Eidgenössische Technische Hochschule Zürich
Swiss Federal Institute of Technology Zurich

Distributed Computing

From Searchers to Proposers: User Behaviour in Ethereum's Proposer-Builder Separation

Bachelor's Thesis

Antonia Unger

aunger@ethz.ch

Distributed Computing Group
Computer Engineering and Networks Laboratory
ETH Zürich

Supervisors:
Ye Wang, Robin Fritsch
Prof. Dr. Roger Wattenhofer

November 15, 2023

Acknowledgements

I would like to express my appreciation to everyone who has offered their support and contributed to the successful completion of my thesis. Special thanks go to Ye Wang and Robin Fritsch for their consistent guidance and valuable insights throughout this journey. Their expertise has been instrumental in shaping my research and the final outcome of the thesis. I'm also grateful to Prof. Dr. Wattenhofer and the ETH DISCO group for giving me the opportunity to work on this interesting project. Lastly, I would like to express my deepest gratitude to my parents for their unwavering support in all my endeavours.

Abstract

In September 2022, as the Ethereum blockchain shifted from a proof-of-work to a proof-of-stake consensus mechanism, the block generation process has undergone a remarkable transition with the widespread adoption of the Proposer-Builder Separation (PBS) mechanism. This separation allocates the roles of block building and block validation to distinct stakeholders (builders and proposers). Additionally, entities called relays act as intermediaries and connect builders to proposers. Finally, another key stakeholder, the searchers, try to capitalize on profit opportunities by dispatching transaction bundles to builders. This establishes a 3 tier layered structure (searchers, builders, and proposers), where each tier may optimize for maximum profit. In this study, we delve into the dynamics and interplay amongst the PBS market stakeholders and asses the current status of the PBS mechanism compared to its original objectives. Our findings reveal that both the relay and builder segment are characterized by an oligopolistic structure and are dominated by a handful of players. Further, we show that proposers consistently hold the largest profit share, followed by searchers and then builders. In contrast, a more granular analysis of individual players' weekly profits suggests a potential disproportionate favoring towards builders as opposed to individual proposers and searchers. Additionally, we conceptualize the bidding auctions where builders compete against each other for their block to be chosen for chain inclusion. On one hand, our research suggests high potential for improvement of the builders' bidding strategies and, on the other hand, reveals that notable percentage of builders use a reactive bidding strategy as opposed to bidding solely based on a block's value.

Contents

CHAPTER 1

Introduction

Over the last decade, the Ethereum blockchain has undergone a remarkable transformation continuously working towards greater decentralization. The process of generating blocks, from individual transactions of users to attested blocks in the chain, has always opened opportunities for different players to participate. In 2019, the concept of Maximum Extractable Value (MEV) emerged [1], which refers to the potential profits miners can extract by including, excluding, or reordering transactions. This has opened the playing field for many miners to compete against each other for MEV opportunities. However, malicious MEV extraction such as sandwich attacks can have detrimental effects on users [1] and to mitigate such negative implications, the block creation mechanism has fundamentally changed in September 2022. As Ethereum shifted from a proof-of-work (PoW) to a proof-of-stake consensus mechanism (PoS), the Proposer-Builder Separation mechanism has found widespread adoption. PBS distinguishes the roles of block producers and block proposers, whereas in PoW both roles were held by miners. While validators remain responsible for proposing and voting on blocks, a new entity called *block builders* have been introduced to order transactions and assemble blocks. Block builders and block proposers both connect to entities called *relays*, which create a channel between the two. From all the blocks that are forwarded from the relays, the proposer picks the most lucrative one. Finally, there exists another key player in the ecosystem - the *searchers*. Their role is to identify MEV opportunities, and capitalize on them by creating transaction bundles which are handed to the block builders. During this entire process, on one end, searchers bribe builders to include their transaction bundles in the block and, on the other end, builders bribe proposers to select their block. This structure gives rise to intriguing dynamics among the various stakeholders. In this study, we aim to shed light on how different players in this newly emerged PBS market interact with each other and evaluate how the current status of the PBS mechanism aligns with its original goals.

The following are the main contributions of this study:

- A longitudinal time analysis of the market dynamics of PBS. This includes an overview of the wide spread adoption of PBS, as well the different players (builders, relays, proposers) and how they share the market.

- A profitability analysis comparing the profit shares between the searchers, builders and proposers. While past papers have focused on the profits of individual or 2 of the players (e.g. [2], [3]), this study provides a generalized approach by illustrating the entire flow of tokens from searchers to builders to proposers.

- An analysis of the auctions that arise when different builders compete for their block to be chosen by proposers. This includes the analysis of newly collected auction data of over 350'000 auctions which has not yet be thoroughly explored, as well as a deeper dive into two different bidding strategies: fraction-based bidding based on block value and reactive counterbidding.

Chapter 2

Background

In this section, we introduce background information and concepts that are relevant to understand the analysis conducted in this thesis.

2.1 Ethereum 2.0 Architecture

In September 2022 the Ethereum network, the most popular decentralized blockchain, underwent a significant transition to a proof of stake (PoS) consensus mechanism. This transition known as "The Merge" involved connecting the original execution layer with its new proof-of-stake consensus layer, the Beacon Chain [4].

2.1.1 Execution Client, Consensus Client and Validators

Figure 2.1 illustrates the relationship between the execution and consensus clients in PoS Ethereum. Both clients connect to their respective peer-to-peer (P2P)

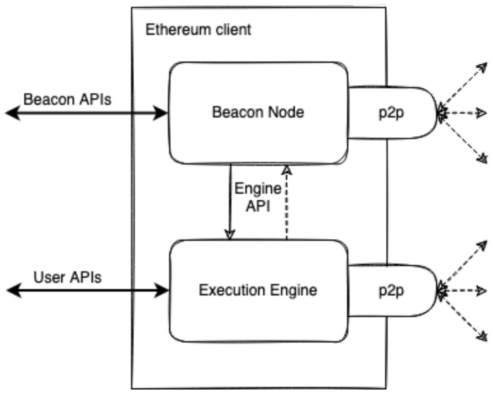

Figure 2.1: Ethereum 2.0 Client Architecture [5]

networks. The Execution Client manages transaction handling, transaction gossip, state management and supporting the Ethereum Virtual Machine (EVM). It serves as a gateway for users to interact with the blockchain, enabling them to submit queries and transactions, and deploy smart contracts. In contrast, the consensus client is responsible for the synchronization logic necessary to keep nodes up to date with the Ethereum network. This involves receiving blocks from peers and running a fork choice algorithm to ensure the node consistently follows the chain with the highest accumulation of attestations. The consensus client itself does not participate in attesting or proposing blocks. These actions are carried out by validators, which are an optional extension of the consensus client. Node operators can add a validator by depositing 32 ETH. Running the validator software also makes a node eligible to be selected to propose a new block.

