

Analyzing the Fairness of Proof of Stake Ethereum

Stefano Bistarelli¹, Cosimo Laneve², Ivan Mercanti¹ and Adele Veschetti³

¹University of Perugia, Italy

²Department of Computer Science and Engineering, University of Bologna, Italy

³Department of Computer Science, TU Darmstadt, Germany

Abstract

The paper presents an analysis of fairness within the Proof of Stake (PoS) Ethereum protocol Gasper, utilizing simulations conducted with the PRISM+ modelling tool. Our study focuses on stake analyses, revealing that stake growth dynamics are tangled to the initial distribution of the stake. Furthermore, we investigate the impact of malicious validators on stake growth dynamics and fairness property in the Gasper protocol. Through our analyses, we highlight the significant influence of the initial stake distribution on stake growth and fairness metrics. Our findings contribute to a deeper understanding of the Gasper protocol's performance under varying conditions, aiding in developing strategies to enhance fairness and security within PoS Ethereum networks.

Keywords

Gasper, Proof of Stake, stochastic modeling and analysis

1. Introduction

Blockchain technology offers a decentralized and immutable ledger system that enables secure and transparent peer-to-peer transactions without the need for intermediaries. By leveraging cryptographic techniques and consensus mechanisms, blockchains provide a tamper-proof record of digital assets and transactions. In the previous years, blockchain technology has found success in various applications, including managing cryptocurrencies (with Bitcoin being a prominent example [1]), facilitating decentralized applications (such as Ethereum smart contracts [2]), deploying voting systems [3], and enabling Decentralized Finance (DeFi) initiatives [4].

Traditionally, inspired by Nakamoto's seminal work [1], blockchain protocols have relied on a probabilistic mechanism known as Proof of Work (PoW), requiring nodes to solve complex computational puzzles to update the ledger. However, PoW's significant drawback lies in its high demand for computational resources and energy consumption [5]. In response, alternative proposals, notably Proof of Stake (PoS), have emerged. In PoS, nodes can update the ledger based on the quantity of cryptocurrency they have invested, referred to as their *stake*.

Fairness in blockchain protocols is a critical consideration encompassing various aspects of decentralized systems, including participation opportunities, reward distribution, and governance mechanisms. Ensuring fairness in blockchain protocols involves designing mechanisms that promote equitable access and representation for all participants, regardless of their resources or status. This includes addressing issues such as concentration of mining or staking power, economic incentives that may favor certain stakeholders, and the impact of protocol parameters on network decentralization.

In this paper, we present an analysis of the fairness of the proof of stake protocol employed within the Ethereum blockchain ecosystem, called Gasper, focusing on the stake distribution among the network nodes. Following the approach of [6, 7, 8], we use PRISM+¹, an extension of PRISM² with dynamic data types, allowing for precise modelling of complex protocols with varying stake dynamics. We then present simulation results derived from our model, illustrating that the validators consistently maintain their positions with respect to the wealth in terms of stake. By focusing on these aspects, our study

DLT 2024: 6th Distributed Ledger Technologies Workshop, May, 14-15 2024 – Turin, Italy

✉ stefano.bistarelli@unipg.it (S. Bistarelli); cosimo.laneve@unibo.it (C. Laneve); ivan.mercanti@unipg.it (I. Mercanti); adele.veschetti@tu-darmstadt.de (A. Veschetti)



© 2024 Copyright for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).

¹<https://github.com/adeleveschetti/ethereum-analysis>

²<https://www.prismmodelchecker.org/>

contributes to a deeper understanding of PoS protocols and their role in fostering a fair and resilient blockchain ecosystem. We notice that, in this context, fairness may also address transaction inclusion, ensuring that transactions are selected to be inserted in blocks without favoring any specific group of validators or users. Our analysis overlooks this issue because the PRISM+ model considers blocks as a whole, disregarding the inner transactions.

The paper is structured as follows: Section 2 provides background information on the proof of stake protocol used in Ethereum. Section 3 presents our PRISM+ model, detailing its construction and key features. In Section 4, we present the results of simulations conducted using our model, showcasing the dynamics of validator stake distribution. Section 5 reviews related works in the field of blockchain protocol fairness analysis. Finally, Section 6 presents the conclusion of our study, summarizing key findings and outlining avenues for future research.

