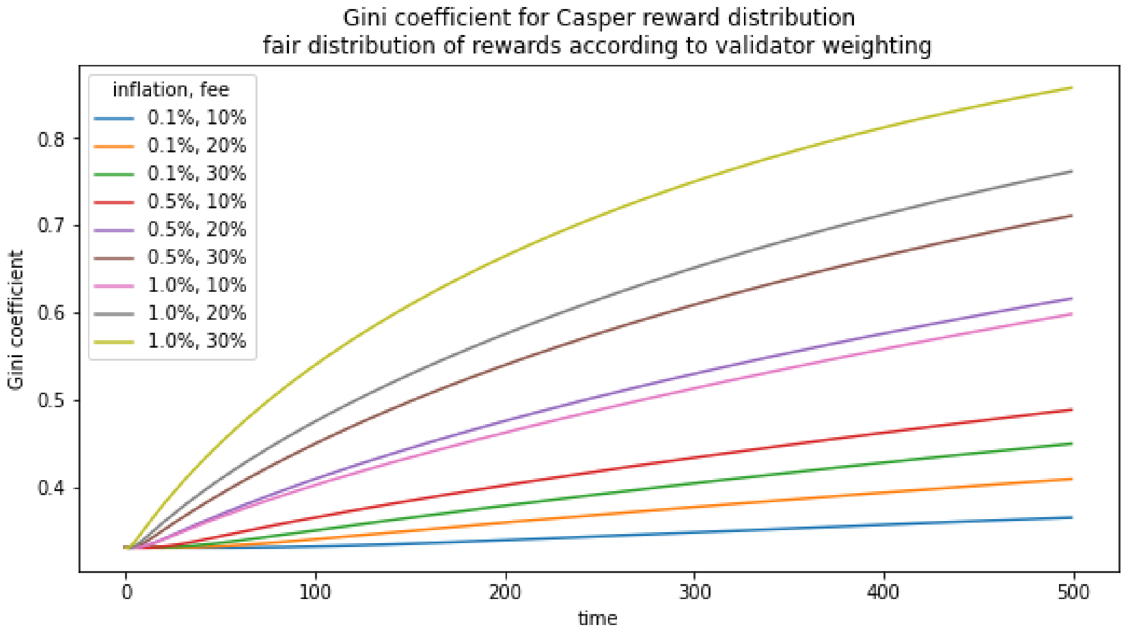

Figure 3 Gini coefficient for Casper reward distribution simulation for 500 time steps

The simulation of the Casper reward distribution shows that for all selected inflation rates and validator fee rates, the Gini coefficient, starting at 0.33, increases continuously over the time steps in a flattening curve (figure 3).

Higher inflation rates are associated with a steeper curve and thus with higher Gini values in similar time steps. The level of validator fees amplifies this effect of the inflation rate on the Gini coefficient. A closer look at the first time steps for the respective values also shows that the Gini value does not increase sharply immediately, but remains stable or even decreases slightly at the beginning (figure 4). This phase with little or no growth in the Gini coefficient lasts longer, the lower the inflation rate or the validator fee.

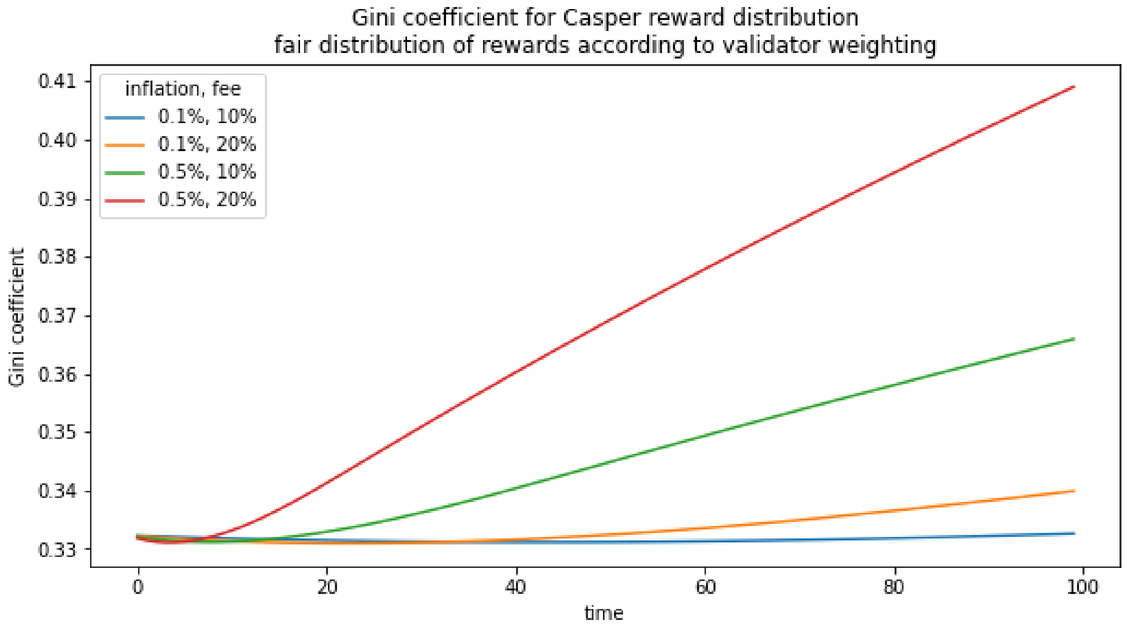

Figure 4 Gini coefficient for Casper reward distribution simulation for 100 time steps

A separate simulation over 25’000 time steps also shows that the Gini coefficient, over a longer time period, converges continuously to the value 1 of full inequality even for low values of inflation and validator fee (figure 5).

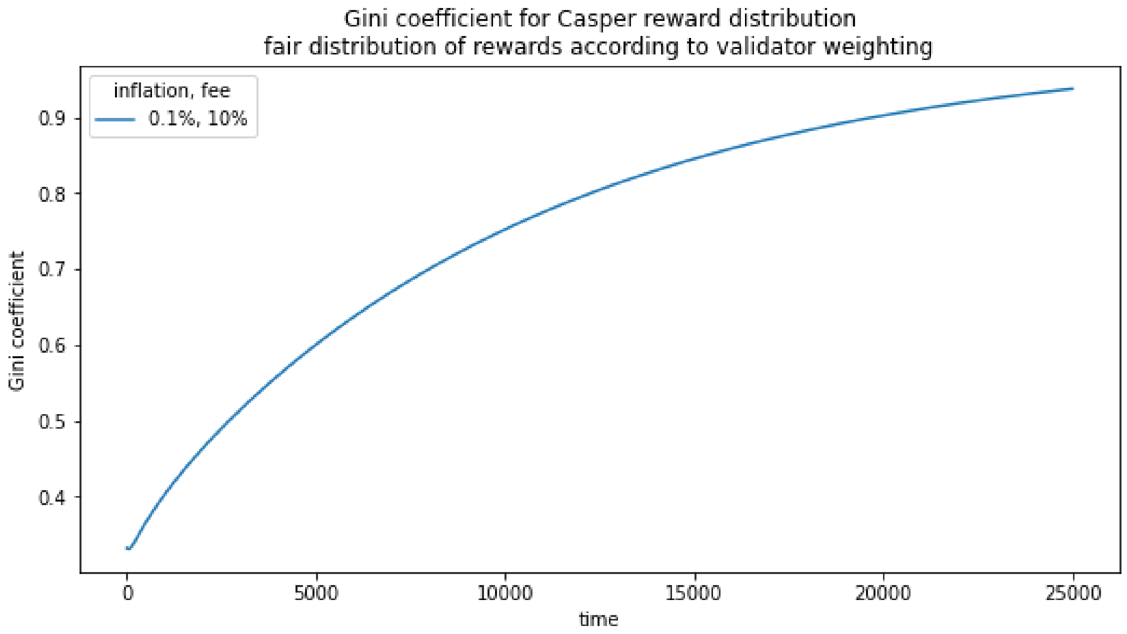

Figure 5 Gini coefficient for Casper reward distribution simulation for 25'000 time steps

4.3.2 Basic PoS Reward Distribution

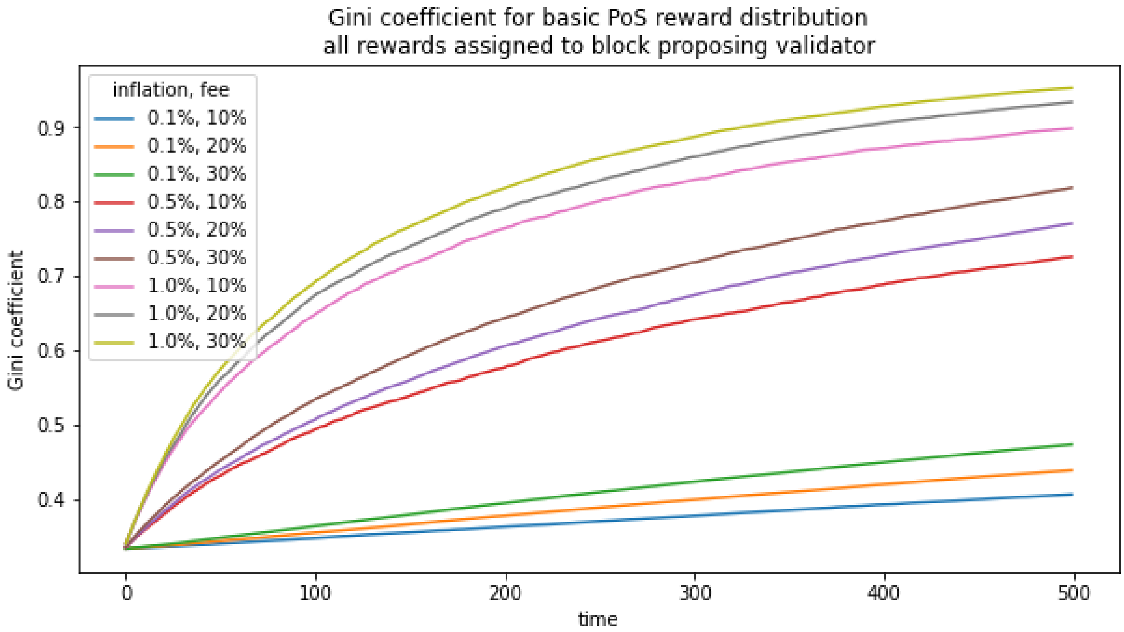

Figure 6 Gini coefficient for basic PoS reward distribution simulation for 500 time steps

For basic PoS distribution, as for Casper, a general trend can be observed that higher inflation rates and validator fees correlate with faster rising Gini values. In contrast to the Casper reward distribution, however, the increase in the Gini value already begins with the first time steps, as figure 6 illustrates.