In the Ethereum network, a block can be added to the blockchain in each slot, where every slot lasts 12 seconds. During each slot, a single validator is pseudo-randomly selected to propose a block (using the RANDAO algorithm) [6].

2.2 MEV, PBS and MEV-Boost

2.2.1 Maximal Extractable Value (MEV)

Maximal extractable value (MEV) has gained considerable attention since it was introduced by Daian et al. [1]. It refers to the maximum profit that can be attained through the manipulation of transaction ordering, in addition to the standard block reward and gas fees. By including, excluding, and changing the order of transactions, users gain the ability to extract MEV. This extraction can be accomplished through various methods, with the most common types being arbitrage, front-running and sandwich-attacks.

In PoS Ethereum, MEV in theory accrues entirely to validators since they are the only parties capable of guaranteeing the execution of a profitable MEV opportunity. However, in practice, a substantial portion of MEV is extracted by independent network participants known as *searchers* [6].

There already exists extensive literature on the security implications of MEV bidding strategies: Malicious MEV extraction can cause users to directly loose money [7], can lead to on-chain bidding wars resulting in network congestion and high gas fees [1], and has been argued to be a centralizing force because finding the most profitable MEV opportunities requires significant resources which only big players can afford [8]. In order to mitigate such negative implications, several MEV countermeasures may be considered [9], including in-protocol Proposer-Builder Separation.

2.2.2 Proposer-Builder Separation (PBS)

Proposer-Builder Separation (PBS) is a proposed solution aimed at reducing the negative impact of MEV on Ethereum's decentralization and security. PBS introduces a significant change to the consensus protocol, primarily by separating the block producer and block proposer roles. While validators remain responsible for proposing and voting on blocks, a new class of specialized entities called "block builders" is introduced to order transactions and construct blocks. This separation of roles has the goal of promoting greater competition and decentralization, as any user can become a block builder and extract MEV. Although PBS is not yet implemented in-protocol, there are ongoing efforts to introduce in-protocol support in the future [6].

2.2.3 MEV-Boost

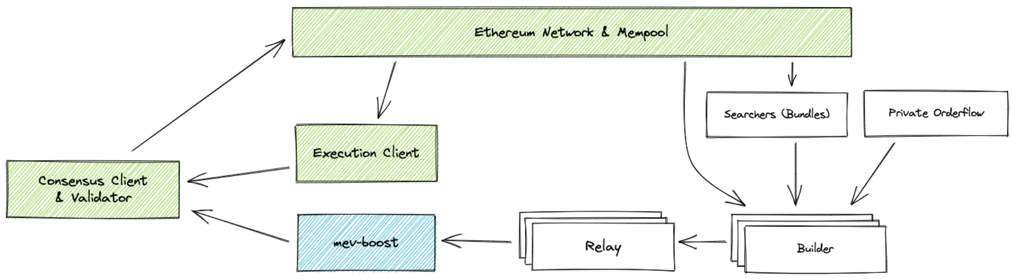

Figure 2.2: MEV-Boost Architecture [10]

MEV-Boost serves as an intermediate realization of PBS - a separate piece of open source software, which outsources block-building to a network of builders. As illustrated in Figure 2.2, MEV-Boost operates as follows [11]:

- **Builders** assemble full blocks by collecting transactions from users/searchers, the public mempool, and by inserting own transactions to extract MEV. They then submit their blocks to relays, containing promised payments to block proposers. Notable examples of block builders include Flashbots, Eden, BloXroute and Blocknative. All builders which were labelled in this analysis (and their respective addresses) are included in Tab. A.2 in the Appendix.
- **Relays** aggregate blocks from multiple builders in order to select the block with the highest fee and propagate the received blocks to listening proposers (validators). There are different types of relays which distinguish themselves with regard to how they connect to builders and what censorship they deploy. Relays may either connect to their own internal builders

(e.g. Blocknative and Flashbots) and/or connect to external builders in a permissioned (e.g. Eden and Bloxroute) or permissionless manner (e.g. Flashbots and Manifold). Furthermore, some relays censor certain transactions to be OFAC-compliant [12]. Table 2.1 provides an overview of the relays analyzed in this study and their characteristics.

- **Validators** (i.e. proposers in PoS) ultimately select the most profitable block and propose it to the Ethereum network for attestation and block inclusion.

- **Searchers** are another critical type of participant in the Ethereum block-building funnel who extract a large portion of MEV. They run complex algorithms (usually bots) to detect profitable MEV opportunities and submit a bundle of transactions through a private channel to a builder (such as Flashbots, Eden, Bloxroute, and many more). Additionally, searchers include a bidding amount to the builders, in order for the builder to include their transaction bundle in the block. All transactions in a bundle will be carried out atomically and as they are not revealed to the public mempool the searchers do not run the risk of becoming the victim of a front-running attack.

Relay Name	**Builder Connections**	**Censorship**
Flashbots	internal & permissionless	OFAC-compliant
bloXroute (E)	internal & permissioned	x
bloXroute (M)	internal & permissioned	x
bloXroute (R)	internal & permissioned	OFAC-compliant
Eden	internal & permissioned	OFAC-compliant
UltraSound	permissionless	x
Agnostic Gnosis	permissionless	x
Blocknative	internal	OFAC-compliant
Aestus	permissionless	x
Manifold	permissionless	x

Table 2.1: The table above summarizes the different types of relays and their main characteristics. The 2$^{\text{nd}}$ column outlines to what types of builders the relays connect to and the 3$^{\text{rd}}$ column indicates if they are are regulated under OFAC and censor certain transactions. We obtained this information directly from the relay websites A.1 as well as [3].

Chapter 3

Data Analysis

3.1 Market Dynamics

In the following we present a high-level overview of the PBS market dynamics and provide the reader with a first intuition on how the market is structured, what different players exist and how they interact with each other. Overall, the findings align with those presented in previous papers, for example [2] and [3]. Compared to previous papers, this analysis covers a longer time period, which provides a more holistic view of the current ecosystem and illustrates potential risks associated with network centralization.

3.1.1 Data Retrieval and Methodology

In order to conduct this analysis we require data on MEV blocks, in particular the different players which have participated in the block building process. For each MEV block, the MEV-Boost Relay Data API [13] allowed us to retrieve the payloads that were delivered from each relay to all proposers. This includes information on the block builder, relay and proposer which allowed for the analysis presented in the next section. Data was collected from September 15$^{\text{th}}$ 2022 until June 15$^{\text{th}}$ 2023.

3.1.2 Results

MEV-Boost has found rapid adoption since its introduction on September 15$^{\text{th}}$ 2022 and has been proposers' preferred choice of block-building: the share of blocks using the MEV-Boost software ('PBS blocks' or 'MEV-boosted blocks') reached 86% on November 3$^{\text{rd}}$ 2022, and has been relatively stable between 85-94% since then (Figure 3.1). On April 13$^{\text{th}}$ one can observe a drastic decrease in the share of MEV-boosted blocks (reaching under 66%), which was likely due to the Ethereum Shapella upgrade on April 12$^{\text{th}}$. Capella changes required proposers, relays and block builders to update there software stack which may have incurred delays and temporary unavailability of some relays and builders [14].