2. Background

The Gasper ledger is structured as a tree of blocks, with a pointer to a leaf block at the maximal depth, referred to as the "handle". The sequence of blocks from the handle to the root is known as the "blockchain", with the initial block being the "genesis block". Each block in the ledger is assigned a height, representing the length of the path from that block to the genesis block.

Gasper's operation involves three primary steps: (i) block creation: the selection of the validator responsible for proposing a new block, followed by the creation of the block and its addition to the ledger; (ii) finalization mechanism: the finalization of block storage in the blockchain, ensuring irreversibility, and the initiation of the voting process for blocks (iii) incentives and penalties: the implementation of the incentive mechanism, which rewards honest validators and penalizes those who misbehave. In the following paragraphs we discuss in some detail the foregoing steps.

Block creation. In Gasper, the process of block creation involves validators, which are nodes selected to create new blocks and validate transactions based on the amount of coins they hold as collateral in the network, known as the *stake*. Validators are required to stake at least 32 ETH to register in the Gasper smart contract and receive a unique index.

During each fixed time interval of 12 seconds, referred to as a *slot*, a single validator is chosen to propose a block. If the designated proposer is offline during their slot, no new block will be generated for that slot, causing other validators to wait until the end of the 12-second interval to proceed with a new block. Otherwise, the proposer validator creates a block by collecting transactions and a block seed, derived from the hash of the validator index and the epoch seed.

While the time frame for proposing a block is limited to the current slot, network latency may result in several blocks being received in the same slot, leading to a fork in the ledger with multiple blocks at the same height.

An epoch consists of 32 consecutive slots, potentially containing fewer than 32 blocks. The blocks at the beginning of an epoch are called *checkpoints* or *epoch boundary blocks*.

The validator selection process is defined by the RANDAO smart contract³. The process, at the start of each epoch, combines the seeds stored by validators in the blocks at epoch e with the epoch number e to create the epoch seed for epoch $e + 1$. Then, this value is utilized by the RANDAO smart contract at epoch $e + 1$ to generate a pseudo-random sequence of 32 validator indexes with stakes greater than 32 ETH, which will propose the blocks for epoch $e + 3$.

Finalization mechanism. Finality is a crucial attribute of checkpoints that ensures their permanence, extending this characteristic to all blocks in the blockchain with lower heights. This is achieved through a two-step process. Initially, a checkpoint must be *justified* by gaining approval from at least two-thirds of the total staked ETH, indicating a high level of confidence in its inclusion in the canonical chain. Subsequently, the process of "finalization" occurs when another checkpoint is justified on top of a

³<https://github.com/randao/randao>

previously justified block. This action solidifies the commitment to include the ancestor block in the canonical chain, rendering it irreversible.

Additionally, each validator’s vote for a checkpoint, contained within a message called an *attestation*, also includes a vote for a block. This vote aids in resolving forks in the ledger. Specifically, for each block vote, a weight proportional to the validator’s stake is added to every block in the chain that has the voted block as a descendant. In the event of forks, the LMD-Ghost algorithm identifies the main chain by selecting the one with the highest weight.

Incentives and penalties. Validators earn rewards through activities such as consistent voting alignment with the majority of other validators during checkpoints and block proposals. The reward values are determined within each epoch based on the `base_reward`, representing the average reward a validator would receive under optimal conditions per epoch. This unit is proportional to the validator’s effective balance and inversely proportional to the total number of validators on the network.

Validators incur penalties if they vote differently from the majority or fail to send attestations within the inclusion delay, resulting in the same penalty amount as the corresponding reward. For instance, a validator receives a penalty if he misses the inclusion delay of 32 slots (384 seconds). Votes beyond the same checkpoint epoch are not allowed and are penalized.

Gaspar also implements a harsher penalty mechanism known as slashing, which results in expelling a validator from the network and confiscating their stake. Validators face slashing in two scenarios: proposing and endorsing two different blocks for the same slot or engaging in double voting for checkpoints. Upon detection of these actions, the validator incurs a slashing penalty of 1/32 of their stake, immediate burning of the penalty, and a ban from the network for 36 days, during which additional penalties may be applied.

