Time Moves Faster When There is Nothing You Anticipate: The Role of Time in MEV Rewards

Burak Öz
burak.oez@tum.de
Technical University of Munich
Munich, Germany

Benjamin Kraner
benjamin.kraner@uzh.ch
University of Zurich
Zurich, Switzerland

Nicolò Vallarano
nicolo.vallarano@uzh.ch
University of Zurich
Zurich, Switzerland

Bingle Stegmann Kruger
krgbin001@myuct.ac.za
University of Cape Town
Cape Town, South Africa

Florian Matthes
matthes@tum.de
Technical University of Munich
Munich, Germany

Claudio Juan Tessone
claudio.tessone@uzh.ch
University of Zurich
Zurich, Switzerland

ABSTRACT

We present a novel analysis of a competitive dynamic present on Ethereum known as "waiting games", where validators can use their distinct monopoly position in their assigned slots to delay block proposals in order to optimize returns through Maximal Extractable Value (MEV) payments, a type of incentive outside the Proof-of-Stake incentive scheme. However, this strategy risks block exclusion due to missed slots or potential orphaning. Our analysis reveals evidence that, although there are substantial incentives to undertaking the risks, validators are not capitalizing on waiting games, leaving potential profits unrealized. Moreover, we present an agent-based model to test the eventual consensus disruption caused by waiting games under different settings, arguing that such disruption only occurs with significant delay strategies. Ultimately, this research provides in-depth insights into Ethereum's waiting games, illuminating the trade-offs and potential profit opportunities for validators in this evolving blockchain landscape.

CCS CONCEPTS

• **Security and privacy → Economics of security and privacy**.

KEYWORDS

blockchain, incentives, consensus, maximal extractable value

ACM Reference Format:
Burak Öz, Benjamin Kraner, Nicolò Vallarano, Bingle Stegmann Kruger, Florian Matthes, and Claudio Juan Tessone. 2023. Time Moves Faster When There is Nothing You Anticipate: The Role of Time in MEV Rewards. In *Proceedings of the 2023 Workshop on Decentralized Finance and Security (DeFi '23), November 30, 2023, Copenhagen, Denmark.* ACM, New York, NY, USA, 8 pages. https://doi.org/10.1145/3605768.3623563

DeFi '23, November 30, 2023, Copenhagen, Denmark

ACM ISBN 979-8-4007-0261-7/23/11…$15.00
https://doi.org/10.1145/3605768.3623563

1 INTRODUCTION

The transition from a Proof-of-Work (PoW)-based block proposer selection mechanism to a Proof-of-Stake (PoS) one, combined with the introduction of the Gasper consensus protocol [15], introduced new dynamics for incentives on Ethereum. Specifically, the block proposal rewards for validators, the consensus participants, underwent a significant decrease, as they no longer require the high incentives previously needed by PoW-Ethereum miners. As a result, validators now rely even more on transaction fees and exogenous rewards such as Maximal Extractable Value (MEV), to maintain their incentives and profitability.

Currently, the PoS-Ethereum protocol enforces a block release schedule of 12-second long slots, during which the designated validator has a total monopoly on the block release. This feature introduces new opportunities for MEV extraction through strategic *waiting games*. To illustrate, we hypothesize that the block proposer may exploit their monopolistic position and wait as long as possible within the slot limits to release the block in order to extract additional MEV. However, if a large proportion of validators on the Ethereum network engage in these delayed releases, it could result in instability within the consensus of the network, as delays in block release may lead to forked blocks.

This paper explores the dynamics of these waiting games and their impact on Ethereum's consensus and validator behavior. We examine the profitability of waiting games considering the relative significance of MEV rewards, assess the extent of current validator participation in these strategic dynamics, evaluate whether validators can safely engage without risking exclusion from the canonical chain, and investigate the potential of waiting games to destablize Ethereum's consensus along with the specific conditions that might trigger such instability. Our findings contributes to the research field by showing evidence of an untapped profit source for strategic validators and provide a simplified model presenting the safety to engage in waiting games from a protocol standpoint.

2 BACKGROUND AND RELATED WORK

This section delivers necessary background knowledge for the reader regarding Ethereum consensus, MEV, and the Proposer-Builder-Separation (PBS) design behind the MEV-Boost marketplace. Additionally, we provide a succinct review of prior research on waiting games in PoS systems.

2.1 PoS-Ethereum Consensus

Ethereum's Proof-of-Stake (PoS) network uses epochs and slots to manage consensus rounds, with each epoch comprising 32 slots, and each slot lasting 12 seconds. Network members who staked 32 Ether, known as validators, are randomly [11] distributed to slots for each epoch, either as block proposers or committee members who attest to the validity of the proposed block. Validators are incentivized by the rewards they receive for honest conduct, but can be penalized for misbehavior, such as proposing multiple blocks in a single slot. A block is finalized after it receives a super-majority of attestations in two epochs [15].

Validators follow the Latest Message Driven Greediest Heaviest Observed SubTree (LMD-GHOST) fork choice rule for determining the canonical chain [33], with the consensus mechanism operating in two phases: LMD-GHOST at the slot level and Casper Friendly Finality Gadget (Casper FFG) at the epoch scale [14]. The latter is a Byzantine fault-tolerant consensus protocol that finalizes blocks and ensures security even during temporary network partitions, with the confirmation rule outputting the most recent finalized block and its prefix.

2.2 Maximal Extractable Value

Maximal Extractable Value (MEV) has emerged as a prominent incentive within the Ethereum blockchain, particularly with the surge in Decentralized Finance (DeFi) domain since 2020 [42]. MEV represents all value that can be extracted within a blockchain network, which privileged actors can take advantage of by manipulating transaction sets and order in a block, beyond standard protocol incentives [19, 28]. Network participants can extract MEV by influencing transaction order through payments to block proposers, facilitated by markets like Flashbots Auction [10], and leveraging DeFi instruments and ordering techniques inspired by traditional finance [21, 29, 35, 41].

While current estimates of MEV are lower-bound due to heuristic limitations [7, 8, 28, 38], over $675 million worth of MEV had been extracted by September 2022 according to MEV Explore [8], and approximately $9.5 million in a single month from May 2023 to June 2023 according to EigenPhi [7].