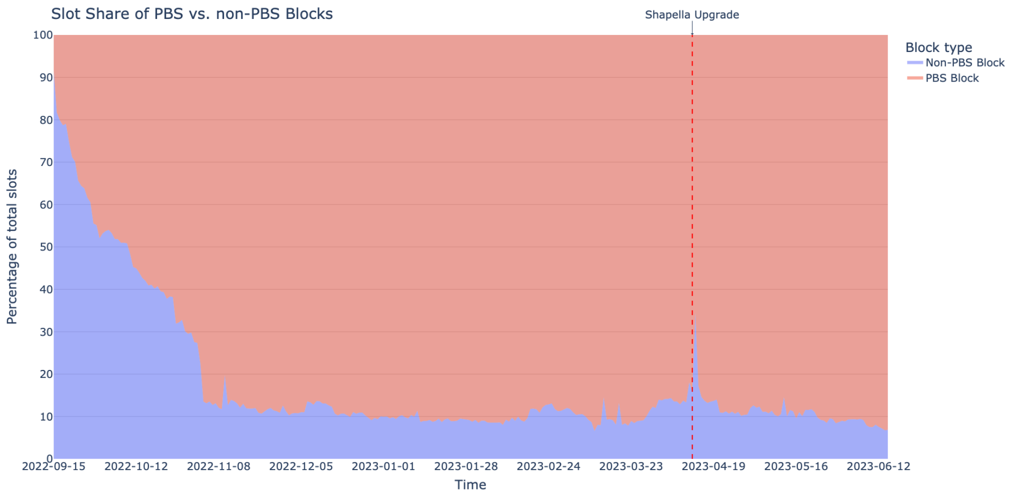

Figure 3.1: The above figure illustrates the proportion MEV-boosted blocks over time. 'MEV-boosted blocks' or 'PBS blocks' refer to blocks which have gone through the MEV-boost block creation funnel (builder to relay to proposer). In contrast, non-PBS blocks are blocks which have been built by the block proposer itself and were not submitted through a relay.

Figure 3.2 illustrates the shares of PBS blocks corresponding to the different relays. A total of 11 relays exist (10 are displayed in the graph) creating a strong dependence on very few players. This leads to a high network fragility, and an up-time below 100% can quickly lead to missed slots, which may have caused of the dip in PBS blocks on April 13$^{\text{th}}$ during the Shapella upgrade. We observe, that the Flashbots relay was initially dominating the market (with over 80% market share). However, over time other relays gained higher market shares and decentralization has increased. Nevertheless the top 3 relays (Flashbots, Agnostic Gnosis and Ultrasound) still accrue for the majority of PBS blocks and the relay landscape continues to be highly centralized.

Figure 3.3 illustrates the share of PBS blocks built by different builders. Similar to the relay landscape, we can observe that Flashbots was initially the single dominant builder (building over 40% of PBS blocks), and other builders gained increasing market share over time. In contrast to the relay landscape, many small builders exist which are summarized in Fig. 3.3 in the "Other" category, however, they don't account for a large percentage of blocks (typically between 10-20%).

We conclude that the concentration of both the builder and relay market decreased with time, however, in particular the relay landscape remains highly

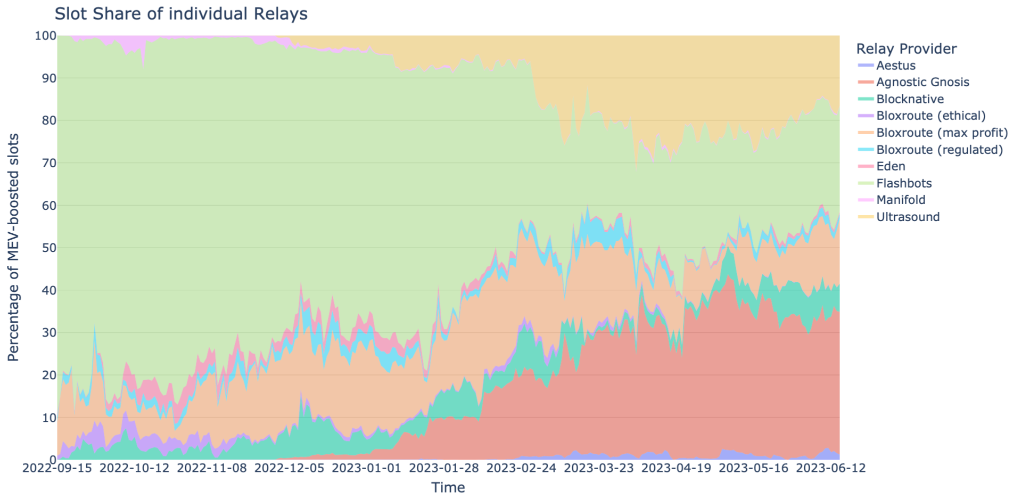

Figure 3.2: The above figure illustrates the proportion of mev-boosted blocks (y-axis) over time (x-axis) split into individual relays. Each color represents one relay, and we plot the daily share of blocks of each relay.

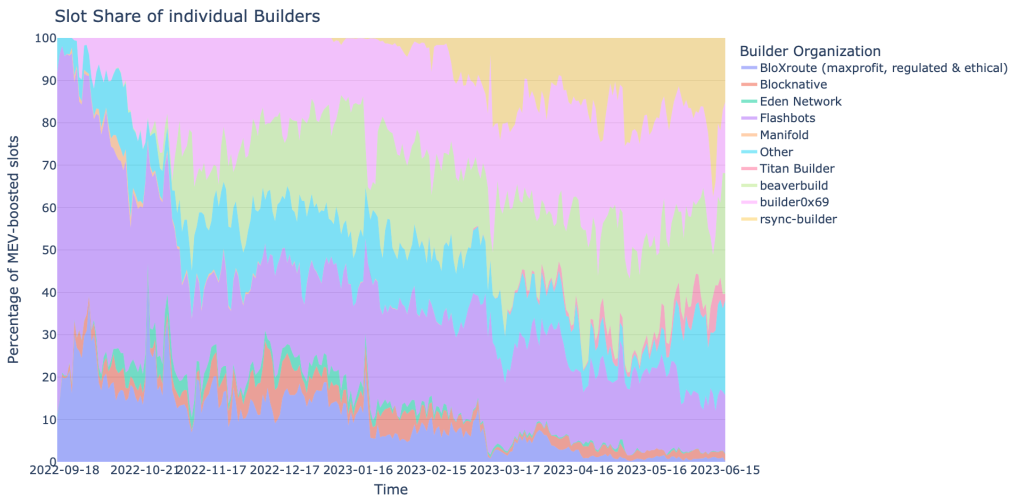

Figure 3.3: The above figure illustrates the proportion of mev-boosted blocks (y-axis) over time (x-axis) split into individual builders. Each color represents one builder, and we plot the daily share of blocks of each builder. Only the most well-known builders are labelled inidivdually, the remaining ones fall into the category "Other".

centralized. The builder market is slightly more competitive and decentralized as many small builders exist, but nevertheless, the market is strongly dominated by a few large players. Further analysis on the market concentration of relays and builders can be found in [3], where the Herfindahl-Hirschman Index is calculated as a measure of market concentration.