Fairness. A key issue in PoS protocols is to have a *fair distribution of the possibility of generating blocks* among validators that is resistant to attacks such as the “nothing-at-stake” [9] and the “branching process attack” [10]. In the current version of Gaspar, the ability to generate blocks is proportional to the wealth on stakes of a validator. Therefore a validator with higher stake has more chance to be selected as a proposer. On the other hand, the validator’s mining increases his own stake, thus making it higher and higher. As a consequence, the distance in wealthiness augments, thus reducing the ability of participants with smaller stakes to create blocks. This paper tries to measure this proportion by using a technique based on stochastic model checking.

3. The Proof of Stake Model

In [6], we presented PRISM+, an extension of the model checker PRISM with native support for expressing and manipulating dynamic data types, such as lists and trees, and data types specifically designed for modeling blockchain protocols such as block and ledger. In our model, blocks are triples $(v^n; p; h)$, where v is the *name* of the validator v who created the block; n is a unique numeric label; p is the name of the ancestor block which $(v^n; p; h)$ points to; h is the *height* of the block in the ledger.

```

1  module Vote_Manager
2      Stakes : map {} ;
3      Votes : map {} ;
4      epoch = 0 ;
5      for i from 0 to N:
6          Stakes[validator_i] : [0..MAX_STAKE] init STAKE_i;
7      for i from 0 to N:
8          [voteB_i] -> 1 : Votes' = addVote(Votes, b_i, Validator_i);
9          [finB_i] (height(lastF(L_i)) > epoch) ->
10             1: epoch' = height(lastF(L_i)) Stakes' = updateS(Stakes, Votes, lastF(L_i));
11          [finB_i] (height(lastF(L_i)) <= epoch) -> 1: ;
12 endmodule

```

Listing 1: The code of `Vote_Manager`.

In turn, ledger are tuples $\langle T; f; p \rangle$ where T is a tree of blocks; f is the name of a block in T (the *last finalized block* of L) and p is the *handle* of L , *i.e.* the name of a leaf block at maximal height in the subtree rooted at f .

We modelled Gasper in PRISM+ as a parallel composition of different modules: `Validators`, `Network`, `Vote_Manager`, `Randao`, `RandaoSelection` and `Global`. Each module plays a critical role in simulating and analyzing the behavior of our system, contributing to the overall robustness and reliability of our implementation.

In particular, the `Network` module manages the entire part relating to adding or eliminating a block to the main blockchain, considering the addition or removal delays foreseen by the protocol used. Moreover, the `Randao` and `RandaoSelection` model the random selection of the next block proposers. We offer a succinct overview that emphasizes essential points of the modules that we take into account: the `Validator` and the `Vote_Manager` (Listing 1).

The `Validator` module may undertake one of the following actions: (i) create a new block if the RANDAO algorithm has selected it; (ii) if selected to vote in this slot, transition to a voting phase; (iii) receive a new block from the network. The `Vote_Manager` module is used to initialize, track and update the stakes of each validator belonging to the network.

The transitions involved in this module are repeated for each validator, as can be got from the pseudocode. The transition `voteB_i` (line 8 of Listing 1) adds a vote inside the hash map for that specific block passed as input. Much more complex is the `finB_i` transition (line 9), which, after a careful analysis of the consistency regarding the height of that block in relation to the other checkpoints or finalized blocks, allows you to update the stake value of the single validator and consequently also that of the total stake of the network.

4. Simulations

The section reports the experiments performed in order to test the stake growth for the Gasper protocol. The experiments are conducted with 13 validators, which is the same number used in Hybrid Casper tests [6]. We notice that this limited number of validators does not affect our experiments in a sensible way: when the number increases, the overall trend does not change (in case of 16 validators, the differences are in the order of 10^{-3}). For all the upcoming experiments, we consider the tie interval needed to create ~ 1200 blocks (12 s for block) and a reward of 10% of the initial stake.




Figure 1: Stake growth

Figure 1 reports the analysis of the total stake growth within the Gasper system. We examine two distinct scenarios: one where the network operates without any malicious validators (blue line), and another where 40% of validators exhibit malicious behavior (red line), *i.e.* they vote for a different block than the one the majority of validators has voted for. The visualization clearly illustrates that the presence of attackers leads to a slower growth rate of the stake. These results are consistent with those presented in [11], where it was observed that, in the absence of misbehaving validators, the final stake increased by a factor of 21 from its initial value, whereas in the latter scenario, the increase was limited to a factor of 5.