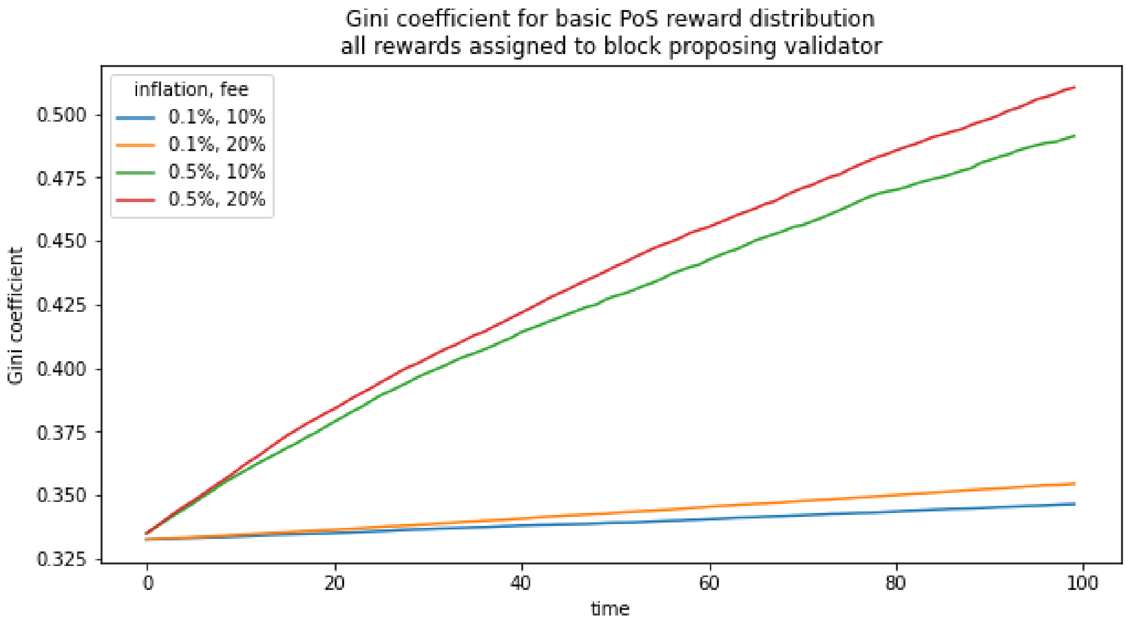

Figure 7 Gini coefficient for basic PoS reward distribution simulation for 100 time steps

Overall, increasingly higher Gini values are observed in the basic PoS reward distribution for all inflation and validator fee rates in all time steps compared with the Casper reward distribution. Accordingly, in such a system, higher inequality is reached after a shorter time, and convergence to the value 1 of the maximum inequality occurs in a shorter time (figure 7).

5 Data Analysis

5.1 Data Acquisition

For data acquisition, a locally installed Casper client was used, which could query various information about individual blocks from connected Casper validators. Since the Casper blockchain only stores information about rewards in the last block at the end of each era in so-called “switch blocks”, the first step was to identify which blocks are switch blocks. A simple brute force approach proved to be ineffective, as the nodes showed increasingly higher response times with a high number of requests for every single block height. By using a public, experimental API (Casper Metrics, n.d.), which collects data from the Casper blockchain, it was finally possible to identify all switch block heights for the observation period. Using the Casper client, data on detailed staking balances and seigniorage allocations (paid out staking rewards) for all delegators and validators from era 0 (genesis) to era 5446 (most recent era at the time of the final data acquisition on 29.06.2022) could be retrieved by specifying the corresponding switch block height for each request. Furthermore, data on validators, their weights and total supply were retrieved.

5.2 Descriptive and Exploratory Data Analysis

The observation period covers 15 months, data since the launch of the mainnet on March 31, 2021 (era 0) until June 29, 2022 (era 5446) was considered. Within this period, a total number of 900’380 blocks were finalized.

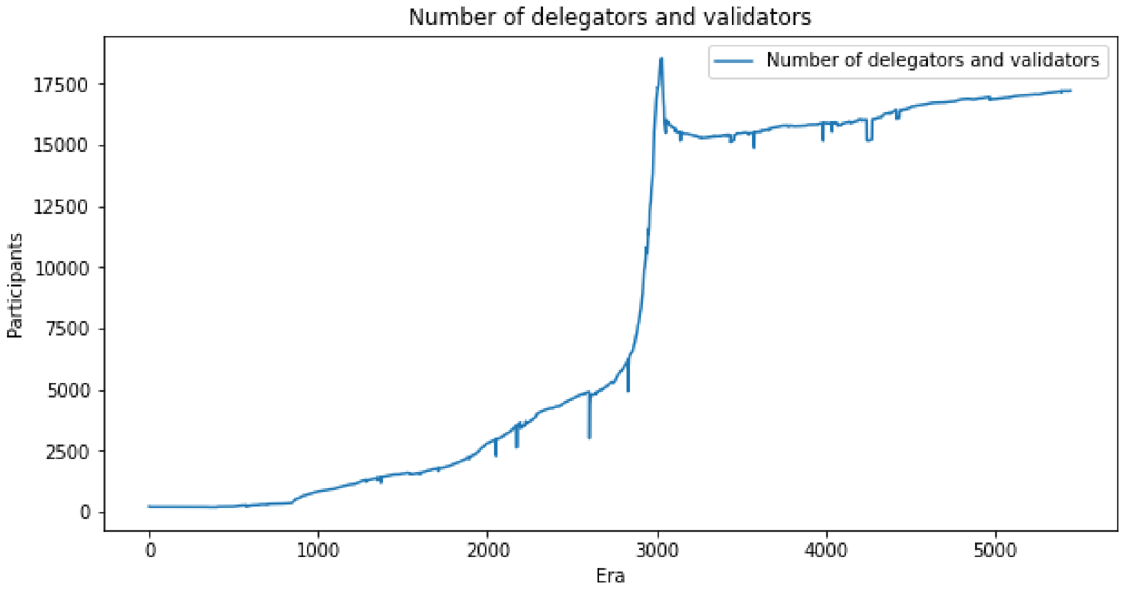

Figure 8 Development of the number of delegators and validators over all eras in the Casper network

As shown in figure 8, the number of participants (active delegators and validators) in the Casper network increases steadily over time. In the early days, during the first 1000 eras, the number of participants remains relatively stable and then starts to gradually increase. A very rapid increase can be observed starting at era 2900 (around November 29, 2021). Within a few days the number of participants almost tripled and for a short time reached a peak of 18'548 unique participants at era 3030. The sharp increase in the number of participants, which began at the end of November is followed by a sudden drop at the mid of December and coincides with the price development of cryptocurrencies, including the CSPR token, on the market (CoinMarketCap, n.d.). Presumably, the sudden drop in cryptocurrency prices also dampened the interest of investors, causing some participants to terminate their participation in the Casper network. In the first 500 eras after the crash, the number of participants stabilized and has been growing more slowly but continuously ever since.

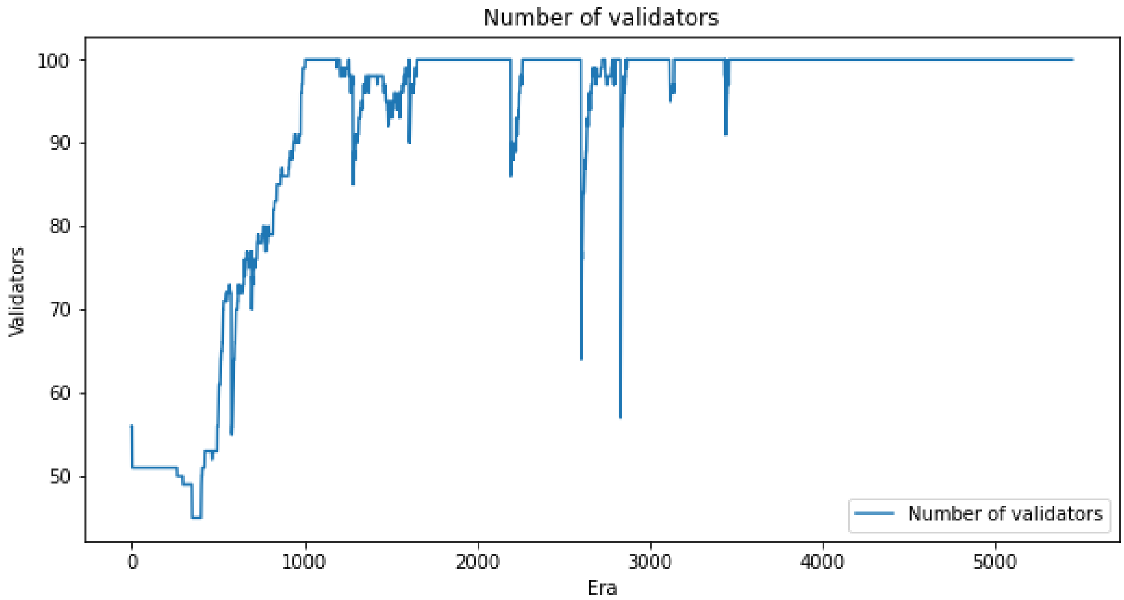

Figure 9 Development of the number of validators over all eras in the Casper network

Initially, the number of validators fluctuates around 50 during the first 500 eras. The number continues to increase up to era 1000 and then stagnates at 100 validators. This is due to the fact that the number of validators is technically limited to a maximum of 100. Usually, between 150 and 180 bidders compete for the 100 validator places, which are awarded to those with the highest number of stakes. Since the number of bids in almost every era is greater than the available slots, the number of validators usually remains constant over time at the maximum of 100. It can be observed in figure 9 that a failure of up to 10% of the validators occurred from about era 1300 onwards. This trend lasts for almost 300 eras, during which the number of validators continues to fluctuate. One cause for this incident could be the update 1.3.1, which was implemented in era 1346. Some validators may had problems with the transition or encountered network problems. A second update (1.3.2) in era 1605 seems to have resolved this issue. Similarly, other major deviations also occur at the same time as a new updates (update 1.3.4 in era 2193 and 1.4.1 in era 2600) were rolled out. In addition, two network outages occurred, one in era 2595 during the changeover to a new update, and the second in era 2819, both lasting for several hours. An explanation for smaller, mostly short-term downward deviations could be due to network problems of individual validators or certain regions.