3.2 Distribution of Profits

In this section we present the results of a profitability analysis, comparing the profit shares between the searchers, builders and proposers. Past papers have so far either focused on comparing builder and proposer profits (e.g. [2], [3]), or comparing searcher profits to the bribes they send to builders. In this analysis we generalized the previous approaches by illustrating the entire flow of currencies from searchers to builders to proposers. This extension allows us to offer a more comprehensive understanding of how MEV is shared within the Ethereum ecosystem.

3.2.1 Data Retrieval and Methodology

To conduct this analysis we require three metrics: (1)*searcher profit*, (2) *builder profit* and (3) *proposer profit.*

For a transaction to be included in a block the user (and thus also searcher) pays a fee (*total fee*), which is the sum of the *base fee* and the *priority fee.* The *base fee* is paid per unit of gas and depends on the computational complexity of the transaction. It is dynamically set by the network according to demand and burned by the system. The *priority fee*, on the other hand, is set by the user and serves as a bribe for the transaction to be included in the block. Additionally, users can tip the block creator via direct transfers to the fee recipient address. Correspondingly, the searchers costs are also composed of the transaction costs (*base fee* and *priority fee*) and any direct transfers. The revenue of searchers is simply the revenue from the arbitrage and sandwich-attack transactions ("MEV revenue"). This gives us the following *searcher profit*:

$$\boldsymbol{Searcher\ Profit} = \text{MEV Revenue} - (\text{Transaction Costs} + \text{Direct Transfers})$$

Based on the the *priority fees* and *direct transfers* which users pay for their transactions to be included in a block, we define the total *block reward* for builders and validators as follows:

$$\boldsymbol{Block\ reward} = \text{Sum of } \textit{priority fees} + \text{Sum of } \textit{direct transfers}$$

For non-PBS blocks this block reward is accrues entirely to the validator. For PBS Blocks, which we focus on in this analysis, the block builder receives the block reward and pays the proposer a separate reward by including a direct transfer to the proposer as the last transaction in the block. This transaction can be easily traced and constitutes the ***Proposer Profit***. Consequently, the *builders profit* is defined as follows:

$$\boldsymbol{\mathit{Builder\ Profit}} = \mathit{Block\ Reward} - \mathit{Proposer\ Profit}$$

After having theoretically defined the searcher, builder and proposer profit, we now consider how they were computed in practice.

To compute the *proposer profit* we retrieved PBS block data from the MEV-Boost Relay Data API [13] which directly includes this information for each block. For the *builder profit*, we calculated the sum of priority fees of all transaction of a block and added any direct transfers to the block builder.

In order to compute the *searcher profit* one would need to trace all MEV-transactions from the public mempool and from private transaction pools (as searchers often submit their transaction bundles via private channels). Flashbots, who operate one of the largest private transaction pools, encourage transparency of the ecosystem and make their private transactions openly available. For the scope of this analysis, we focus solely on transactions within Flashbots private transaction pool. To ensure a fair comparison between the profit of searchers, builders and proposers, we compute the *builder profit* and *proposer profit* exclusively from blocks built by the 5 Flashbots builders which include private transactions (see list of builders in Appendix A.2). In future research, one could generalize this study to encompass all searchers.

To calculate the MEV revenue of searchers we trace each Flashbot private transaction and compute the profit of all arbitrage and sandwich attack transactions[1]. Since the MEV revenue is not necessarily in Ether, we use the Binance Klines API endpoint [15] to translate the revenue value into Ether at the time of the transaction. This allowed us to convert 88.5% of all transactions and the remaining 11.5% were disregarded.[2] Finally, we subtracted the transaction costs ($\mathit{gas\ fee} * \mathit{gas\ used}$) and any direct transfers to builders from the revenue giving us the *searcher profit*. Note, that other types of MEV transactions exist (e.g. liquidations and NFT MEV [16]) which where not considered in this study. As a result, the effective searchers profit may be higher than the value calculated in this study.

[1]For arbitrage transactions (/ sandwich attacks), we calculate the searcher's balance difference before the execution of the first arbitrage (/front-run) transaction and after the execution of the last arbitrage (/back-run) transaction.

[2]The Binance Klines API endpoint [15] does not contain all currency exchange rates, especially those of smaller tokens are often missing.

3.2.2 Results

In Figure 3.4 we illustrate the daily profit distribution among searchers, builders, and proposers, and in Figure 3.5 their weekly sum of profits. Proposers consistently hold the most substantial portion of profits throughout the studied period and typically captured 75-90% of profits. Searchers claim the next significant portion, with a profit share between 5-20%. Finally, the builder profit constitutes the smallest share rarely exceeding 10%.

Notably, the searcher and builder profits show a high volatility. We could not identify direct causes for their profit peaks and dips, and it seems it is mainly due to their dynamic nature in the PBS ecosystem. Many factors such as dynamically changing MEV opportunities, fluctuating transaction costs and bidding auctions between searchers may influence this volatility. In order to check whether searchers, builders or proposers are favored during moments of high volatility on Ethereum, we conducted correlation tests between ETH-USD price changes and their weekly profits.[3] Data points of ETH to USD price changes were downloaded from Etherscan [18] and we tested for correlations of this data against the profits of searchers, builders and proposers. However, no statistically significant correlations were measured (see Table B.1 in the Appendix). As we used rather coarse-grained data points of the ETH to USD price changes (daily), one may re-evaluate this correlation in future research using smaller time intervals (e.g. by calculating the volatility of each 12 second slot).

[3]This notion originated from a study by Gupta et al. which finds that certain builders are more likely to win in the PBS auctions when volatility is high [17].