MEV has been linked to user value loss, network congestion, consensus destabilizing attacks, and centralization risks [16, 19, 34, 39]. To mitigate these issues, solutions have emerged across system layers, including fair transaction ordering protocols, privacy-preserving mechanisms, efficient MEV extraction designs, and MEV-aware applications [1, 3, 4, 9, 12, 22–24, 39, 40].

2.3 Proposer-Builder Separation and MEV-Boost

Proposer-Builder Separation (PBS) is a design concept that mitigates the sophistication of validators in block construction by dividing tasks between proposing blocks and constructing them [13]. Validators profit from MEV without dealing with complex block building, as this competition is shifted to builders who are further away from the consensus layer. This setup, however, depends on addressing trust issues and commitments between proposers and builders [26]. In response, Flashbots introduced MEV-Boost [3], using relays as intermediaries and counteracting MEV's centralizing effects, albeit bearing censorship risks from relays withholding builder bids. However, the public auditability of relay actions fosters honest behavior.

The MEV-Boost architecture comprises validators, relays, and builders. Validators, controlling proposers, register with relays to receive execution payloads from builders, who prepare these using transactions from the public mempool and their private order flow [5]. Validators receive the highest-value header from their registered relays via MEV-Boost middleware, sign it and send it back to the relevant relay. Upon verifying the validator's signature, the relay propagates the complete block to the peer-to-peer network.

2.4 Related Work

As far as we are aware, [32] is the sole study exclusively focusing on waiting or timing games in PoS-Ethereum, although [37] also notes a positive correlation between bid arrival time and value, while exploring the intricate details about the block construction market available on Ethereum, MEV-Boost [3]. In [32], the authors present a model where honest-but-rational consensus participants can delay their block proposals to maximize MEV capture, while still ensuring their inclusion in the blockchain in a timely manner. Despite noting that timing games are worthwhile, they find these strategies not currently being exploited by consensus participants. In PoW protocols, a case of block release delay strategy is represented by selfish mining [30, 31] which have substantial differences to waiting games in PoS. While in selfish mining, the miner undermines the status of the ledger by purposely hiding blocks that eventually end up in the canonical chain, waiting games do not necessarily imply the substitution of any block. In addition, selfish mining usually takes place at a larger time scale than the PoS waiting games we study in the present work; while selfish mining spans on a multi-block level, PoS waiting games are played within individual slots.

3 DATA COLLECTION

We collected data from the MEV-Boost protocol and Ethereum consensus layer to study the impact of waiting games on validator payments, attestations, and consensus stability. Utilizing the public data endpoints provided by the MEV-Boost relays[1 2], we extracted information about builder bids and proposed blocks. We primarily concentrated on three dominant relays [36]: Ultra Sound, Flashbots[3], and Agnostic, and procured bid and proposed block data spanning a slot range from 6,087,501 to 6,100,000, equating to 12,500 slots during March 27 to 28, 2023. For consensus-related data such as attestations, forked blocks, and consensus rewards, we utilized the API endpoints provided by beaconcha.in[4]. Finally, to identify the builders, we used the builder overview on mevboost.pics [36]. We noted that some builders submit identical bids repetitively until the block auction ends and also to multiple relays to increase their selection likelihood. In our time-based bid value analysis, we focused on unique bids, discarding duplicates and recognizing the first arrival time of a bid on any relay as the initial contribution of that specific bid.

[1] https://flashbots.github.io/relay-specs/#/Data/getReceivedBids
[2] https://flashbots.github.io/relay-specs/#/Data/getDeliveredPayloads
[3] For Flashbots bids and blocks, we also used the data dumps available at https://flashbots-boost-relay-public.s3.us-east-2.amazonaws.com/index.html
[4] https://beaconcha.in/api/v1/docs/index.html

4 RESULTS

Our results contain two sections: initially, we delve into the value of waiting, questioning the worthiness of engaging in waiting games and whether validators already participate. Subsequently, we investigate the inherent risks associated with it and adopt an agent-based simulation model specifically designed for PoS-Ethereum to inspect how waiting strategies can influence the emerging consensus properties.

4.1 The Value of Waiting

4.1.1 The Evolution of Value Over Time. Figure 1 dissects the distribution of bids for slot 6 093 815, which are represented as dots and color-coded based on the submitting builder. The relation between their arrival time at the relay and their associated value is examined. The study uncovers a distinct upward trend in bid values over time. It is important to note that we measure the arrival time relative to the start time of the slot; therefore, any negative time value indicates that the bid arrived during the previous slot. With this in mind, the earliest recorded bid arrived at -10905ms with a value of 0.014 ETH. On the other hand, the winning bid chosen at 299ms had a much higher value of 0.046 ETH. This represents a substantial increase of approximately 228% from the initial bid, demonstrating the significant potential advantages of waiting. It is observed that different builders adopt distinct strategies, such as Flashbots builders submitting approximately every 0.5s since the start time of the previous slot, in contrast to rsync-builder.xyz or Bob the Builder, who only commence submissions after a specific point in time.

This increase in value is attributable to the expanding public transaction pool and potentially private order flow that builders observe over time. As regular users and MEV searchers successively submit new transactions and bundles, builders can access more opportunities to construct their blocks. As a result, they can offer larger payments to the validators. The escalating number of transactions included in the builder blocks further substantiates this (see Figure 7 in Appendix A).

In order to measure the incremental value gained for each millisecond of delay, we collected unique bids from each relay across 750 slots, yielding slightly over 480k unique bids. To accurately capture value progression, we residualized the bid values against slot and builder fixed effects, which might cause artificial value fluctuations due to high or low MEV regimes, as discussed in [32].

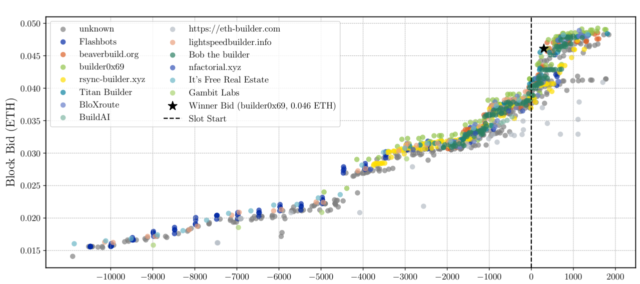

Figure 1: Distribution of unique builder blocks for slot 6 093 815 based on their arrival time to the relay and bid value, across Ultra Sound, Flashbots, and Agnostic relays. Each color denotes a unique builder.