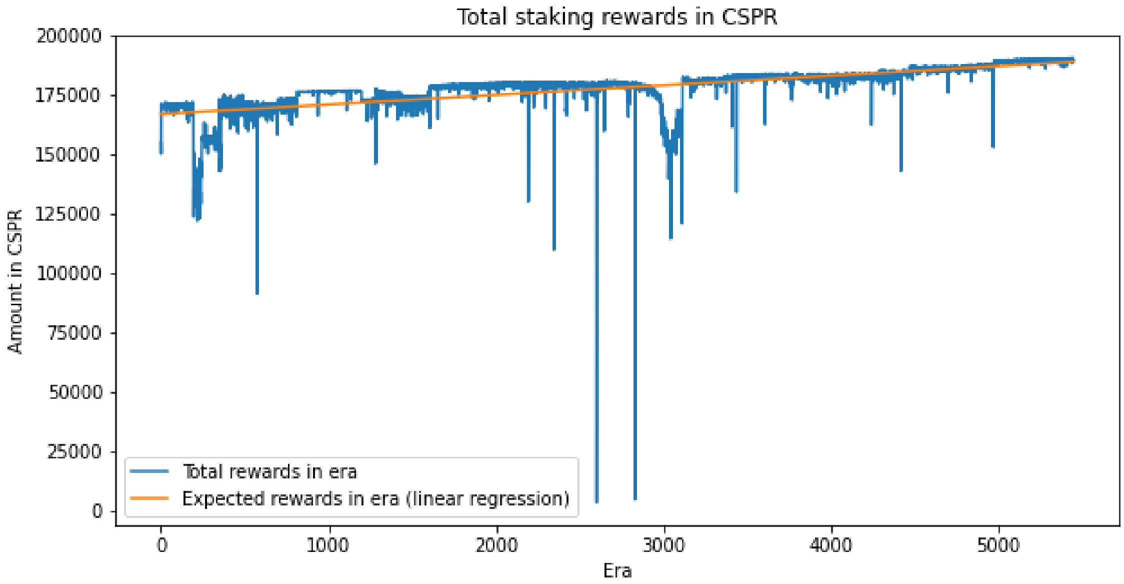

Figure 10 Total rewards for validators and delegators per era

Figure 10 shows that the total rewards increase steadily over the eras. The expected rewards assumed by a linear regression are plotted in the graph with the orange line. The development of the total rewards is in line with statements made by the Casper team, which claims that the distribution of rewards is based on the total supply and that 8% of the total supply is minted each year and distributed as a reward to the validators and delegators after each era. In the event of a failure of many validators, the total reward that can be paid out in this era is multiplied by the percentage of the remaining active validators and only this reduced amount is subsequently distributed. Nevertheless, each validator receives a reward as a percentage of his staked balances, but each of them receives less. This leads to the fact that in some eras there is almost no reward or for some time, the total reward is significantly lower than the expected reward for the validators. The steep downward spikes in era 2595 and 2819 can be attributed to a system crash in the Casper network.

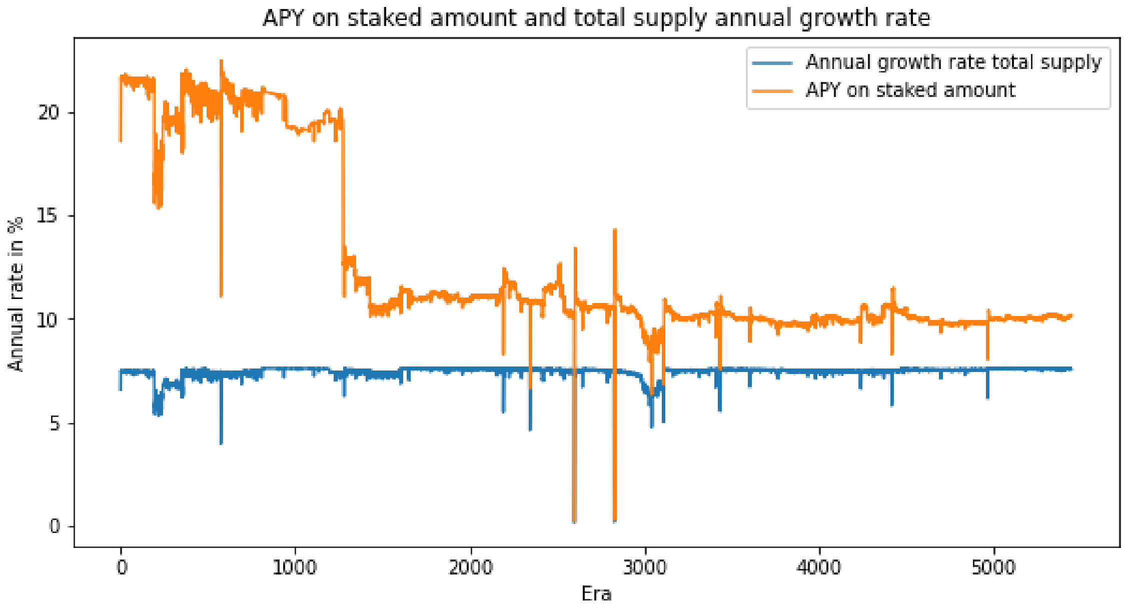

Figure 11 APY and annual growth rate per era

The annual growth rate of the total supply, shown in blue in figure 11, has been relatively consistent throughout the year, ranging from 7 to 8%. For the whole observation period, the annual growth rate of the total supply was on average 7.43%, as shown in table 2. The APY, on the other hand, fluctuated between 15 and 22 percent for most of the first 1300 eras. This is due to the fact that at the beginning only approx. 38% of the total supply was staked, which results in a higher return for the individual participants. As the number of participants increases from era 1000 onward and the percentage of total staking balances to total supply rises sharply by almost 50% in era 1300 (see figure 12), this causes the APY on staked amount to correct itself downward to about 10%. Since era 1300, the APY has remained relatively stable, with large downward swings occurring only in rare cases such as system crashes (era 2595 and era 2819). Over the entire observation period, the average APY was 12.72%.

Table 2 Annual growth rate and APY metrics

	Annual Growth Rate	APY
Mean	0.074360	0.127237
Std	0.003420	0.042821
Min	0.001375	0.001897
25%	0.074333	0.100280
50%	0.075142	0.105836
75%	0.075711	0.124427
Max	0.076348	0.225031

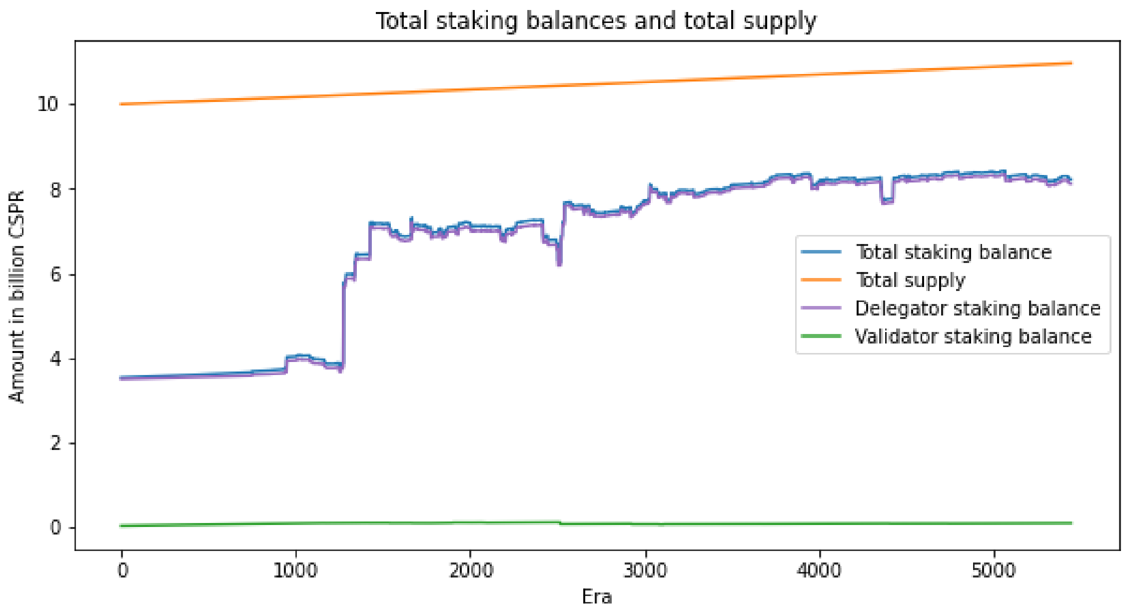

Figure 12 Total staking balance and total supply per era

Figure 12 shows the development of all staking balances together with the total supply. Initially, only about 38% of the total supply of the Casper network was staked. Over the first 1000 eras this has increased to about 40%. In era 1300, however, there is an extremely strong increase in the staked balances to around 60% and only a short time later in era 1500 it rises up to about 70%. A possible reason for this sudden increase is that the Casper Association started to deposit their own assets in the network in order to collect rewards for the future. As can be seen in figure 13, the Casper Association owns 14.3% of the initial token allocation. Insiders, which include the team, the advisors and Casper Labs Holding AG, own a further 40% of all initial tokens. Unfortunately, Casper did not provide any information regarding this strong increase in staked tokens. However, the curve of the total staking balances is moving towards a constant growth rate after this increase. Since the slopes of the total supply and total staking balances curves are quite similar, it suggests that the staking rewards participants receive for their staked tokens remain on their staking account and are not withdrawn. Over the entire time period, the curve of the delegators' staking balances is slightly below the total stakes, but the trend is almost identical. Turning to the staking balance curve of the validators, one can see that the curve is almost constant and at a very low level over the entire duration of eras. These facts indicate that the validators are staking their assets and continuous income not with the public key of the node, but with another account as delegator.