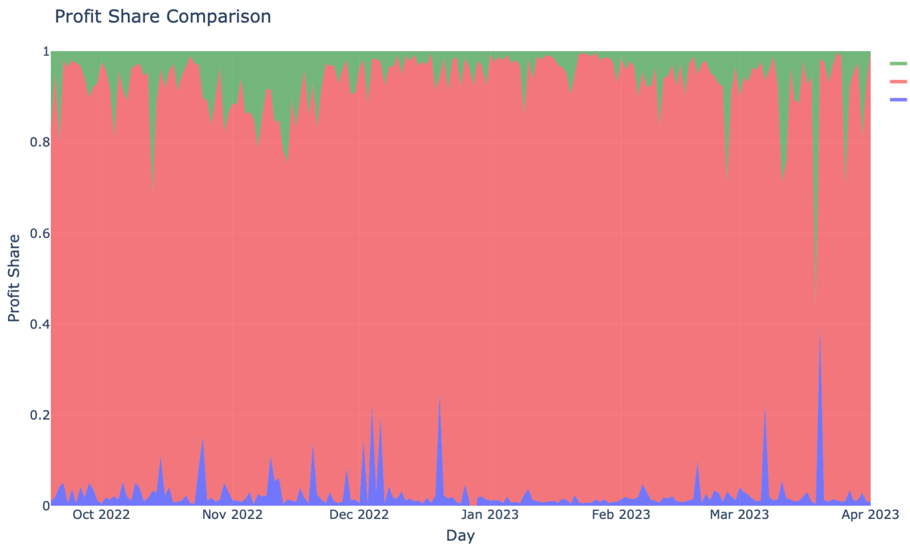

Figure 3.4: The figure above illustrate the profit shares of block builders, block proposers and searchers. The x-axis represents time and the y-axis profit shares of the entire profit. The dataset only includes blocks from Flashbots builders.

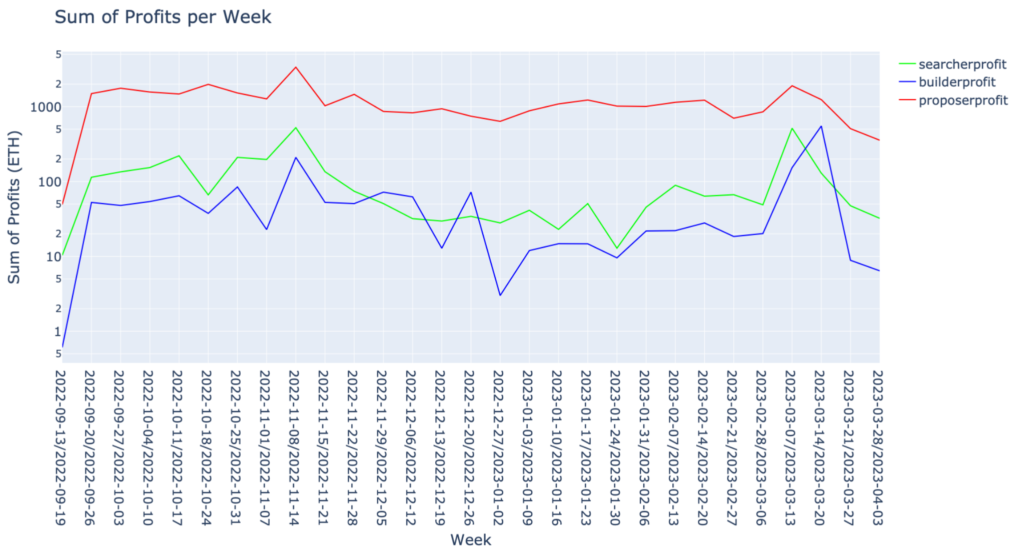

Figure 3.5: The figure above illustrates the weekly sum of profits of searchers, builders and proposers. The x-axis refers to the time (i.e. week) and the y-axis to the sum of profits in ETH.

At this point it seems like the proposers are the most profitable party, however, we must also consider the number of users this profit is split between and how it is distributed: The builder profit accrues entirely to the 5 Flashbot builders due to the selected dataset. In contrast, the profit earned by searchers is dispersed amongst approximately 50 to 100 distinct addresses per day, and for proposers the number varies between 100-200 addresses (driven by the pseudo-random selection algorithm)[4]. In table 3.1, we proceed to compare the per player profits of searchers, builders and proposers.

	Weekly profit	
Flashbot Builders	3 - 551 ETH	
	Mean weekly profit	**Median weekly profit**
Searchers	0.15 - 4.6 ETH	0.01 - 0.17 ETH
Proposers	0.9 - 3.2 ETH	0.09 - 0.21 ETH

Table 3.1: Summary of profit per player of builders, searchers and proposers. A more comprehensive viusaliation of the mentioned data point can be found in Fig. 3.5, Fig. 3.6 and Fig. 3.7

As already previously mentioned, the Flashbots builder profit per week fluctuates significantly which is why we can only indicate a broad range of between 3-551 ETH. For proposers and builders we have a more constant median weekly profit of all players of between 0.01-0.17 ETH for searchers and 0.09 - 0.21 ETH for proposers. In Figure 3.6 and 3.7 we illustrate the distribution of weekly profit per player of searchers and proposers. For searchers we observe that the third quantile point is typically between 0.1-0.8 ETH. However, there are some outliers which achieve a much higher profit and thus increase the mean value. This would indicate that there are a few searchers outsmarting the majority and as a result achieve a large proportion of the profit. A similar trend is present for the proposers, where the third quantile point is typically between 0.2-0.8 ETH, and again, the mean value is pulled up by a few players achieving a high weekly profit. However, this is due to staking organizations like Lido [19] or stakefish [20], which control a significant share of the validator market[5] and thus receive the largest profit share. Note, that these organizations themselves operate a network users who stake Ether and to whom this profit is passed on to, so we cannot jump to the same conclusion as for searchers.

In general, it is not possible to give a definite answer regarding which stakeholder is most profitable, as the variations between the individual proposers and searchers are significant. However, as the entire builder profit accrues to Flashbots (See Fig. 3.5), this analysis seems to indicate that builders achieve an over proportional share of the entire profit compared to individual searchers and pro-

[4]See Appendix B.1 for the precise values.

[5]As of September 2023, Lido controls over 30% of Ether stake [19]

posers. However, this study only provides an isolated view of Flashbots builders and accordingly, this notion would need to be validated for more builders. This may be a focus of future research.

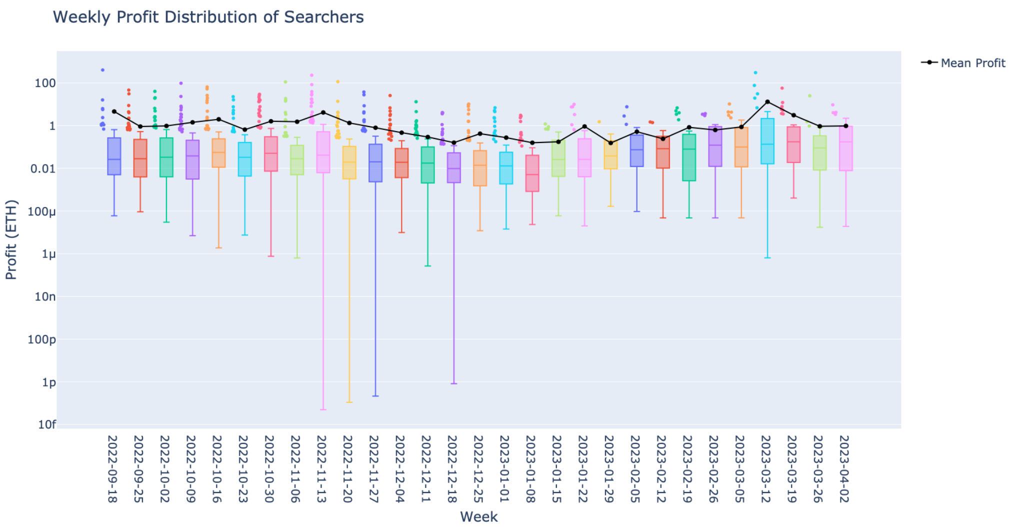

Figure 3.6: The above figure illustrates how the searcher profit is shared amongst individual searchers. The x-axis refers to the time (i.e. week) and the y-axis to the profit amount in ETH. Each box represents the interquartile range (IQR) of individual searchers profits and the line within a box the median. Individual dots mark outliers which are outside the whisker limit. Note that not all outliers are visible on this y-scale.