Following this, we performed a linear regression on these residualized bid values relative to time, revealing a positive marginal value of 5.71×10^{-6} ETH/ms. This supports our initial observation from a single slot, reaffirming that by prolonging their wait time, validators can enhance their MEV payments.

4.1.2 The Significance of Waiting. Although we noted a positive marginal value of delay, if the rewards gained from waiting are insignificant compared to the consensus layer rewards issued for block proposals, then it might not be worth risking the consensus in the first place. To scrutinize the relationship between the rewards from waiting (i.e., MEV rewards) and the proposal rewards from consensus which comprises attestation and sync aggregate inclusion rewards [18], we examined 5,726 proposed blocks from the relays we have monitored. These blocks belong to three different epoch ranges between June 1, 2023, and June 11, 2023. While we used the MEV reward data provided by the relays, we fetched the consensus rewards information from the beaconcha.in API.

A prominent observation for all three epoch ranges is the dominance of the MEV rewards over proposal rewards (see Figure 8 in Appendix A). Overall, the MEV rewards accounted for 58.32% of all the rewards, while the median MEV rewards for analyzed epoch ranges were 0.067 ETH, 0.038 ETH, and 0.042 ETH, respectively, culminating in an overall median of 0.048 ETH. Meanwhile, the proposal rewards consistently stayed at 0.034 ETH. This led to a median difference of 0.013 ETH per proposed block, demonstrating the significance of MEV rewards over consensus rewards for the inspected slots.

While the analyzed epoch ranges highlight the prevailing value of MEV rewards, to draw more generalized conclusions, we compared the two rewards starting from the transition of Ethereum to PoS in September 2022 as this is when the waiting games on Ethereum emerged. We utilized the open data sets available at mevboost.pics [36] to find the individual MEV-Boost blocks' reward data. Given that proposal rewards are anticipated to remain steady at 0.034 ETH, we discovered approximately 1,290,528 blocks (59.07% of all MEV-Boost blocks) yielding larger MEV rewards than consensus rewards, with a median value of 0.053 ETH. Hence, we conclude that, as the generalized results also support our initial analysis of certain epoch ranges, the potential gains from waiting outweigh protocol rewards, affirming the value of risking consensus.

4.1.3 Unrealized Value. Schwarz-Schilling et al. [32] established that, currently, validators do not actively participate in the waiting games, and any delay observed primarily stems from the complex signing processes utilized by certain staking entities and validator clients. Our research supports their findings as we analyze the arrival time of winning builder bids in comparison to the highest value bid observed in that slot. Figure 2 portrays the distribution of both early winners (colored in green), those whose bids arrived prior to the highest bid, and late winners (colored in yellow), those whose bids arrived after the highest bid. This distribution is presented in relation to time and the value difference between the winning bid and the highest bid.

Our findings reveal that out of the 8,121 unique winning blocks processed by the relays we have studied, 7,672 (94.47%) had early winners and 269 (3.31%) of them had late winners. The remaining 180 blocks (2.22%) constituted the highest bid itself. The early

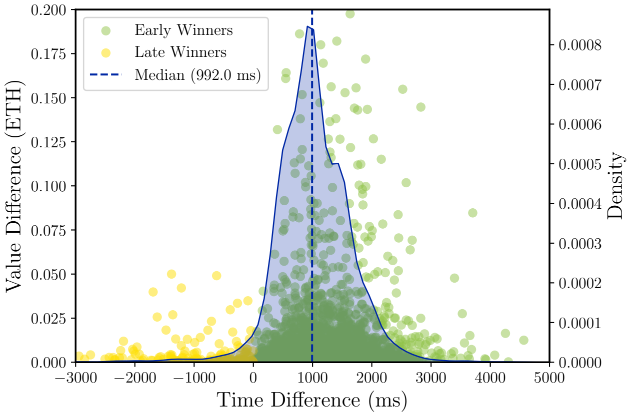

Figure 2: The scatter plot shows early (green) and late (yellow) winning bids' time and value deviations from the highest bid of the same slot. The density plot reveals a median time difference of 992ms when the winning bid was not the highest.

winners, on median, arrived 1001ms before the highest bid, albeit with a median value of 0.001 ETH less. The late winners, contrarily, arrived 392ms post the highest bid but still delivered 0.0008 ETH less value. In total, across the 7,941 blocks that did not capitalize on the highest bid, with a positive median time difference of 992ms (indicating the winning bid arrived first), a remarkable 931.27 ETH remained unrealized as the highest bid was not available when the winning bid had arrived. Conversely, 27.76 ETH was realizable as the highest bid had already arrived before the winning bid but still went uncaptured. These results reinforce our contention that validators are not engaging in the waiting games and often act prematurely in selecting the winning bid.

For future work, it would be beneficial to incorporate the *getHeader* call timestamp from relays, which indicates the exact moment the proposer requested the block header from the relay. By comparing this time with the arrival time of the highest value bid, we could obtain more accurate data about the unrealized value and the time difference between winning and highest value bids. As this data is not publicly accessible, we resorted to using bid arrival times (*receivedAt* timestamp), which still provide significant insight.

4.1.4 Playing the Game Rationally. Our research thus far has uncovered considerable incentives for rational validators within the Ethereum network to participate in waiting games. However, the manner in which these strategies are employed differs from relay to relay. In the default relay implementation [6], each builder bid undergoes a simulation to confirm its validity before it is made accessible to the proposer. This process introduces an average latency of 140ms [17], shortening the duration of the block auction and reducing the number of competing bids. To counter this latency, the optimistic relay design has been proposed [27], and adopted by the Ultra Sound Relay. Under the optimistic approach, it is presumed that builders are submitting valid blocks, which are instantaneously made available while the validation is delayed. However, this strategy necessitates that builders deposit funds upfront, securing payment for the validator if the builder fails to deliver the promised block or payment.

In examining 12,500 slots, we verified the positive influence of optimistic relaying on latency reduction. Out of these slots, 9,250

Table 1: Summary of Relay Performance Metrics for the Slot Range 6 087 501-6 100 000.