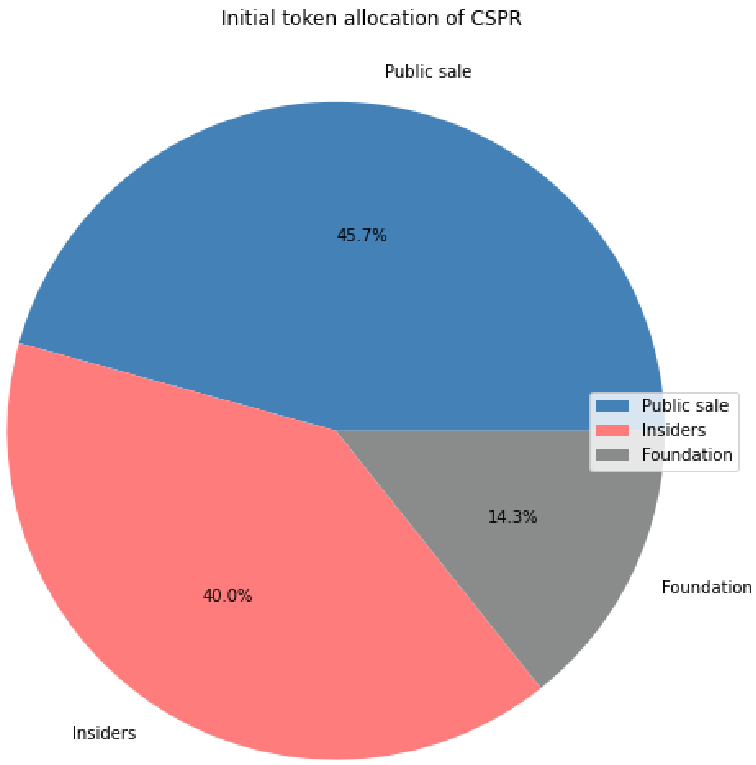

Figure 13 Initial token distribution of Casper

The initial distribution of tokens can have a significant impact on the decentralization of a blockchain. Especially in PoS, where the amount of the native cryptocurrency of a participant also expresses his power, it is interesting to know how the initial distribution took place. As shown in figure 13 above, the share of tokens distributed in the public sale is the largest (see CasperStats Docs, n.d.; AscendEx, 2020). All tokens that were available for public participation during pre-launch sales (or "lock-drop" allocations) are included in this category. In the case of Casper, nearly 16% of the tokens were distributed during Coinlist sales and another 29.7% through validator sales. Nevertheless, 40% of all initial tokens were granted to insiders, including the team, developers and the company itself, such as the CasperLabs Holding AG. A final portion of the distribution goes to foundations and community governed grant pools, some of which operate independently of the founding company, including the Casper Association. However, when the non-profit Casper Association and the whole group of insiders are taken together, their share exceeds 50%, which means that the majority of the of power is in the hands of a few groups.

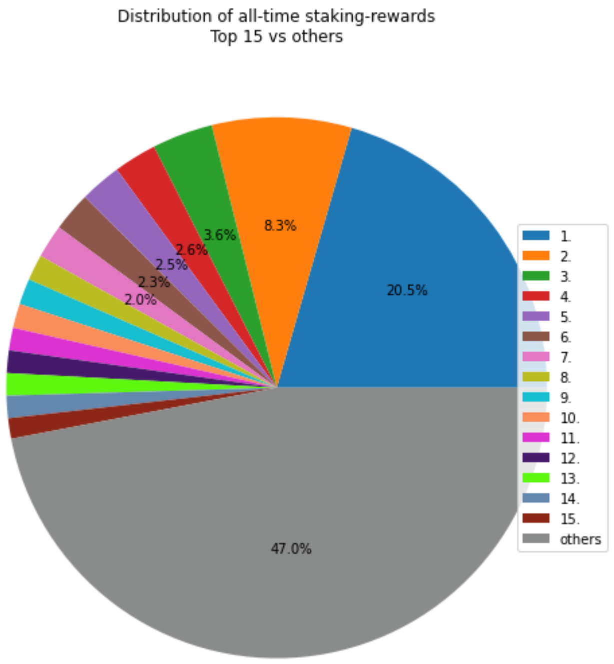

Figure 14 Share of staking rewards summed over all eras of validators and delegators

Figure 14 shows the distribution of staking rewards per public key summed up over the entire observation period. The pie chart reveals that the top two participants received more than 25% of all staking rewards. Furthermore, the top 15 received more than half of all staking rewards. Although Casper insiders and the foundation as a whole own more than 50% of all coins and have at least partially staked them, there is no delegator or validator that is of such a large size. Therefore, it can be assumed that members of the Casper team as well as other large investors have distributed their tokens over several wallet addresses.

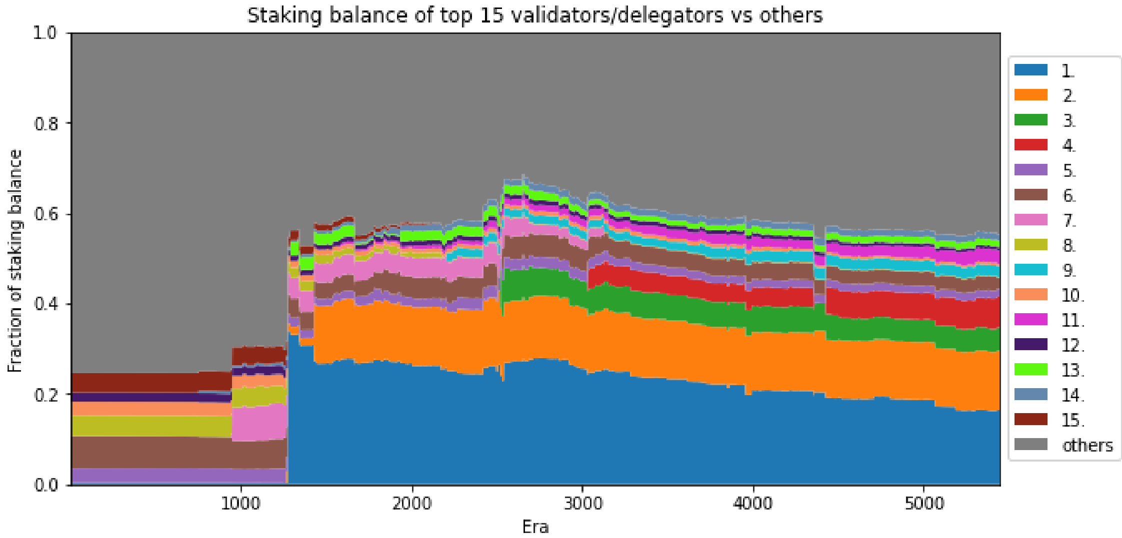

Figure 15 Share of staking balances per era among validators and delegators

In figure 15, we can compare the changes in fraction of the total staking balances of the top 15 validators and delegators in staking rewards determined in figure 14. In particular, around era 1300, two major players enter the game and from then on dominate the first two ranks with a share of about 23% and 17%, respectively. In some eras, these two together account for 40% of the total balances. The percentage of staking rewards attributed to the two largest balance holders, which can be seen in figure 14, is smaller, since the two have only staked their coins from era 1300. Based solely on the initial token allocation and if the possibility of massive transfers is excluded, there are only a few participants who could possess such a large amount of CSPR, namely the Casper Association (14.3%), CasperLabs Holding AG (10%), the team (8%), the developers (16%) and the advisors (6%).

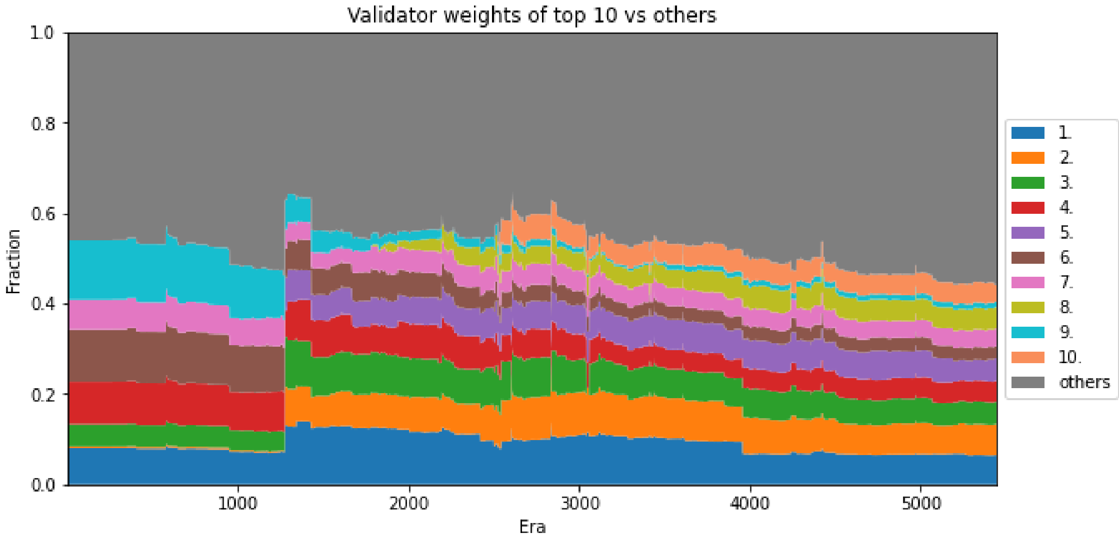

Figure 16 Validator share of staking balances per era

In figure 16 above, we can compare the average weights of the 10 largest against all other validators. The size of a validator is the sum of all balances that have been staked with him, both from individual delegators and from staking pools. A sudden increase in the staking balances of all top 10 validators occurs at era 1300, increasing their weight compared to the other validators. Compared to figure 15, none of the validators in figure 16 reach a total weight of 20 percent or more. Given this fact, it can be assumed that the top 2 delegators have not allocated all of their assets to a single validator, but rather have spread them over several validators. The top 2 delegators that joined from era 1300 onwards have distributed their tokens mainly among validators 1, 2, 3 and 5, as those have increased significantly in weight compared to the other 6 validators, which was confirmed by a closer look at the underlying data. In the beginning, the top 10 validators collectively covered almost 55% of the total staked assets; this share increases to over 60% as of era 1300 and then gradually diminishes thereafter, settling at around 50% at the end of the observation period, in era 5446.