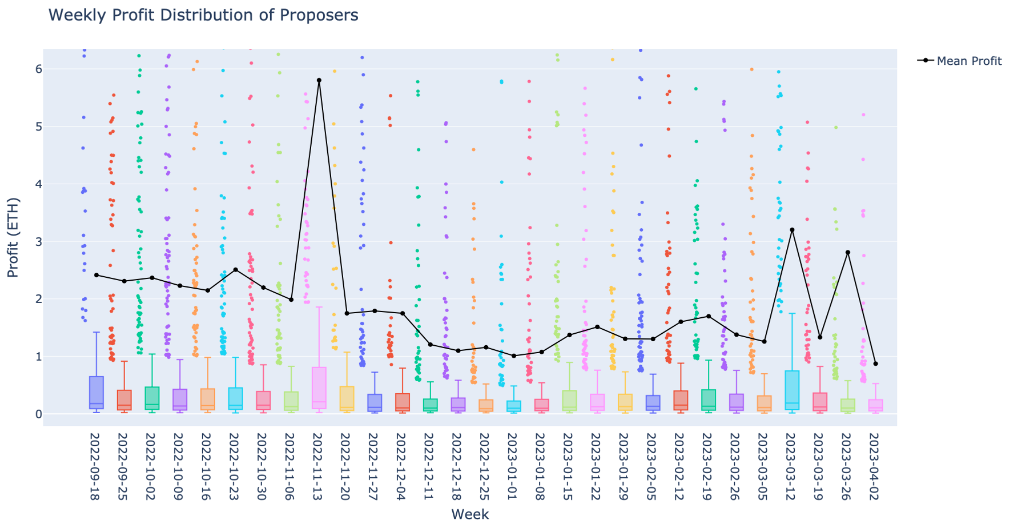

Figure 3.7: The above figure illustrates how the proposer profit is shared amongst individual proposer. The x-axis refers to the time (i.e. week) and the y-axis to the profit amount in ETH. Each box represents the interquartile range (IQR) of individual proposers profits and the line within a box the median. Individual dots mark outliers which are outside the whisker limit. Note that not all outliers are visible on this y-scale.

3.3 Builder Auctions

This section focuses on analyzing bidding behaviors of builders: After a builder orders transactions to form a block (from searchers directly or from the public mempool), they determine a *bidding value* to be sent with the block header to one or multiple relays. The relays connected to the current proposer, will forward the highest received bid together with the block header to the proposer, which chooses one of received proposals or alternatively assembles a block themselves.

This bidding process is fundamentally related to the goal of PBS to decentralize MEV profits as, in theory, anyone can become a block builder and thus compete for MEV opportunities.

In the following section we introduce the most important specifications and constraints of a builder auction. In the context of this notion, we present overall data characteristics, analyze several bidding heuristics and aim to shed light on different types of bidding strategies.

3.3.1 Data Retrieval and Methodology

To conduct this analysis we collected two datasets from the MEV-Boost Relay Data API [13]:

1. Firstly, we collected the auction winners for each MEV-boosted block from the *proposer payload delivered* endpoint. We enriched this dataset with the previously computed *builder profit* and the *proposer profit* for each block.[6]

2. Secondly, for each relay[7], we collected all bids that were received for each block auction from the *builder blocks received* endpoint. Each relay was queried individually and the specific endpoints are included in Table A.1 in the Appendix. During the querying process we faced several failure status codes (for example *Status 429* errors with the Flashbots endpoint and *Status 503* with the Blocknative endpoint), resulting is missing data points for some relays. In order to prevent false results from incomplete data, we decided to continue the analysis exclusively with blocks from the relays Eden, Agnostic Gnosis and Ultrasound which had the highest data completeness ($> 99\%$). In further research, on could generalize this analysis to data from all relays.

When combining both datasets we noticed that for some transactions (apprx. 5%) the *proposer profit* equals 0 although the bid value of the winning bid does not equal 0. During correct execution, these two values match and thus we

[6]Refer to section 3.2.1 for the definition derivation of builder and proposer profit.

[7]Flashbots, Blocknative, Bloxroute (ethical), Bloxroute (max profit), Bloxroute (regulated), Manifold, Eden, Ultrasound, Agnostic Gnosis, Aestus

removed these transactions. The final dataset consisted of 398'700 blocks from Sep-15-2022 - Jun-15-2023.

3.3.2 Auction Game Specification and Constraints

A builder auction can be viewed as a sequential game among a set of bidders (i.e. block builders) who bid against one another to obtain a share of the *block reward*. The game has the following properties:

a) **Open auction start time:** Players can decide freely when to start submitting bids, even before the start time of a slot. This behavior is commonly observed in practice and beneficial, as relays can immediately send suitable blocks to the proposer when the slot starts [2].

b) **Continuous bidding**: Players can submit bids continuously, rather than in discrete rounds (as the bidding process is asynchronous).

c) **Partially revealed information:** Players can see one another's bids by querying the MEV Relay API [13]. However, bidding information is subject to network latency, resulting in varying delays among players. This latency discrepancy can potentially confer a competitive advantage to some players.

d) **Rate-limited bidding:** Players have the freedom to raise and lower their bids by arbitrary increments, however, they are constrained by a submission rate limit. For the Flashbots relay, this rate limit is 2 submissions / second / IP address [21]. Although we couldn't find information on the other relays, we expect the rate limit to be of similar magnitude. Any new bid of a player will overwrite a previous one, even if it is less profitable [13].

e) **Variable block rewards between players**: Players have the potential to earn varying block rewards, based on the transactions they include in the block and the block building strategies they employ. For instance, builders who receive transactions through private channels alongside those from the public mempool may have an advantage over builders relying solely on public mempool transactions. Additionally, builders utilizing advanced algorithms to exploit MEV opportunities may also gain an edge in earning higher rewards.

f) **Evolving block reward of each player:** Block rewards for players can change and increase as the auction progresses. This is due to the players composing and submitting updated blocks, for example, based on new transactions that are submitted to the public mempool or sent to the player directly.

g) **Unknown auction duration**: The auction does not have a predetermined or fixed duration. Instead, the proposer decides how long to wait for new incoming bids and terminates the auction.[8] All bids submitted after termination will be disregarded even if they are higher than the selected bid.