Relay	Blocks	Avg. Bids	Med. Time (ms)	Med. Value (ETH)
Ultra Sound	3,539	799	159.0	0.0389
Flashbots	3,202	549	-450.5	0.0348
Agnostic	2,509	641	107.0	0.0386

were relayed using at least one of the three relays we studied[5]. Table 1 displays the distribution of block deliveries across the relays. Notably, Ultra Sound Relay, known for its adoption of optimistic relaying, garnered the most bids per slot and delivered the highest quantity of blocks with the largest median value. Additionally, Ultra Sound reported the latest median winning bid arrival time. A more detailed distribution of winning bid timings can be found in Figure 9 in Appendix A.

These results suggest that by diminishing the block simulation latency, optimistic relaying enables relays to consider more bids and encourages builders to dispatch blocks later in the slot duration. The ultimate consequence is an increasing trend of block auctions being clinched by higher-value, late-submitted bids, thereby augmenting the rewards for validators. While no direct evidence of validators partaking in waiting games has been identified [32], our analysis of 12,500 slots, along with the extensive historical data provided in [36], confirms that validators strategically act to maximize their profits from MEV-Boost by registering with relays which deliver the most value. Currently, Ultra Sound Relay leads the pack, delivering around 30% of all blocks relayed through MEV-Boost, and offering the highest median value of 0.06 ETH to the validators [36].

4.2 The Risks of Waiting

4.2.1 Attestation Shares. Expanding on Schwarz-Schilling et al.'s [32] attestation share analysis, we examine the interplay of a block's winning bid arrival time, its attestation share, and potential fork vulnerability due to the proposer-boost mechanism [2]. Recalling that LMD-GHOST computes a block's weight as the cumulative effective balances of validators who attested to the block in a prior slot (remember that an attestation for a block at slot n is only included starting from slot $n + 1$), plus the weight coming from the parent block, we assess the effect of the proposer-boost mechanism. We investigate how this mechanism, awarding a timely block (released within a slot's initial 4 seconds) with a 40% committee weight, impacts potential past block re-organizations (re-org). This weight, derived from the slot's total assigned validator balance, allows for re-org vulnerability assessment tied to the winning block's arrival time. We compute each block's subsequent slot attestation share by calculating its weight (assuming a uniform 32 ETH validator effective balance) and normalizing against the slot's committee weight, resulting in a block's accrued attestation share. Consequently, a block with a below 40% subsequent slot attestation share is considered re-org susceptible via the proposer-boost.

Our results highlight that, on average, blocks gather 98% of attestation shares in the following slot. Out of the 8,121 distinct blocks proposed by the relays under our analysis, only 27 blocks acquired under 40% shares, with a median arrival time for winning bids of 1223ms. For the rest of the blocks, the median time was

[5]A builder may have submitted the same winning bid to multiple relays. In such instances, we categorize all relays featuring that bid as relayers of the winning block.

-22ms. Remarkably, we observed 75 blocks that arrived later than the 1223ms mark yet accrued sufficient shares to be unaffected by the proposer-boost, with an average share of 90%. This analysis is visually represented in Figure 3, where each blue dot stands for a block, and those within the red region identify blocks vulnerable to a re-org using the proposer-boost, based on their attestation share in the subsequent slot.

We further examined the relationship between the arrival times of winning bids in consecutive slots and the accrued attestation shares as the lateness of a block is intrinsically tied to the timeliness of its subsequent block. Figure 4 reveals a tendency for attestation shares to decrease when a block's winning bid is delayed while the succeeding slot's winning bid arrives in a timely manner. Notably, we identified 23 blocks vulnerable to the proposer-boost due to their below 40% attestation shares and a following slot's timely bid. Nevertheless, these blocks remained in the canonical chain, implying that the potential proposer-boost re-org was unexploited.

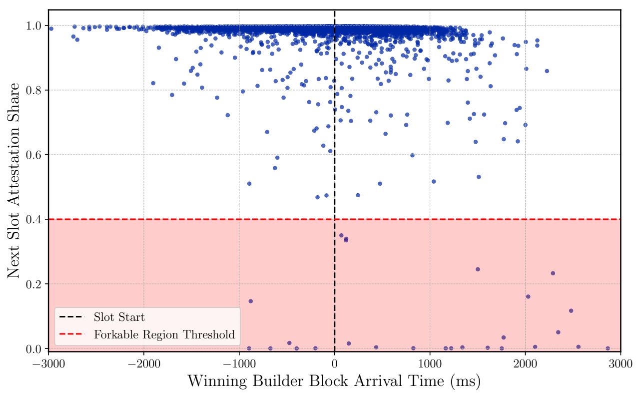

Figure 3: The figure plots blocks' winning bid arrival times versus their subsequent slot attestation shares. Blue dots represent unique blocks, and dots in the red zone indicate blocks with less than 40% attestation share, implying a forking risk.

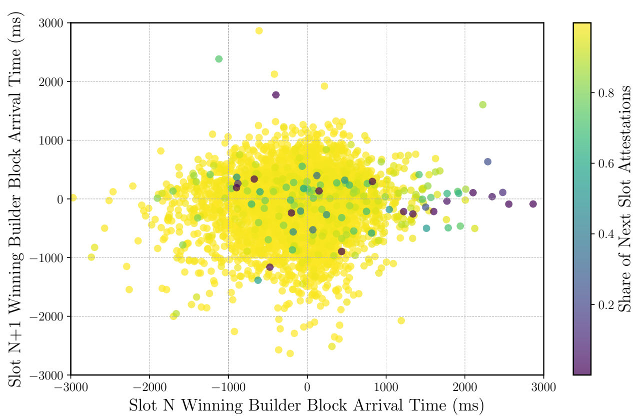

Figure 4: The figure maps consecutive slots' winning bid times with attestation shares. Dots denote blocks, with coordinates reflecting current (x-axis) and next slot (y-axis) bid times. The color gradient indicates the attestation share in the subsequent slot.

4.2.2 Orphaned Blocks. In our final empirical data analysis, we delved into the occurrence of orphaned blocks within our slot range and its correlation with the timing of the winning bid arrivals. We discovered 151 slots, amounting to roughly 1.2% of all slots we have analyzed, where a block failed to be included in the canonical chain. Among these, 123 slots had missed proposals, while 28 slots featured orphaned blocks, half of which stemmed from MEV-Boost and the other half were locally built by validators.