5.3 Results

5.3.1 Gini Coefficient

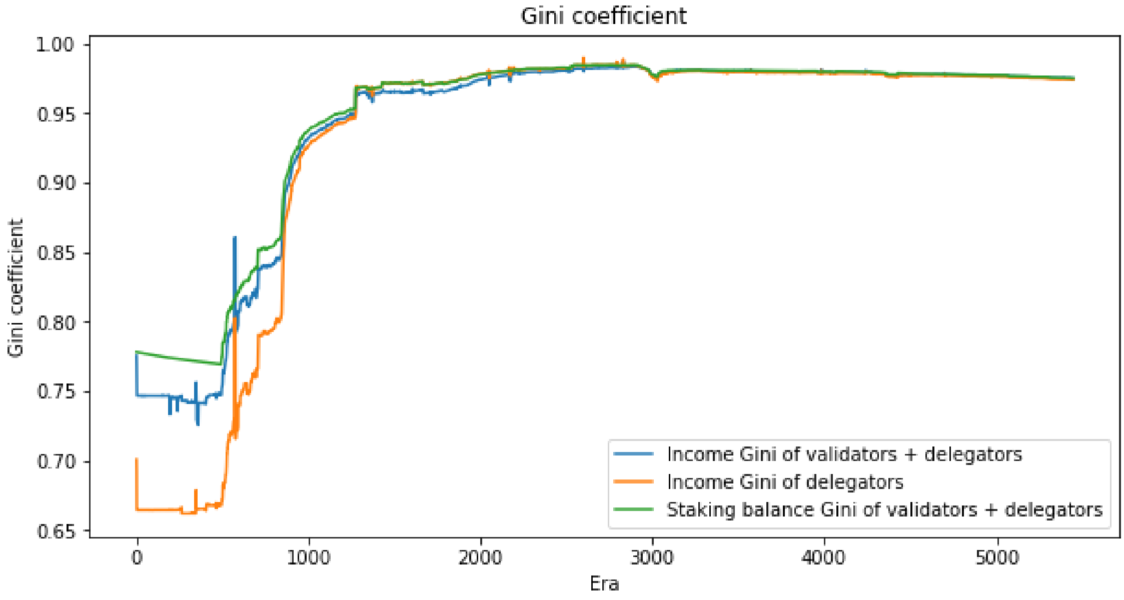

Figure 17 Gini coefficient of staking rewards for all delegators, validators and validators + delegators per era

Figure 17 illustrates the development of the Gini coefficient for staking rewards (income) and staking balance (wealth). It can be seen that the coefficient already starts at high values between 0.66 and 0.78 and, following an s-shaped curve, stabilizes at a very high value of approx. 0.97 after about 2000 eras. Since era 3000, a slight decrease of the inequality can be observed. Furthermore, it can be observed that the three curves (Gini of validator and delegator income, of delegator income only and of validator and delegator staking balance) show significantly different values only at the beginning and are practically identical towards the end. On the one hand, this indicates that the validators do not restake their income via their own address, which could already be deduced from the constantly very low staking balance of the validators (see figure 12), and that, overall, a large part of the income generated is directly staked again, since changes in the Gini coefficient are almost identical in all curves, especially from Era 1000 onward. For the two income Gini curves (delegators only and delegators plus validators), a sharp temporary increase in the Gini is observed in the short term at Era 694, which can be explained by the first major update 1.2.0. Due to the relatively small number of participants at the time, the short-term omission of some validators who had difficulties with a timely update (see figure 9, 10 and 11), led to relatively large fluctuations in the rewards paid out, since a significant fraction of participants received little or no rewards in the short term, which makes the income appear particularly unfair for these periods. Furthermore, the network outage at block 2819 can also be seen here, which also caused some larger participants with presumably high individual stakes (see figures 15 and 16) to suffer a longer outage, so that the inequality decreased slightly in the short term.

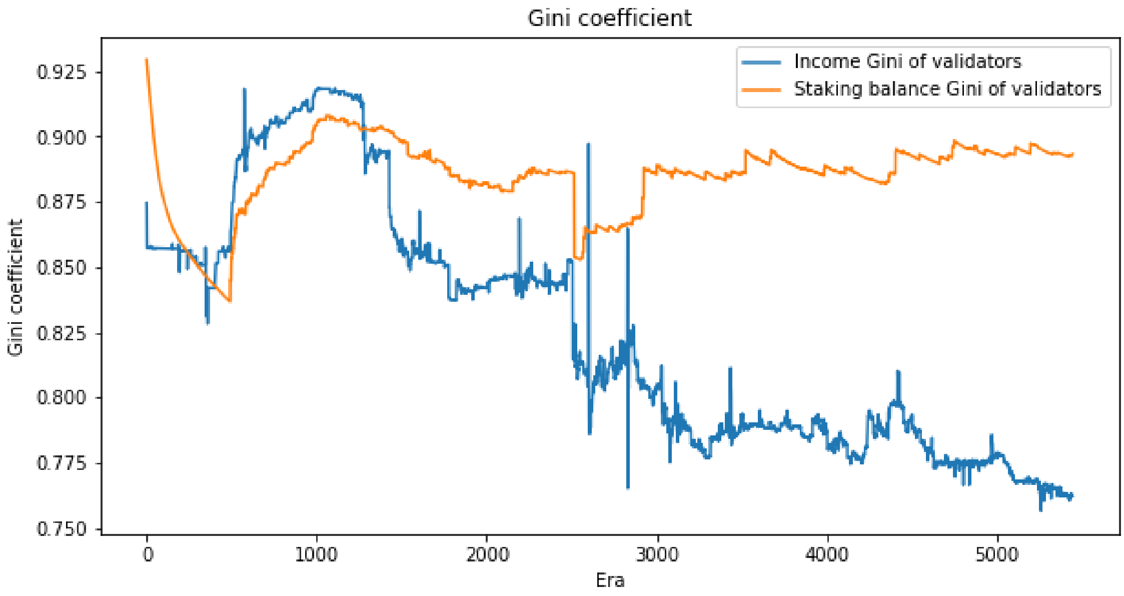

Figure 18 Gini coefficient of staking balances for all delegators, validators and validators + delegators per era

When plotting the Gini of validator balances and rewards (figure 18), it is noticeable that the staking balance Gini is relatively stable. However, since the stakes of the validators are constantly at a very low value anyway (see figure 12), the significance of the values is low. Of greater interest, however, is the Gini coefficient of the income of validators, which, in addition to staking rewards, consists in particular of the commission from validator fees. After initially high values between 0.83 and 0.92, the Gini value decreases and is only at 0.76 at the end of the period under review. This indicates signs of decentralization, since over time the stakes of the delegators were increasingly distributed from a few large validators to several smaller validators, whereby validator fee income is also distributed more fairly. This observation is consistent with the descriptive data in figure 16, where a decreasing dominance of the largest validators by weighting can be seen. Special mention should be made in this context of the upgrade 1.4.5, which temporarily sets a limit of 953 delegators per validator from era 4417 on. This meant that the largest validators by weight, which already had over 953 delegators, were no longer able to accept further delegators. Potential new delegators had to choose smaller validators accordingly, which is one explanation for the increasingly fair distribution. It therefore remains to be seen whether this trend will continue when the limit on delegators is lifted again.

5.3.2 Nakamoto Coefficient

The Casper Highway protocol requires that at least 2/3 of all stakes are needed to reach consensus. Conversely, a certain number of validators are needed to gather more than 33.33% of all stakes to sabotage the protocol. A decentralization measure that determines how many validators would have to collude to reach the required threshold at which consensus can be prevented is the Nakamoto index. The higher the Nakamoto coefficient, the more validators are required, and the more secure and decentralized the network is. Figure 19 shows the Nakamoto coefficient as the number of validators per era. Unlike other blockchains such as Bitcoin or Algorand, where an unlimited number of miners or validators can participate, Casper has a unique feature in that the maximum number of validators is currently limited to 100. The graph shows that the Nakamoto coefficient increases steadily over time, but remains at a very low level until era 5446. Such low coefficients for the Nakamoto as in Casper are not very uncommon. The study by Campajola et al. (2022) revealed that cryptocurrencies such as Dogecoin, Feathercoin, Monacoin and Bitcoin Cash all have a very low Nakamoto index and even fall to 1 occasionally. Particularly surprising are the findings regarding Ethereum, the second largest cryptocurrency, where the Nakamoto index has predominantly fluctuated between 2 and 3 over the entire time. The only exceptions were Litecoin and Bitcoin, which seem to be comparatively less centralized.