As an example, Figure 3.8 illustrates the bidding action that occurs during the auction of block number 16345867 from point of view of the PBS network. We observe 17 builders submitting bids to the relay Agnostic Gnosis. Eventually, builder $0xb194b$(...) wins the auction with a bid value of 0.0467 ETH.

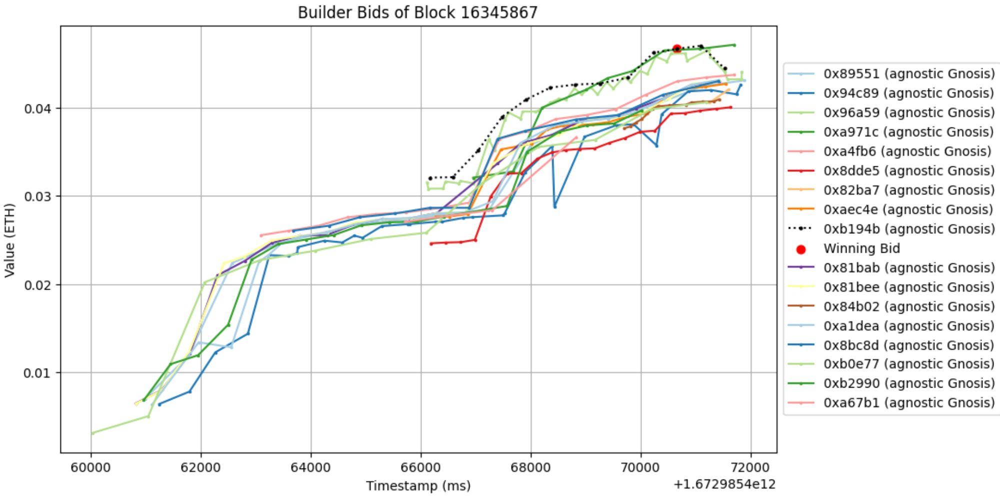

Figure 3.8: Example auction that was observed for block number 16345867. The above figure illustrates all bids received by relay agnostic Gnosis where each line represents one builder.[10]The x-axis indicates the elapsed time and the y-axis the bid value in ETH. Each bid is illustrated as a dot, the winning bid is highlighted in red and the winning builder as the dashed black line.

3.3.3 Overall Auction Analysis

In this section we present general insights regarding the auction dataset in order to shed light on overall market characteristics and dynamics. Table 3.2 contains several statistics of the dataset. We observe that the relay Eden is distinct compared to the other two relays in terms of number of bids per auction and number of builders per auction. One plausible reason for this is that the Eden relay does

[8] If the waiting time is too long, the block may not receive a sufficient number of attestations from the validator network, resulting in the original block being discarded and the proposer losing out on profits. Thus, proposers essentially balance profitability and risk of failure [2].

[10] The endpoints of the relays Eden and Ultrasound did not send any bids for this slot, which is common when proposers only choose to connect to some relays.

no connect to external builders (although anyone can apply to become a builder [22]). In contrast, Ultrasound and Agnostic Gnosis operate on a permissionless model, welcoming any builder to connect. Notably, the mean and median winning bid value is relatively similar for all 3 relays.

		Overall	Eden	Ultrasound	Agnostic Gnosis
Number of bids / auction	Mean	547	123	561	610
	Median	481	32	486	535
Number of builders / auction	Mean	43	7	44	44
	Median	49	4	47	48
Winning bid value	Mean	0.153 ETH	0.176 ETH	0.147 ETH	0.159 ETH
	Median	0.063 ETH	0.065 ETH	0.063 ETH	0.062 ETH

Table 3.2: General data characteristics of builder bidding auction datset

Number of bids per auction

Fig. 3.9 visualizes the number of bids submitted by builders for each auction. We observe significant differences over time as well as between the 3 relays. Specifically, auctions won by the Eden relay typically register fewer bids per auction than those by the Ultrasound and Agnostic Gnosis relays. As previously mentioned, one plausible reason for this is that the Eden relay does no connect to external builders while Ultrasound and Agnostic Gnosis operate on a permissionless model.

Interestingly, the number of bids of the Ultrasound and Agnostic Gnosis relays follow a similar trend over time: starting at approximately block number 16100000 (Dec-02-2022), the number of bids started increasing significantly, regularly surpassing 1500 bids by mid-January. At around block number 17000000 (Apr-07-2023) this trend reversed and the number of bids per auction dropped back down to between 50-500 bids. In order to assess whether the sharp drops in the number of bids at approx. block 16700000 and block 17000000 are either (1) due to changes in the number of bidders per auction, or (2) due to the quantity of bids per bidder per auction, we conducted two correlations tests: firstly between the mean number of bids per builder per auction vs. the total number of bids per auction, and secondly, between the number of bidders per auction vs. the total number of bids per auction. Both tests resulted in a positive Spearman Rank Correlation Coefficient (0.785 and 0.568 respectively) and a small p-value, indicating a statistically significant positive correlation for both. Refer to Table C.1 in the Appendix for the specific values of the correlation tests. As both correlations are positive, this approach does not yet tell us what caused the observed dips in Fig. 3.9, and thus we additionally plotted the data points from the correlation tests in scatter graphs (see Fig. C.2 and Fig. C.1 in the Appendix). Fig. C.2 shows that number of builders per auction significantly drops around block numbers 16700000 and 17000000 similar to the total number of bids, while the mean number of bids per builder (Fig. C.1) does not demonstrate equally notable changes. This would indicate that the sharp drops are primarily due to a decrease in the number of builders participating in the auctions, however, it remains unclear what caused this change in the number of participating builders.

To explore whether auctions with a higher block reward[11] lead to increased competition amongst bidders, we examined correlations between (1) the block reward and the number of bids, and (2) the block reward and the number of bidders. Our findings indicate a weak negative correlation across all relays: the number of bids per auction and block reward correlate with a Spearman rank coefficient of -0.140, while the number of bidders per auction and block reward have a coefficient of -0.070. However, given that the coefficients are very close to zero, the measured relationships are relatively insubstantial. All results, including tests for individual relays, are included in Table C.1 in the Appendix. In conclusion, these findings do not support our initial hypothesis, suggesting that a higher block reward does not necessarily intensify bidding competition.