When we compared the timings of the winning bids of orphaned and canonical blocks in our data set, Figure 5 reveals that the earliest winning bid for an orphaned block was submitted 271ms prior to its slot's beginning, while the median time was 1115ms. Meanwhile, in the finalized chain, we observed 5,631 blocks with winning bids submitted later than the earliest orphaned block, with an overall 155ms median time. Interestingly, 130 of these canonical blocks had winning bids received even later than the median time of the orphaned blocks.

These observations lead us to two key insights: firstly, orphaning is a relatively rare event, occurring in approximately 0.22% of the 12,500 slots we examined. Secondly, orphaned blocks are not necessarily late arrivals, as some canonical blocks with later-arriving winning bids avoided orphaning. Thus, we conclude that delaying the selection of the winning bid, and playing the waiting games, does not necessarily result in a block's orphaning.

4.2.3 The Effects on Consensus: Agent-based Simulation Results. Until here, we based our observations on the empirical data. In this section, we present an agent-based model to study a scenario where a consistent share of the validators follows the same delay strategy. Specifically, whenever a delaying validator is selected as a block proposer, they wait until a certain uniform time into the slot to release their block. This strategy aims to maximize their MEV, as our prior sections suggest a positive correlation between MEV rewards and the interval between the release of successive blocks.

Our simulation here is not focused on estimating the actual profit increase these delayer validators would accrue - our earlier analysis already shows this to be positive. Instead, we aim to understand the

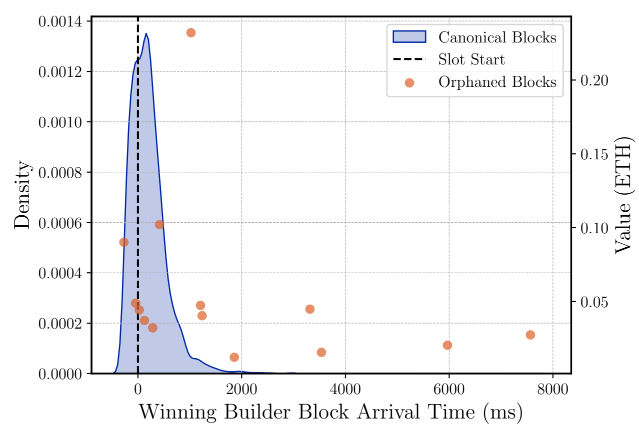

Figure 5: The density plot represents the distribution of canonical blocks that had a winning bid arriving after the earliest orphaned block. Individual orphaned blocks are indicated by scatter points, reflecting their unique arrival times and values.

broader impacts of such strategies on the overall consensus stability. We are interested in observing if such strategies may decrease the number of blocks included in the mainchain, which we assume to be an indicator of a blockchain's consensus efficiency.

The framework we refer to is essentially the same presented in [25]: the agents are validators, connected on a peer-to-peer network generated following the Erdos-Renyi random model [20]. This choice reflects the assumption of the peer-to-peer network to be maximally random in order to refrain from introducing any additional bias; further exploration of the actual structure of peer-to-peer topologies would certainly improve the model power. Time is assumed to be continuous and divided into slots; validators are randomly selected to be block proposers for each slot, and because they are assumed to be fully honest, they release the block exactly at the beginning of the assigned slot. The remaining validators are selected as attestors: they release an attestation to certify on the blockchain that they received the block from the block proposer and that it is valid. If they do not receive the block before the 4 seconds threshold, they are allowed to attest for the previous head of the chain. The only two random events that may happen are block gossiping and attestation gossiping, which follow exponential waiting times of parameters τ_{block} and $\tau_{attestation}$. Gossiping events happen when an agent is randomly picked to communicate the information about the blocks/attestations they received to one of their neighbor agents, picked at random as well. In [25], the authors showed how a phase transition in consensus efficiency is observed because of τ_{block}: when the value becomes larger than a certain threshold depending on the topological properties of the peer-to-peer network, the number of blocks included in the mainchain declines abruptly.

In the present work, we increase the number of parameters to take into account the delay strategy as part of the waiting games. We introduce $x^d \in [0, 1]$ the share of agents who are actively following the delay strategy, and $t^d \in [0, 12]$ which is the actual delay the delayers wait from the start of the slot before releasing the block. Our goal here is to introduce a toy model, an experimental setting to show the resiliency of Ethereum consensus: in this context, validators follow a very simplistic strategy where all delayers (a share of the total number of validators) follow the same delay time. This strategy implies on one hand that all delayers follow a common strategy, and on the other hand that neither delayers nor non-delayers do not adapt dynamically to other agents' strategies.

We proceed by generating a sample of simulations where we vary x^d and t^d, while we keep fixed $\tau_{block} = \tau_{attestation} = 3$ and the peer-to-peer topology, as well as the simulation time.

In order to estimate consensus efficiency we consider the mainchain Rate μ, from [25], formally defined as:

$$\mu = \frac{|M|}{|B|} \tag{1}$$

where $|M|$ is the number of blocks included in the canonical mainchain (M) while $|B|$ is the total number of blocks produced during the simulation.

Simulations results are plotted in Figure 6 and the results are intuitive: the effect of the time delay exercised by the delayers does not significantly affect consensus until it becomes larger than the

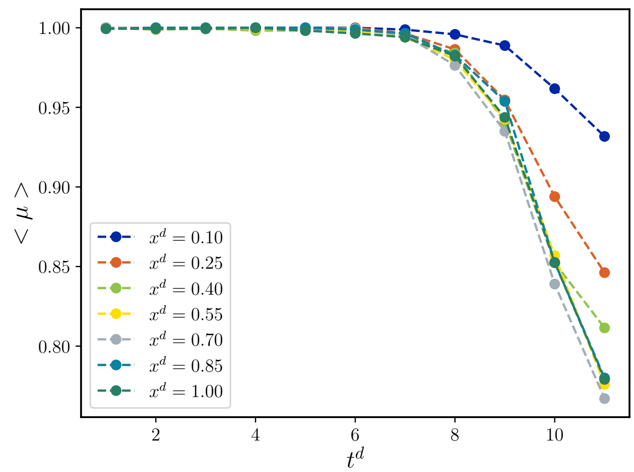

Figure 6: The Mainchain rate, eq. 1, averaged over 20 simulations running for 1000 seconds on an ER[20] graph of $N = 128$ and $< d >= 8$. On the x-axis t^d is the time delayers wait for release, while the curves colours represent the share of delayers wrt the total number of validators.

slot time minus the latency time: $12 - 3 = 9$ (the latency also represents the average time between two consecutive gossip interactions between the same two nodes). We also observe that the effect of the share of delayer stops increasing around $x^d = 0.5$, at which the number of consecutive mixed blocks (delayer blocks followed by honest blocks) is maximum.