We analyze the fairness of the Gasper protocol by means of simulations of the stake growth. The

aim is to highlight how the distribution of stakes evolves and its impact on equity among nodes of the network. Figure 2a illustrates the growth of the stake for a validator initially equipped with the maximum percentage of stake (52 ETH); the line in Figure 2b describes the same analysis for a validator that initially owns the minimum amount of stake (32 ETH). It is worth to observe that the initial stake allocations influence in a relevant way the stake growth dynamics within the system, with richer validators exerting a more significant influence on the growth of the stake. This behaviour is corroborated by the corresponding Gini coefficient [12], which yields a value of 0.366 in this instance. Since this coefficient measures wealth inequality in the system, with higher values indicating greater




Figure 2: Growth of the stake for the validator with the highest (a) and lowest (b) initial stake amount

inequality. The Gini coefficient of 0.366 suggests a moderate level of wealth inequality among validators and, thus, of unfairness.

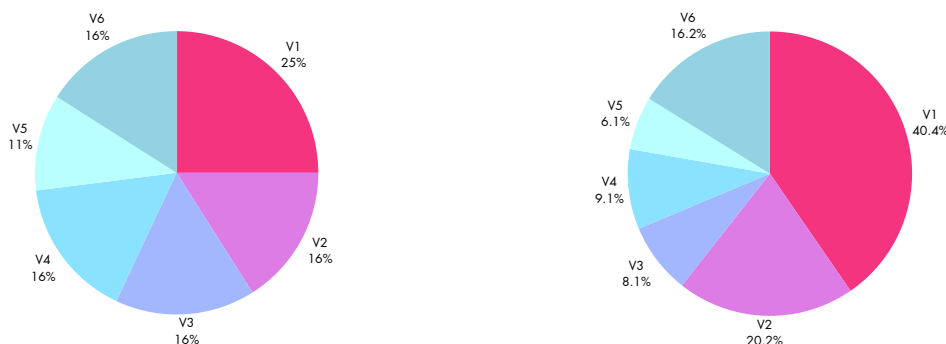


Figure 3: Initial (left) and final (right) stake distribution

The contrast between the initial and the final stake distribution is illustrated in Figure 3 where we report the percentage of stake owned by 6 validators. The trend is clear: the wealthiest (V_1) and the poorest (V_5) validators maintain their respective positions at the conclusion of the execution. Furthermore, it is noteworthy that validators V_2 , V_3 , V_4 , and V_6 , which initially had identical stake distributions, exhibit variations in their stake distributions at the end of the execution. This divergence can be attributed to their voting behaviour and the delay in receiving blocks.

The fairness is then analyzed in Figure 4, which shows up the disparity in block creation between validators with varying stake levels across multiple epochs. The x -axis represents the number of epochs, while the y -axis displays the number of blocks created. The two lines illustrate the trend: one depicting the number of blocks created by validators with more stake, and the other indicating the number of blocks created by validators with less stake. The figure demonstrates a consistent pattern wherein, across epochs, validators with higher stakes consistently produce more blocks compared to validators with lower stakes.




Figure 4: Number of blocks created by the wealthiest validator (V_0) and the poorest (V_5)

5. Related Works

In blockchain systems the concept of fairness is intricately tied to the wealth distribution among network nodes. Therefore the analysis of miners' returns helps understanding the equilibrium between investment capacity and equal opportunities. For this reason, several studies have investigated the wealth distribution among the wealthiest users of PoW and PoS-based blockchains [13, 14, 15]. Notably, in [16], it was found that the wealth of top Bitcoin holders increases at a faster rate compared to accounts with lower balances, a phenomenon commonly known as preferential attachment.

A detailed analysis of the wealth distribution among the wealthiest accounts in different blockchain systems is done in [17]. In particular, the article examines the temporal evolution of statistical metrics and points out differences in wealth centralization, indicating that tokens tend to be more centralized than coins.

A different approach is taken in [11]. Instead of analyzing existing data on cryptocurrency in blockchains, the study examines the conditions under which condition a PoS-based consensus algorithm can achieve a fair wealth distribution over time. By "fair", it is meant that individuals with more wealth have a higher chance of being selected as validators, but their wealth does not increase (or decrease) solely due to a validation activity. In other words, validation alone should not affect anyone's cryptocurrency holdings. To achieve this, the study looks beyond the top 30-50-100 richest cryptocurrency holders and considers wealth distribution among all blockchain users, specifically those aspiring to be selected as block validators. Leporati analyses two different settings: one with 1% and another with 40% of corrupted nodes. In the first case, the simulations showed that, at the end of the run, both the richest and the poorest peers have increased their cryptocurrency holdings proportionally to their initial stake (as a consequence, the risk of diminishing participation interest in the protocol is sensible). In the second setting, the analysis revealed that the only factor draining cryptocurrency from the system is penalizing corrupted peers. Allowing this trend to continue could lead to corrupted peers exhausting their funds, potentially reinstating the initial issue of wealth distribution among honest peers.