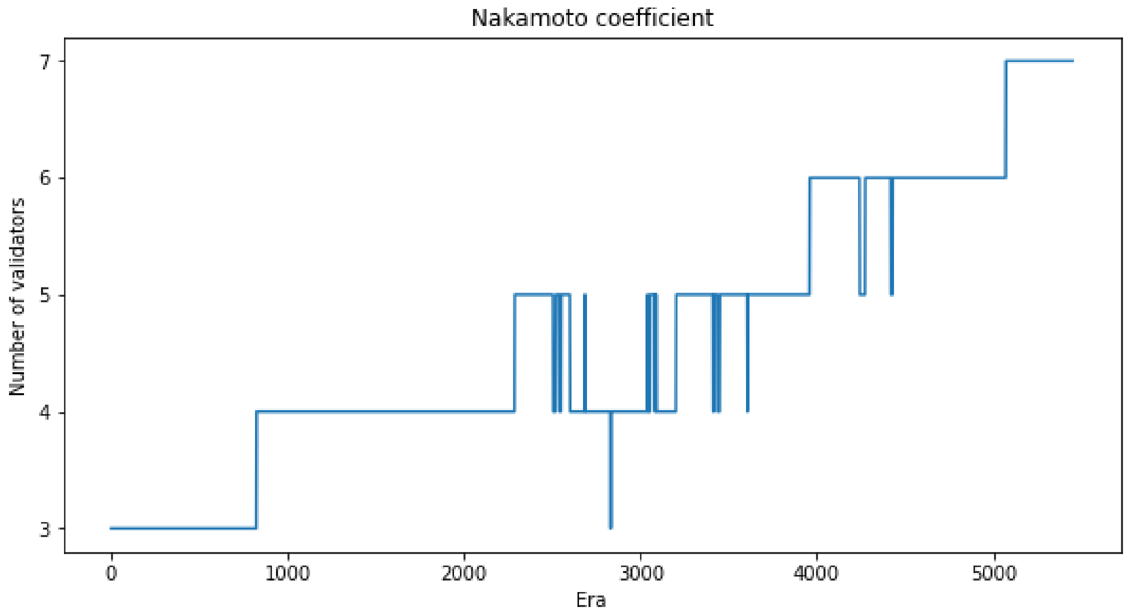

Figure 19 Nakamoto coefficient in number of validators needed in order to reach 33.33% threshold

5.3.3 Expectational Fairness

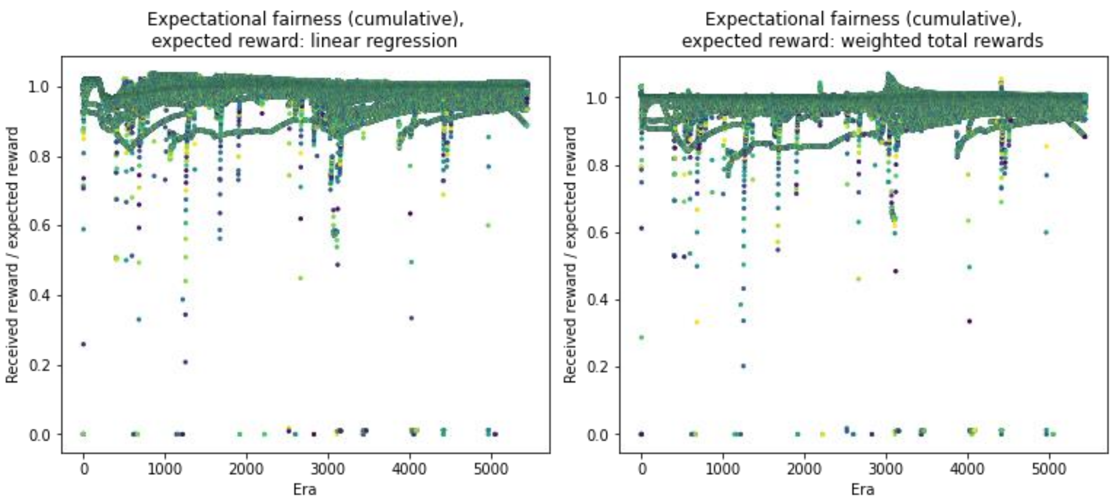

Figure 20 Expectational fairness calculated with cumulated rewards, with the calculation of expected rewards being based either on weighted linear regression prediction or on validator weight * total paid out era rewards

The expectational fairness here indicates how much a validator pool (validator considered as a whole with all its delegator stakes) has effectively received in rewards compared to the rewards that were expected. A value of 1 therefore means that a validator pool received exactly as many rewards as expected, and a value of 0.5 means that only half as many as expected were paid out.

Two different methods were used for the expected value of the rewards. On the one hand, a linear regression based on the historical data was applied. The expected value of the rewards for an individual validator pool is then equal to its weight multiplied by the expected value of the total rewards for the era. This procedure smooths out the outliers and results in an expected reward slightly below the computed APY. The assumption of such expectations seems plausible, since the Casper network is largely stable and only isolated failures have led to a short-term drop in rewards. By comparing the cumulative expected rewards with the cumulative rewards actually received, figure 20 (left) thus shows whether the validators have been assigned those rewards over time that were expected by the total supply growth rate communicated by Casper and the resulting APY. Each point on the graph symbolizes the cumulative expectational fairness value of a validator at the respective time. Overall, the vast majority of the data points are concentrated between 0.95 and 1, which indicates that largely the rewards expected by the validators were paid out. Furthermore, it can be seen that relevant deviations only occur downwards and are followed by a convergence to the value 1. Downward deviations are largely consistent with the dates of network failures and upgrades (see table 1 in 2.5.1), which indicates exceptions and not a systematic disadvantage of individual validators.

Furthermore, on figure 20 (right), a different method was used for the expected rewards: in this case, instead of using the expected value from the linear regression, the weighting was multiplied by the total rewards paid out for the period. Thus, the data points on the graph indicate whether the expected percentage share of total rewards for the period was paid out, which should correspond exactly to the weighting of the individual validator. This makes it possible to read on the graph whether in certain situations some validators received more of the total share than their weighting would account for, regardless of whether total rewards have suffered a sudden downturn during the period. Especially in cases where the absolute total reward dropped abruptly (e.g. the network update of era 2819), it can be seen that many validators received less than their weighting, while some validators received slightly more than their weighting. Here too, however, a rapid reconvergence to the value 1 can be observed.

For both methods (linear regression and expected fraction of total rewards), we further plotted the individual values in figure 21 instead of the cumulative values. The individual points thus describe how much a validator pool received compared to the expected reward within the period, independent of the previous periods. From this we can see how strong the deviations up and down within a period were and whether it happened that several times in a row higher rewards were achieved by a single validator pool, which would indicate selfish behavior. The graphs show that although values up to almost 1.75 of the expected share of the rewards were achieved (figure 21 right), the absolute rewards received in CSPR were not significantly higher than expected and were in most cases even significantly below the usual values, so that in absolute terms no single validator profited excessively from such incidents. These two observations lead us to the conclusion that there is no evidence of selfish behavior and that the reward system largely pays out rewards that are close to the expected values, especially in the long run.

A possible explanation for the accumulation of data points at 0 would be that short-term failures of a validator, e.g. due to hardware, network or software problems, led to the fact that a validator pool did not receive a reward for isolated eras. Individual values significantly below 1 and above 0 can also be triggered by short-term failures within an era for a few blocks or by large new stakes, which increase the weighting of the validator, but are only taken into account in the distribution of rewards after a few hours.

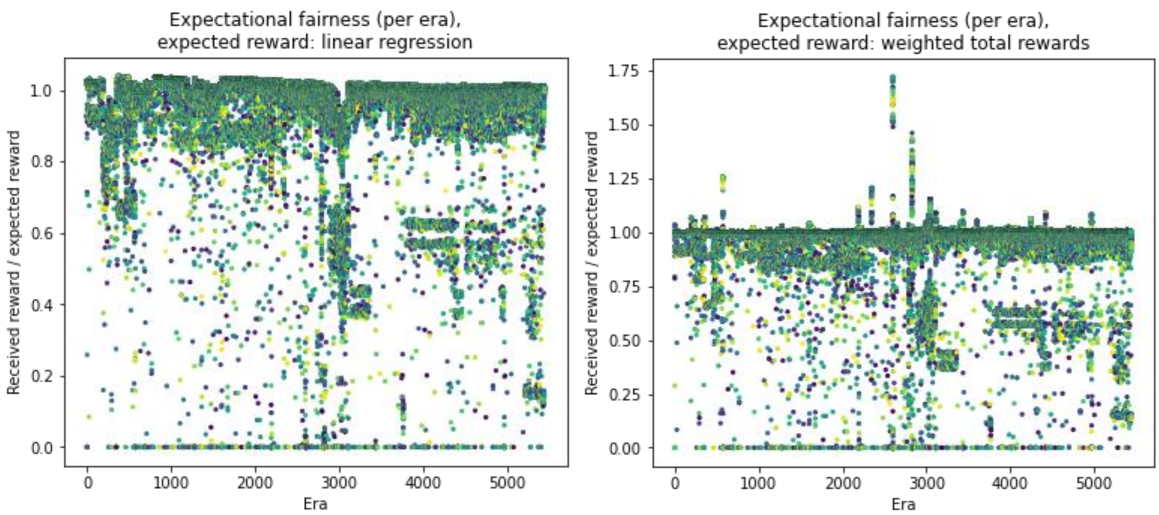

Figure 21 Expectational fairness calculated per validator per era, with the calculation of expected rewards being based either on weighted linear regression prediction or on validator weight * total paid out era rewards