While more research is needed to interpret these results, we believe that the results support the hypothesis that a delayer strategy supported by enough validators can be profitable and does not lead to consensus degradation for a range of delay times way more extensive than the ones we observed in the previous sections, in line with the theoretical results on the equilibrium of the waiting games described in [32].

5 CONCLUSION

In this study, we have investigated the dynamics of strategic waiting games within Ethereum's PoS-based consensus protocol. Our empirical data analysis and simulation results indicate that the potential economic benefits of engaging in waiting games by delaying block proposals outweigh the associated risks of getting excluded from the canonical chain, given the current dynamics of Ethereum's PoS mechanism and MEV. Nevertheless, validators seem to be not capitalizing on these benefits, as our study of unrealized value reveals. While the reasons remain unclear, we suspect large staking entities may be reluctant to participate due to the fear of reputation damage that could arise from such seemingly self-serving activities, potentially undermining consensus stability. Moving forward, we aim to optimize waiting games for validators and promote a more competitive block auction environment for builders by studying the latency in the interactions of these actors of the block production pipeline.

REFERENCES

[1] 2023. CoW Protocol. https://cow.fi Accessed: June 14, 2023.
[2] 2023. Ethereum Proof-of-Stake Consensus Specifications. https://github.com/ethereum/consensus-specs Accessed: May 20, 2023.

[3] 2023. flashbots/mev-boost. https://github.com/flashbots/mev-boost Accessed: June 13, 2023.

[4] 2023. MEV Blocker - Everyday protection from harmful MEV. https://mevblocker.io/ Accessed: June 14, 2023.

[5] 2023. MEV-Boost Block Proposal | Flashbots Docs. https://docs.flashbots.net/flashbots-mev-boost/architecture-overview/block-proposal Accessed: June 14, 2023.

[6] 2023. MEV-Boost Relay. https://github.com/flashbots/mev-boost-relay Accessed: June 14, 2023.

[7] 2023. MEV Data | EigenPhi | Wisdom of DeFi. https://eigenphi.io/ Accessed: June 13, 2023.

[8] 2023. MEV Explore. https://explore.flashbots.net/ Accessed: June 13, 2023.

[9] 2023. MEV-Share | Flashbots Docs. https://docs.flashbots.net/flashbots-protect/rpc/mev-share Accessed: June 14, 2023.

[10] 2023. Overview | Flashbots Docs. https://docs.flashbots.net/flashbots-auction/overview Accessed: June 13, 2023.

[11] 2023. randao/randao. https://github.com/randao/randao Accessed: June 15, 2023.

[12] Leemon Baird, Atul Luykx, and Paul Madsen. 2020. Hedera Technical Insights: Fair Timestamping and Fair Ordering of Transactions. https://hedera.com/blog/fair-timestamping-and-fair-ordering-of-transactions Accessed: June 14, 2023.

[13] Vitalik Buterin. 2021. Proposer/block builder separation-friendly fee market designs - Economics. https://ethresear.ch/t/proposer-block-builder-separation-friendly-fee-market-designs/9725 Accessed: June 14, 2023.

[14] Vitalik Buterin and Virgil Griffith. 2019. Casper the Friendly Finality Gadget. arXiv:1710.09437 [cs.CR]

[15] Vitalik Buterin, Diego Hernandez, Thor Kamphefner, Khiem Pham, Zhi Qiao, Danny Ryan, Juhyeok Sin, Ying Wang, and Yan X. Zhang. 2020. Combining GHOST and Casper. *arXiv:2003.03052 [cs]* (May 2020). http://arxiv.org/abs/2003.03052 arXiv: 2003.03052.

[16] Miles Carlsten, Harry Kalodner, S. Matthew Weinberg, and Arvind Narayanan. 2016. On the Instability of Bitcoin Without the Block Reward. In *Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security* (Vienna, Austria) *(CCS '16)*. Association for Computing Machinery, New York, NY, USA, 154–167. https://doi.org/10.1145/2976749.2978408

[17] Ankit Chiplunkar. 2023. Optimistic relays and where to find them. https://frontier.tech/optimistic-relays-and-where-to-find-them Accessed: June 15, 2023.

[18] Mikel Cortes-Goicoechea, Tarun Mohandas-Daryanani, Jose Luis Muñoz-Tapia, and Leonardo Bautista-Gomez. 2023. Autopsy of Ethereum's Post-Merge Reward System. arXiv:2303.09850 [cs.CR]

[19] Philip Daian, Steven Goldfeder, Tyler Kell, Yunqi Li, Xueyuan Zhao, Iddo Bentov, Lorenz Breidenbach, and Ari Juels. 2020. Flash Boys 2.0: Frontrunning in Decentralized Exchanges, Miner Extractable Value, and Consensus Instability. In *2020 IEEE Symposium on Security and Privacy (SP)*. 910–927. https://doi.org/10.1109/SP40000.2020.00040

[20] Paul Erdős, Alfréd Rényi, et al. 1960. On the evolution of random graphs. *Publ. Math. Inst. Hung. Acad. Sci* 5, 1 (1960), 17–60.

[21] Shayan Eskandari, Seyedehmahsa Moosavi, and Jeremy Clark. 2020. SoK: Transparent Dishonesty: Front-Running Attacks on Blockchain *(Financial Cryptography and Data Security)*, Andrea Bracciali, Jeremy Clark, Federico Pintore, Peter B. Rønne, and Massimiliano Sala (Eds.). Springer International Publishing, Cham, 170–189.