Other contributions have suggested substantial changes to the PoS protocol that aim at improving its fairness and long-term sustainability. An extended form of PoS, termed as e-PoS, is presented in [18]. This e-PoS aims at introducing fairness in the blockchain network and resisting to centralization. The system uses a smart contract that is introduced to run atop the blockchain, facilitating a blind block auction. In particular, the smart contract applies policies that extend mining opportunities to a wider set of network peers and that ensure fair reward distribution. Another solution to the fairness problem is studied in [19], where a new Robust Proof of Stake (RPOS) consensus protocol is proposed. This protocol utilizes the amount of coins to select miners and imposes a maximum value on the coin age to effectively mitigate coin age accumulation attacks and Nothing-at-Stake attacks.

6. Conclusions

In this study, we thoroughly analysed the fairness of the proof of stake protocol employed in the Ethereum blockchain. Our investigation focused on examining the equitable distribution of stakes

among validators, using a combination of modelling and simulation techniques.

By developing a PRISM+ model specifically tailored to capture the dynamics of the PoS protocol, we have gained valuable insights into the behaviour of validators and the distribution of stakes over time. Our simulations reveal that the wealthiest validators consistently maintain their advantageous positions while the least affluent validators struggle to improve their standings. The simulations deepen our understanding of fairness in PoS protocols, revealing challenges and disparities in stake distribution. This emphasizes the need to design and implement mechanisms for greater equity and inclusivity in blockchain ecosystems.

Block optimizations might boost (small) validators revenues and therefore impact on the fairness of the protocol. In particular, since validators' revenues depend on the content of the blocks, modelling block optimizations would require a thorough extension of the PRISM+ model. For example, one well-known technique is the MEV-boost [20] that separates the proposers of the blocks from the nodes that build them. Builder nodes utilize advanced algorithms to aggregate transactions and identify most profitable blocks. Extending our model with transactions, revenues, builder nodes, and optimization algorithms is a task for future work.

Moving forward, further research is needed to explore potential solutions that reduce the observed inequalities in stake distributions in different PoS protocols. The overall aim is to identify strategies that promote fairness and sustainability in PoS protocol over the long term.

Acknowledgments

S. Bistarelli and I. Mercanti are members of the Gruppo Nazionale Calcolo Scientifico-Istituto Nazionale di Alta Matematica (GNCS-INdAM). This work has been partially supported by: the ATHENE project "Model-centric Deductive Verification of Smart Contracts"; SERICS project (PE00000014) under the MUR National Recovery and Resilience Plan funded by the European Union – NextGenerationEU; GNCS-INdAM, CUP_E53C23001670001; European Union - Next Generation EU PNRR MUR PRIN - Project J53D23007220006 EPICA: "Empowering Public Interest Communication with Argumentation"; University of Perugia - Fondo Ricerca di Ateneo (2020, 2021, 2022) - Projects BLOCKCHAIN4FOODCHAIN, FICO, AIDMIX, "Civil Safety and Security for Society"; European Union - Next Generation EU NRRP-MUR - Project J97G22000170005 VITALITY: "Innovation, digitalisation and sustainability for the diffused economy in Central Italy"; Piano di Sviluppo e Coesione del Ministero della Salute 2014-2020 - Project I83C22001350001 LIFE: "the itaLian system Wide Frailty nEtwork" (Linea di azione 2.1 "Creazione di una rete nazionale per le malattie ad alto impatto" - Traiettorie 2 "E-Health, diagnostica avanzata, medical devices e mini invasività").