5.3.4 Information Entropy

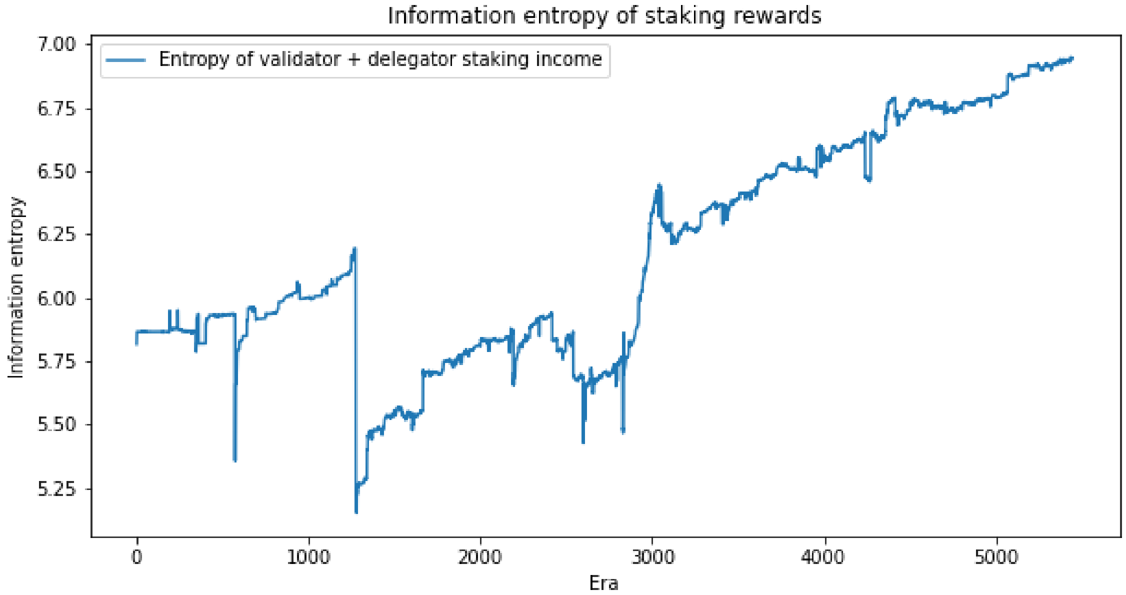

Figure 22 Information entropy of staking income for all delegators, validators and delegators + validators per era

Following the work of Wu et al. (2019) and Lin et al. (2021), the measure of information entropy was used to calculate and quantify the degree of decentralization in Casper. According to the results of the calculation of information entropy above, the value of information entropy increases over time, as can be seen in figure 22. A larger value of entropy indicates more disorder and randomness in the degree of the dataset and thus a higher degree of decentralization. At the time when the delegators recorded an increase of two relatively large new entrants in era 1300 (see figure 15), the value of the information entropy plummeted to its all-time low of 5.15. Following the crash, the information entropy begins to recover and is back at its initial level after approximately 1700 subsequent eras. Not least due to the rapid increase in new rather small delegators, which started around era 3000, there was a sharp rise in information entropy to a new high of 6.3. From this point on, there is a moderate growth in information entropy until the end of the observation period. Concluding from this analysis, we can observe a trend towards more decentralization in the Casper network.

6 Discussion

The simulations conducted showed that Casper's reward distribution also exhibits rich get richer effects, and based on the linear reward function in its current form also contributes to steadily increasing inequality in the system. However, in comparison with the other simulation, it is clearly visible that the increase of inequality under the same conditions is much slower than in a comparable normal PoS network, where the block-proposing validator is assigned all rewards. This can be explained by higher wealth compounding effects, which can arise when a wealthy participant receives all rewards disproportionately often compared to his initial stake due to his steadily increasing stake, which in turn further increases the future probabilities. This finding is significant in that it represents a decisive advantage of the PoS implementation of Casper over PoW blockchains: in the latter, it is not readily possible to pay out block rewards on a proportional basis instead of to just one miner, as this would require detailed and, above all, non-manipulable data on the proportional hash power of all miners. Another finding from the simulations is that higher validator fees and granted interest rates lead to a faster increase in inequality in both the Casper reward distribution system and normal PoS networks, to a greater extent in the latter. This finding emphasizes the importance of motivating a sufficiently large number of validators to participate and to allow them to do so protocol-wise, so that competition can arise and, accordingly, fees can be kept as low as possible through price competition. At the same time, it is advisable to refrain from excessive interest rates, as is still the case with many blockchains in order to attract new participants through high theoretical returns. Besides the fact that these returns are largely relativized by the associated high inflation and corresponding negative effects in the ecosystem, they also lead to a faster increase in inequality, which negatively affects decentralization as one of the main features of a competitive blockchain.

In a second step, the thesis conducted a comprehensive data analysis of real-world data from the Casper blockchain to evaluate how the network has evolved in terms of income distribution and decentralization measures. The initial token distribution suggested a high degree of centralization, as the Casper Association and insiders alone received over 50% of the initial token allocation. However, these initial values are in a similar range as comparable projects (Internet Computer, Avalanche, Flow; see Buterin, 2021). The descriptive and exploratory data analyses based on the data from the blockchain confirm this assumption: thus, over the entire observation period of 5446 eras, only 15 addresses have received over 50% of all rewards generated to date. Furthermore, for the largest 10 validators (out of a maximum of 100 active validators), on average over 50% of the total staking balances are staked.

Due to the high centralization of the initially distributed tokens, a high Gini coefficient of the staking rewards of all participants of approx. 0.75 was already measured in the first 500 eras, which after a steep growth phase has meanwhile stabilized at a very high value of 0.97, indicating a very high inequality in the entire system. The Gini coefficient of validator incomes also began at very high values of up to 0.92, but has been on a steady downward trend since era 1000 and was only 0.76 at the end of the period under review at era 5446, indicating a trend toward decentralization in the system, as delegates increasingly distribute their shares among smaller validators. The Nakamoto coefficient paints a similar picture, rising from an initially very low value of 3 to a value of 7, which also indicates increasing decentralization, although still at a low level. Furthermore, the information entropy of income was calculated, which is a measure of the randomness and disorder in a system. Unlike the Gini coefficient, which evaluates arbitrarily large networks with similar wealth distributions with equal values, the information entropy also allows the size of the system to be taken into account through additional participants, since larger systems are more difficult to coordinate and can therefore be described as more decentralized. This value also shows that the Casper network is becoming increasingly decentralized due to its growth in users over the course of the observed period. Another measure examined was expectational fairness in the Casper network: this describes whether a participant has received a fair share of the rewards for their efforts. Overall, high fairness values were found, since Casper's proportional distribution of rewards largely pays out the rewards to the participants corresponding to their shares of the total staking pool, regardless of who is the block-proposing validator. Furthermore, no evidence was found that identical participants received a disproportionate high number of rewards for longer periods in a row, which is usually an indication of engaging in selfish behavior.

Although the work conducted has provided fundamental insights into the Casper network and has applied several different measures to provide a holistic view, there are several limitations to consider. The applied agent-based models are only able to represent a simplified subset of the real PoS networks. For instance, a constant selection of validators and associated delegators is assumed, although in reality these change frequently.

Furthermore, fixed uniform validator fees are assumed, which contradicts the real dynamics in a free market. Moreover, a closed system is assumed, in which all rewards are restaked and no new stakes from other sources are possible. Finally, a model was described as a basic PoS reward distribution - but a uniform definition of a common PoS system is actually not possible, because the systems differ in many properties and therefore an application to an existing blockchain is only possible with adaptations. In the applied data analysis, despite the use of independently sourced primary data directly from active nodes of the Casper blockchain, the natural properties of the blockchain limit the data quality: the entire analysis is largely based on the strong assumption that a single user or entity owns only a single wallet address. However, because a blockchain does not specify who can create how many addresses and addresses cannot be uniquely assigned to a user, it is likely that, especially for larger token holders such as the Association or venture capitalists, larger amounts of tokens will be split among many smaller amounts on different wallet addresses. Besides privacy concerns, security concerns alone would be conclusive arguments for such an approach. In addition, the measures used are also only indications, which despite the use of several variations can only convey part of the information. Finally, it should be noted that the Casper network has only been in operation for just over a year at the time of the analysis and is therefore still very young. A look at other blockchains (cf. Ethereum or Bitcoin) shows that early benchmarks on decentralization and developments do not necessarily determine the long-term future of a blockchain.

Future research could empirically test the findings from the agent-based models in an experiment within a real blockchain. Moreover, the models could be adapted to investigate other reward functions besides the linear reward function, such as geometric rewards, in conjunction with the two reward distribution mechanisms that were modeled. In addition, more in-depth investigations of transaction flows could be undertaken to determine which addresses are most likely to belong to the same users, which would improve data quality and thus validity. Furthermore, only the reward mechanism, the rewards and the staking balances were investigated: by extracting all transactions, all active addresses and corresponding liquid, non-staked funds in each era could be determined, which would enable an analysis of the entire assets in the network.

7 Conclusion

Basic properties of the reward mechanisms of the Casper network and basic PoS networks were elaborated in order to develop agent-based models based on them, which were then used to simulate the exchange processes of the reward system of the consensus mechanism. The results indicate that although the Casper reward mechanism leads to inequality in the long run, it is fairer and leads to a slower centralization of wealth compared to common PoS networks. Furthermore, the work extracted real-world data from the Casper blockchain and examined it in detail using a number of inequality and decentralization measures. The results suggest that while the Casper network as a whole exhibits very high wealth inequality, the network is beginning to show signs of decentralization as the network grows larger and delegators are increasingly distributed among smaller validators. In addition, the Casper network is characterized by high fairness values in the distribution of income in relation to effort, since participants largely receive the share of rewards in the form of interest that they can expect from their investment in form of their staking balance.