[22] Flashbots. 2022. The Future of MEV is SUAVE | Flashbots. https://writings.flashbots.net/the-future-of-mev-is-suave/ Accessed: June 13, 2023.

[23] Chris Hager and Frieder Paape. 2022. Running Geth within SGX: Our Experience, Learnings and Code | Flashbots. https://writings.flashbots.net/geth-inside-sgx Accessed: June 14, 2023.

[24] Mahimna Kelkar, Fan Zhang, Steven Goldfeder, and Ari Juels. 2020. Order-Fairness for Byzantine Consensus. Cryptology ePrint Archive, Paper 2020/269. https://eprint.iacr.org/2020/269 https://eprint.iacr.org/2020/269.

[25] Benjamin Kraner, Nicolò Vallarano, Claudio J. Tessone, and Caspar Schwarz-Schilling. 2023. Agent-Based Modelling of Ethereum Consensus. arXiv:2305.13906 [cs.GT]

[26] Barnabé Monnot. 2022. Notes on Proposer-Builder Separation (PBS). https://barnabe.substack.com/p/pbs?utm_medium=reader2 Accessed: June 13, 2023.

[27] Mike Neuder. 2023. michaelneuder/optimistic-relay-documentation. https://github.com/michaelneuder/optimistic-relay-documentation Accessed: June 15, 2023.

[28] Kaihua Qin, Liyi Zhou, and Arthur Gervais. 2022. Quantifying Blockchain Extractable Value: How dark is the forest?. In *2022 IEEE Symposium on Security and Privacy (SP)*. 198–214. https://doi.org/10.1109/SP46214.2022.9833734

[29] Kaihua Qin, Liyi Zhou, Benjamin Livshits, and Arthur Gervais. 2021. Attacking the DeFi Ecosystem with Flash Loans for Fun and Profit. In *Financial Cryptography and Data Security*, Nikita Borisov and Claudia Diaz (Eds.). Springer Berlin Heidelberg, Berlin, Heidelberg, 3–32.

[30] Caspar Schwarz-Schilling, Sheng-Nan Li, and Claudio J Tessone. 2021. Agent-based modelling of strategic behavior in pow protocols. In *2021 third international conference on blockchain computing and applications (BCCA)*. IEEE, 111–118.

[31] Caspar Schwarz-Schilling, Sheng-Nan Li, and Claudio J Tessone. 2022. Stochastic modelling of selfish mining in proof-of-work protocols. *Journal of Cybersecurity and Privacy* 2, 2 (2022), 292–310.

[32] Caspar Schwarz-Schilling, Fahad Saleh, Thomas Thiery, Jennifer Pan, Nihar Shah, and Barnabé Monnot. 2023. Time is Money: Strategic Timing Games in Proof-of-Stake Protocols. arXiv:2305.09032 [cs.GT]

[33] Yonatan Sompolinsky and Aviv Zohar. 2015. Secure High-Rate Transaction Processing in Bitcoin. In *Financial Cryptography and Data Security*, Rainer Böhme and Tatsuaki Okamoto (Eds.). Springer Berlin Heidelberg, Berlin, Heidelberg, 507–527.

[34] thegostep. 2022. Introducing the MEV Supply Chain. https://flashbots.mirror.xyz/bqCakwfQZkMsq63b50vib-nibo5eKai0QuK7m-Dsxpo Accessed: June 14, 2023.

[35] Christof Ferreira Torres, Ramiro Camino, and Radu State. 2021. Frontrunner Jones and the Raiders of the Dark Forest: An Empirical Study of Frontrunning on the Ethereum Blockchain. In *30th USENIX Security Symposium (USENIX Security 21)*. USENIX Association, 1343–1359. https://www.usenix.org/conference/usenixsecurity21/presentation/torres

[36] Anton Wahrstätter. 2023. MEV-Boost. https://mevboost.pics Accessed: June 14, 2023.

[37] Anton Wahrstätter, Liyi Zhou, Kaihua Qin, Davor Svetinovic, and Arthur Gervais. 2023. Time to Bribe: Measuring Block Construction Market. https://doi.org/10.48550/arXiv.2305.16468 arXiv:2305.16468 [cs].

[38] Ben Weintraub, Christof Ferreira Torres, Cristina Nita-Rotaru, and Radu State. 2022. A Flash(Bot) in the Pan: Measuring Maximal Extractable Value in Private Pools. In *Proceedings of the 22nd ACM Internet Measurement Conference* (Nice, France) *(IMC '22)*. Association for Computing Machinery, New York, NY, USA, 458–471. https://doi.org/10.1145/3517745.3561448

[39] Sen Yang, Fan Zhang, Ken Huang, Xi Chen, Youwei Yang, and Feng Zhu. 2022. SoK: MEV Countermeasures: Theory and Practice. arXiv:2212.05111 [cs.CR]

[40] Haoqian Zhang, Louis-Henri Merino, Vero Estrada-Galiñanes, and Bryan Ford. 2022. F3B: A Low-Latency Commit-and-Reveal Architecture to Mitigate Blockchain Front-Running. *ArXiv* abs/2205.08529 (2022).

[41] Liyi Zhou, Kaihua Qin, Christof Ferreira Torres, Duc Viet Le, and Arthur Gervais. 2021. High-Frequency Trading on Decentralized On-Chain Exchanges. *2021 IEEE Symposium on Security and Privacy (SP)* (2021), 428–445.

[42] Burak Öz and Felix Hoops. 2023. The Advent of a New Incentive - Maximal Extractable Value. In *A Companion to Decentralized Finance, Digital Assets, and Blockchain Technologies*, Pely Cronqvist (Ed.). Edward Elgar Publishing.

A ADDITIONAL FINDINGS

A.1 Progression of Builder Bids

Figure 7 shows the distribution of builder bids across the duration of slot 6 093 815, based on the number of transactions they contain. The results indicate an increase in the number of transactions included in builder bids over time which can be attributed to an influx of transactions in the public mempool and the builders' exclusive order flow. This trend further substantiates the positive marginal value of time observed in Figure 1, as the potential value of bids increases with the accumulation of transactions available for block building.