References

- [1] S. Nakamoto, Bitcoin: A peer-to-peer electronic cash system, <https://bitcoin.org/bitcoin.pdf>, 2008.
- [2] V. Buterin, Ethereum white paper, <https://github.com/ethereum/wiki/wiki/White-Paper>, 2013.
- [3] S. Bistarelli, I. Mercanti, P. Santancini, F. Santini, End-to-end voting with non-permissioned and permissioned ledgers, *J. Grid Comput.* 17 (2019) 97–118. URL: <https://doi.org/10.1007/s10723-019-09478-y>. doi:10.1007/s10723-019-09478-y.
- [4] E. Napoletano, J. Schmidt, Decentralized finance is building a new financial system, <https://www.forbes.com/advisor/investing/defi-decentralized-finance/>, 2021. (last access 2021).
- [5] C. C. for Alternative Finance, Cambridge bitcoin electricity consumption index, <https://cbeci.org/>, 2021. (last access 2021).
- [6] L. Galletta, C. Laneve, I. Mercanti, A. Veschetti, Resilience of hybrid casper under varying values of parameters, *Distributed Ledger Technol. Res. Pract.* 2 (2023) 5:1–5:25.
- [7] S. Bistarelli, R. De Nicola, L. Galletta, C. Laneve, I. Mercanti, A. Veschetti, Stochastic modeling and analysis of the bitcoin protocol in the presence of block communication delays, *Concurrency*

- and Computation: Practice and Experience (2021) e6749. doi:<https://doi.org/10.1002/cpe.6749>.
- [8] C. Laneve, S. Solmonte, A. Veschetti, A Stochastic Analysis of the Gasper Protocol, in: PerCom Workshops, 2024 *forthcoming*.
 - [9] W. Li, S. Andreina, J.-M. Bohli, G. Karame, Securing proof-of-stake blockchain protocols, in: J. Garcia-Alfaro, G. Navarro-Arribas, H. Hartenstein, J. Herrera-Joancomartí (Eds.), Data Privacy Management, Cryptocurrencies and Blockchain Technology, Springer International Publishing, Cham, 2017, pp. 297–315.
 - [10] P. Gaži, A. Kiayias, A. Russell, Stake-bleeding attacks on proof-of-stake blockchains, in: 2018 Crypto Valley Conference on Blockchain Technology (CVCBT), 2018, pp. 85–92. doi:10.1109/CVCBT.2018.00015.
 - [11] A. Leporati, Studying the compounding effect: The role of proof-of-stake parameters on wealth distribution, in: DLT, volume 3460 of *CEUR Workshop Proceedings*, CEUR-WS.org, 2023.
 - [12] J. Hasell, Measuring inequality: What is the gini coefficient?, 2023. <https://ourworldindata.org/what-is-the-gini-coefficient>.
 - [13] N. Dimitri, Monetary dynamics with proof of stake, *Frontiers Blockchain* 4 (2021) 443966.
 - [14] C. Li, B. Palanisamy, Comparison of decentralization in dpos and pow blockchains, in: ICBC, volume 12404 of *Lecture Notes in Computer Science*, Springer, 2020, pp. 18–32.
 - [15] B. Kusmierz, S. Müller, A. Capossele, Committee selection in DAG distributed ledgers and applications, in: SAI (2), volume 284 of *Lecture Notes in Networks and Systems*, Springer, 2021, pp. 840–857.
 - [16] D. Kondor, M. Pósfai, I. Csabai, G. Vattay, Do the rich get richer? an empirical analysis of the bitcoin transaction network, *CoRR* abs/1308.3892 (2013). URL: <http://arxiv.org/abs/1308.3892>. arXiv:1308.3892.
 - [17] B. Kusmierz, R. Overko, How centralized is decentralized? comparison of wealth distribution in coins and tokens, in: COINS, IEEE, 2022, pp. 1–6.
 - [18] M. Saad, Z. Qin, K. Ren, D. Nyang, D. Mohaisen, e-pos: Making proof-of-stake decentralized and fair, *IEEE Transactions on Parallel and Distributed Systems* 32 (2021) 1961–1973. doi:10.1109/TPDS.2020.3048853.
 - [19] A. Li, X. Wei, Z. He, Robust proof of stake: A new consensus protocol for sustainable blockchain systems, *Sustainability* 12 (2020) 2824.
 - [20] A. Wahrstätter, L. Zhou, K. Qin, D. Svetinovic, A. Gervais, Time to bribe: Measuring block construction market, *Cryptology ePrint Archive*, Paper 2023/760, 2023. URL: <https://eprint.iacr.org/2023/760>, <https://eprint.iacr.org/2023/760>.