References

AscendEx. (2020, March 25). *Announcement of Casper validator sale*. https://ascendex.com/en/support/articles/154-announcement-of-casper-clx-validator-sale

Buterin, V. (2021, June 29). *Against overuse of the Gini coefficient*. https://vitalik.ca/general/2021/07/29/gini.html

Campajola, C., Cristodaro, R., De Collibus, F. M., Yan, T., Vallarano, N., & Tessone, C. J. (2022). *The evolution of centralisation on cryptocurrency platforms*. https://doi.org/10.48550/ARXIV.2206.05081

Casper Metrics. (n.d.). *Caspermetrics*. https://caspermetrics.io/api

CasperLabs. (2020, April 10). *Highway: A cartel-resistant consensus protocol*. https://casperlabs.io/blog/highway-a-cartelresistant-consensus-protocol

CasperLabs. (2021, April 5). *New power usage report shows the Casper network's impressive energy efficiency relative to other blockchain protocols*. https://blog.casperlabs.io/new-power-usage-report-shows-the-casper-networks-impressive-energy-efficiency-relative-to-other-blockchain-protocols/

CasperStats Docs. (n.d.). *Token metrics & distribution*. https://docs.casperstats.io/docs/token-metrics-and-distribution

CoinMarketCap. (n.d.). *Casper price today, CSPR to UDS live, marketcap and chart*. https://coinmarketcap.com/de/currencies/casper/

Cosmos, C. (2021, Mai 27). *Why blockchains need cosmos proof-of-stake for a sustainable environment*. https://blog.cosmos.network/why-blockchains-need-cosmos-proof-of-stake-for-a-sustainable-environment-878b3edd2e85

Digiconomist. (2022). *Bitcoin energy consumption index*. https://digiconomist.net/bitcoin-energy-consumption/

Dimitri, N. (2021). Monetary dynamics with proof of stake. *Frontiers in Blockchain, 4*. https://doi.org/10.3389/fbloc.2021.443966

Eyal, I., & Sirer, E. G. (2014). *Majority is not enough: Bitcoin mining is vulnerable* (p. 436–454). Springer. https://doi.org/10.1007/978-3-662-45472-5_28

Fanti, G., Kogan, L., Oh, S., Ruan, K., Viswanath, P., & Wang, G. (2019). *Compounding of wealth in proof-of-stake cryptocurrencies* (I. Goldberg & T. Moore, eds.; vol. 11598, p. 42–61). Springer International Publishing. https://doi.org/10.1007/978-3-030-32101-7_3

Gilad, Y., Hemo, R., Micali, S., Vlachos, G., & Zeldovich, N. (2017). *Algorand: Scaling byzantine agreements for cryptocurrencies*. 51–68. https://doi.org/10.1145/3132747.3132757

Huang, Y., Tang, J., Cong, Q., Lim, A., & Xu, J. (2021). *Do the rich get richer? Fairness analysis for blockchain incentives*. 790–803. https://doi.org/10.1145/3448016.3457285

Jia, Y., Xu, C., Wu, Z., Feng, Z., Chen, Y., & Yang, S. (2022). *Measuring Decentralization in Emerging Public Blockchains*. 137–141. https://doi.org/10.1109/IWCMC55113.2022.9825341

Kane, D., Fackler, A., Gągol, A., & Straszak, D. (2021). *Highway: Efficient consensus with flexible finality*. arXiv. https://doi.org/10.48550/arXiv.2101.02159

Li, S.-N., Yang, Z., & Tessone, C. J. (2020a). *Mining blocks in a row: A statistical study of fairness in Bitcoin mining*. 1–4. https://doi.org/10.1109/ICBC48266.2020.9169436

Li, S.-N., Yang, Z., & Tessone, C. J. (2020b). *Proof-of-Work cryptocurrency mining: A statistical approach to fairness*. 156–161. https://doi.org/10.1109/ICCWorkshops49972.2020.9209934

Lin, Q., Li, C., Zhao, X., & Chen, X. (2021). *Measuring decentralization in bitcoin and ethereum using multiple metrics and granularities*. 80–87. https://doi.org/10.1109/ICDEW53142.2021.00022

MacKenzie, S. (2021, June 15). *China is kicking out more than half the world's bitcoin miners – and a whole lot of them could be headed to Texas*. CNBC. https://www.cnbc.com/2021/06/15/chinas-bitcoin-miner-exodus-.html

Neuder, M., Moroz, D. J., Rao, R., & Parkes, D. C. (2020). *Defending against malicious reorgs in tezos proof-of-stake*. 46–58. https://doi.org/10.1145/3419614.3423265

Nguyen, C. T., Hoang, D. T., Nguyen, D. N., Niyato, D., Nguyen, H. T., & Dutkiewicz, E. (2019). Proof-of-stake consensus mechanisms for future blockchain networks: Fundamentals, applications and opportunities. *IEEE Access*, *7*, 85727–85745. https://doi.org/10.1109/ACCESS.2019.2925010

Roşu, I., & Saleh, F. (2021). Evolution of shares in a proof-of-stake cryptocurrency. *Management Science*, *67*(2), 661–672. https://doi.org/10.1287/mnsc.2020.3791

Sai, A. R., Buckley, J., & Le Gear, A. (2021). Characterizing wealth inequality in cryptocurrencies. *Frontiers in Blockchain*, *4*, 730122. https://doi.org/10.3389/fbloc.2021.730122

Saleh, F. (2021). Blockchain without waste: Proof-of-stake. *The Review of Financial Studies*, *34*(3), 1156–1190. https://doi.org/10.1093/rfs/hhaa075

Srinivasan, B. S., & Lee, L. (2017, July 28). *Quantifying decentralization*. https://news.earn.com/quantifying-decentralization-e39db233c28e

Tendermint. (n.d.). *What is Tendermint*. https://docs.tendermint.com/master/introduction/what-is-tendermint.html

Wang, J., & Ge, L. (2022). *Consensus algorithm of proof-of-stake based on credit model*. 88–94. https://doi.org/10.1145/3532640.3532652

Wang, Y., Yang, G., Bracciali, A., Leung, H., Tian, H., Ke, L., & Yu, X. (2020). Incentive compatible and anti-compounding of wealth in proof-of-stake. *Information Sciences*, *530*, 85–94. https://doi.org/10.1016/j.ins.2020.03.098

Wu, K., Peng, B., Xie, H., & Huang, Z. (2019). *An information entropy method to quantify the degrees of decentralization for blockchain systems*. 1–6. https://doi.org/10.1109/ICEIEC.2019.8784631

Zamfir, V. (2017, February 11). *Casper the friendly ghost: A „correct-by-construction" blockchain consensus protocol*. https://github.com/ethereum/research/blob/master/papers/CasperTFG/CasperTFG.pdf

Appendices

Appendix A Additional Table

Table 3 List of updates

Update	Time	Era	Reference
1.1.0	22/04/2021	347	https://docs.cspr.community/docs/ubuntu/setup-mainnet-validator-from-scratch.html
1.1.2	11/05/2021	574	https://docs.cspr.community/docs/ubuntu/setup-mainnet-validator-from-scratch.html
1.2.0	28/05/2021	694	https://docs.cspr.community/docs/ubuntu/setup-mainnet-validator-from-scratch.html
1.2.1	13/07/2021	1281	https://docs.cspr.community/docs/ubuntu/setup-mainnet-validator-from-scratch.html
1.3.1	30/07/2021	1346	https://github.com/make-software/how-to-casper-labs/blob/master/docs/testnet/upgrade-1_3_1.md
1.3.2	12/08/2021	1605	https://docs.cspr.community/docs/ubuntu/setup-mainnet-validator-from-scratch.html
1.3.4	14/10/2021	2193	https://docs.cspr.community/docs/ubuntu/setup-mainnet-validator-from-scratch.html
Outage	03/11/2021	2595	https://github.com/casper-network/casper-node/wiki/Postmortem#mainnet-outage-2021-11-03-1345-utc
1.4.1	02/11/2021	2600	https://docs.cspr.community/docs/ubuntu/setup-mainnet-validator-from-scratch.html
Outage	22/11/2021	2819	https://github.com/casper-network/casper-node/wiki/Postmortem#mainnet-outage-2021-11-03-1345-utc
1.4.4	14/01/2022	3435	https://docs.cspr.community/docs/ubuntu/setup-mainnet-validator-from-scratch.html
1.4.5	04/04/2022	4417	https://docs.cspr.community/docs/ubuntu/setup-mainnet-validator-from-scratch.html
1.4.6	20/05/2022	4968	https://docs.cspr.community/docs/ubuntu/setup-mainnet-validator-from-scratch.html

Appendix B Code Implementation

The Python programming language was used to implement the data acquisition, the data analysis including all used measures and the and agent-based simulations of the reward mechanisms. The data acquisition from the Casper node was solved via shell scripts retrieved from the Python Jupyter notebook. The notebook "Casper_Analysis.ipynb" contains all steps for data acquisition, database creation, descriptive and explorative data analysis and calculation of the used measures. The notebook "Casper_Simulation.ipynb" contains the implementation of the agent-based model and the executed simulations.

Requirements: In addition to Python and the necessary libraries (defined in the respective Jupyter notebooks under "imports"), a complete installation of the Casper client is necessary in case a new data acquisition is to be carried out. A documentation for the installation of the Casper client is available at: https://docs.casperlabs.io/workflow/setup/

The Casper client is not needed when working exclusively with the already collected raw files or with the finalized database. Below are the links to three project folders, one without collected data (Casper client needed for data collection), and one each with the raw files without database and with the finalized database without raw files. For a direct re-execution of the code for the calculation of the measures, only the project folder with the finalized database without raw files (3.) is necessary.

1. Code only (1.6MB):
   https://gitlab.uzh.ch/thomas.levrand/casper_masterthesis_thomaslevrand/-/blob/master/CSPR_Analysis_nodata.zip
2. Code with raw files, without database (7GB):
   http://thomas.levrand.ch/casper_thesis/CSPR_Analysis_with_rawdata.zip
3. Code with database, without raw files (9GB):
   http://thomas.levrand.ch/casper_thesis/CSPR_Analysis_with_DB.zip

Statutory Declaration

I hereby declare that the thesis with title

"Wealth Inequality in CasperLabs' Proof-of-Stake Blockchain"

has been composed by myself autonomously and that no means other than those declared were used. In every single case, I have marked parts that were taken out of published or unpublished work, either verbatim or in a paraphrased manner, as such through a quotation. This thesis has not been handed in or published before in the same or similar form.

Dübendorf, 15.08.2022 Thomas Levrand