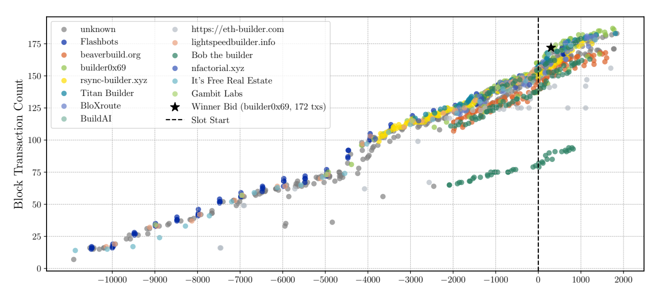

Figure 7: Distribution of unique builder blocks for slot 6 093 815 based on their arrival time to the relay and transaction count, across Ultra Sound, Flashbots, and Agnostic relays. Each color denotes a unique builder.

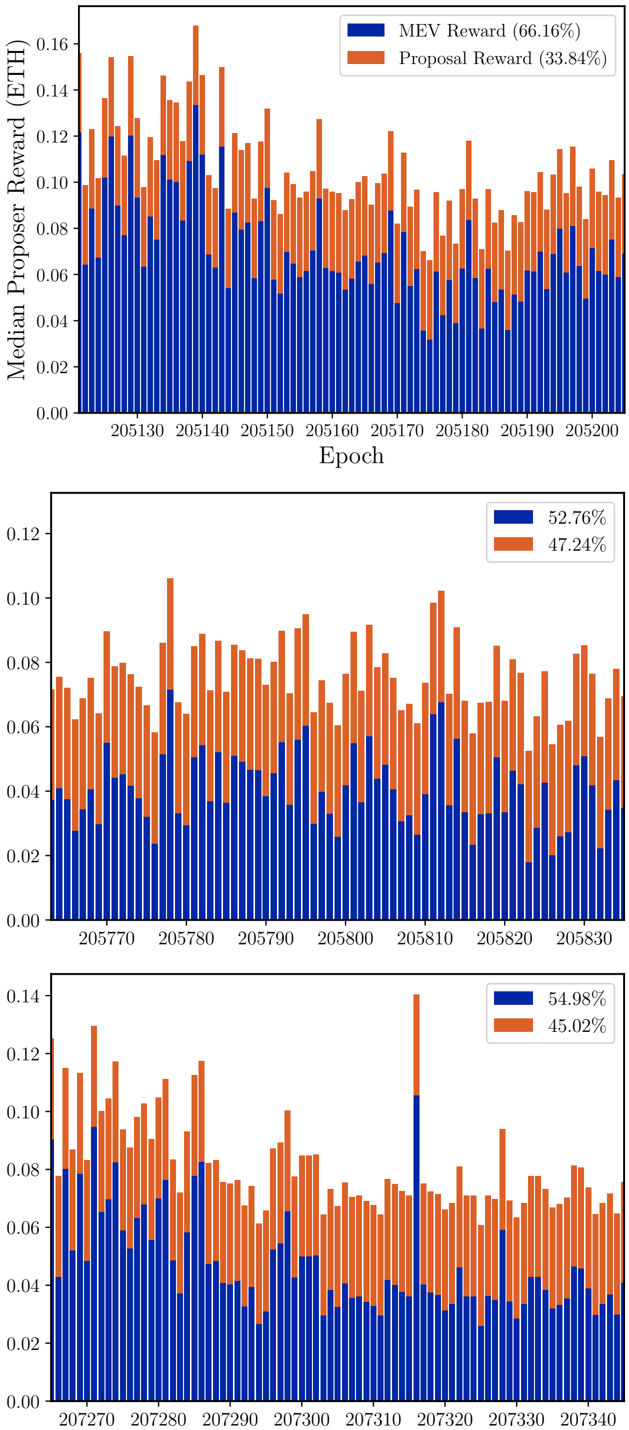

Figure 8: Distribution of median MEV and proposal rewards per epoch across different ranges. Each stacked bar represents an epoch, split into MEV rewards (blue - bottom) and proposal rewards (orange - top).

A.2 Reward Comparison

In Figure 8, we analyze three distinct epoch ranges, each containing around 80 epochs on average. The stack bar diagram distinguishes between two forms of rewards – the blue (bottom) and orange (top) segments. The blue segment represent the median MEV reward that validators, who proposed a block during that epoch via one of our observed relays, received. In contrast, the orange segment represents the median proposal reward issued by the consensus layer to validators. Our findings reveal that MEV rewards consistently surpassed the rewards from the consensus layer. While proposal rewards stayed steady at 0.034 ETH, the median MEV rewards accumulated to an overall median of 0.048 ETH, accounting for 58.32% of all rewards across the analyzed slots.

A.3 Winner Bid Arrival Times by Relays

Figure 9 presents our analysis of winning bid arrival times across different relays, underlining the variations between them. We found that Ultra Sound, recognized for its optimistic nature, and Flashbots, known to be non-optimistic, have distinct median winner arrival times. Interestingly, despite no known claims of being optimistic, Agnostic relay's median winner arrival time is markedly closer to Ultra Sound's than to Flashbots', suggesting potential similarities in their handling of block submissions.

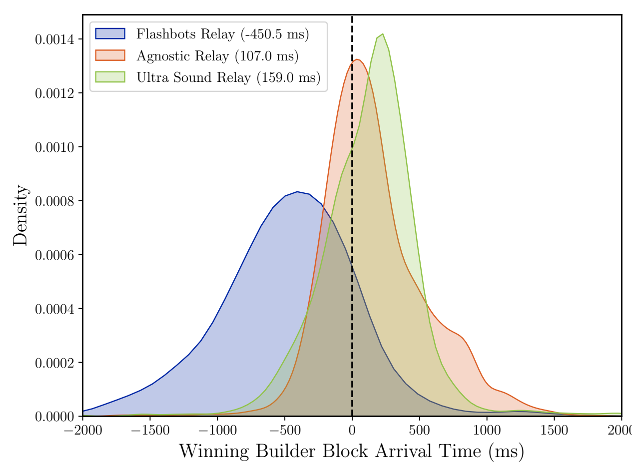

Figure 9: Density plots representing the arrival time distribution of winning bids across Ultra Sound (green), Flashbots (blue), and Agnostic (red) relays. Each plot corresponds to a specific relay, providing a comparative visualization of median arrival times for winning bids.