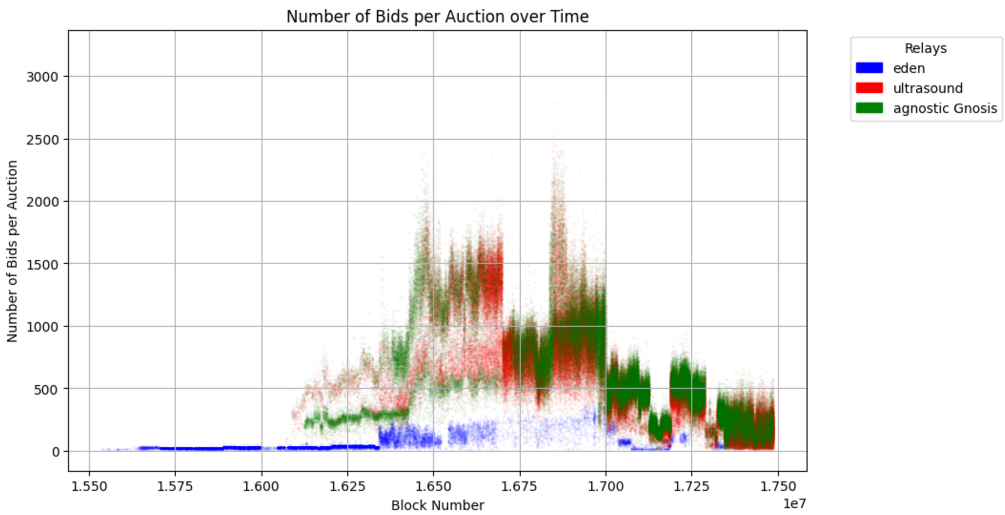

Figure 3.9: The above figure illustrates the number of bids submitted to each auction. The x-axis indicates the auctions' block numbers and the y-axis the number of bids. Each dot in the scatter graph represents the number of bids submitted to a single auction of the relays Eden, Ultrasound or Agnostic Gnosis.

[11]Refer to section 3.2 for the definition and derivation of an auction's block reward

Success frequency of builders

In the following section we analyze the auction success frequency of different builders. The weekly success frequency of a builder is defined as follows:

$$\textbf{Weekly success frequency} = \frac{\text{Number of won auctions per week}}{\text{Number of participated auctions per week}} \tag{3.1}$$

Overall, the mean weekly success frequency of all builders is 2.37% and the median lies at 0.33%. Fig. 3.10 illustrates the auction success frequency of all builders over time. During the earlier phase (up to December 2022), the data points are sparse, primarily because the Ultrasound and Agnostic Gnosis relays had yet to go live and gain market traction, as shown by Fig. 3.2. Overall, we observe a strong discrepancy between a minority of builders that obtain a substantially higher success frequency (up to 35%) in contrast to the majority of builders whose success frequencies are typically below 5%. This distinction is also emphasized by the difference between the smaller median and higher mean lines. Additionally, the success frequency of all builders demonstrates a high variability [12]. In April/Mai 2023 we observe an increase in success frequencies, but the cause of this is unclear. For a break-down of success frequencies into individual relays, one can refer to Appendix Fig. C.4, C.6 and C.5.

[12] This is even more visible on a log scale as depicted in Fig. C.3 in the Appendix.

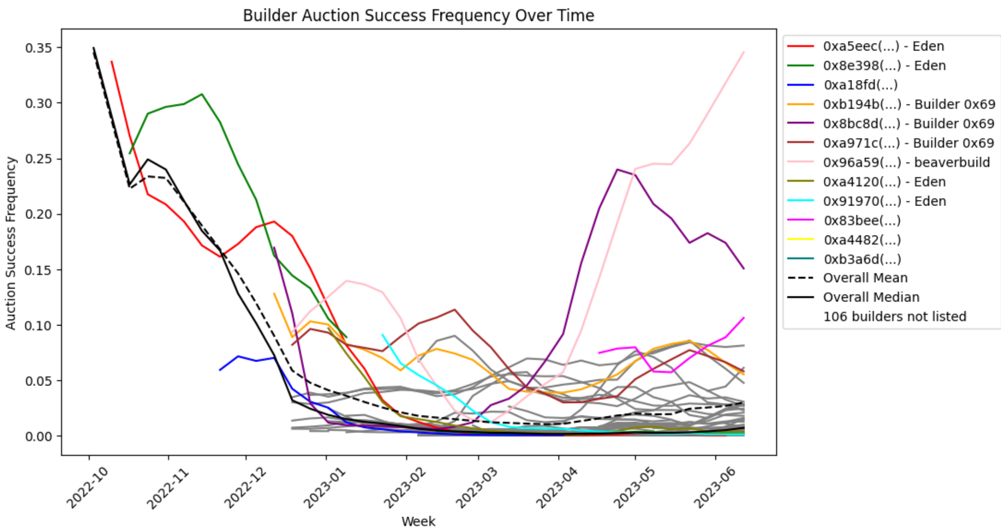

Figure 3.10: The above figure illustrates the auction success frequency of builders over time. The x-axis indicates the elapsed time and the y-axis the weekly success frequency 3.1. Each line represents the 4-week moving average of the weekly success frequency of a builder. Only the top 12 builders are listed in the legend and displayed in color in the figure. For builder addresses which are known (see Table A.2), we have added the corresponding builder organization to the label.

One hypothesis we investigated was, if the number of bids submitted during an auction is relates to the success frequency of a builder. To test this notion, we conducted a correlation test between the mean number of bids of each winning bidder per day and the daily success frequency of each winning bidder. Since neither dataset is normally distributed, we chose the Spearman Rank Correlation Coefficient. The computation yielded a coefficient of 0.266, suggesting of a positive correlation. To further validate this, we used the bootstrap method, which gave a 95% confidence interval for the Spearman correlation between 0.209 and 0.284. This range reaffirms the presence of a positive monotonic trend between the two datasets which confirms our hypothesis, that a higher success frequency is correlated with a higher number of bids. All values of the correlation test can be found in Table C.1 in the Appendix.

3.3.4 Examination of Individual Bidding Strategies

In this section we dive deeper into builders bidding strategies, aiming to understand patterns which guide their bid submissions during auctions. In particular, we focus on differentiating and analyzing the following two strategies:

- **Fraction-based bidding** is a relatively simple, non-adaptive bidding strategy. The bidder simply calculates the bid value as a fixed fraction of the total block value (e.g. 99%). As the block value increases over the course of an auction, the bidder submits increasing bids. This strategy is invariant to other bidder's behavior.

- **Reactive counterbidding** is a bidding strategy where a bidder reacts to the opponents bids and places a higher bid to win the auction. In the analyzed builder auctions, we refer to this strategy when a bidder submits bids based on the block value, and additionally adapts their bid submissions based on the opponents behavior.

The goal of our analysis is to evaluate the prominence of each strategy, and assess if builders can be categorized based on using one or the other strategy.

As a first approach we measured the bid value difference between the highest (winning) bid and the subsequent highest bid (i.e. the highest losing bid). This is based on the notion, that a builder who adopts reactive counterbidding would likely exhibit a smaller increment compared to one using a fraction-based bidding strategy. In Fig. 3.11 we summarize the mean bid difference of each builder in a histogram. The results of all individual auctions are illustrated in Fig. C.7 and Histogram C.8 in the Appendix. The histogram indicates that a certain group of builders achieve a small bid difference (between 0.0-0.0025), which may indicate that these builders use a reactive counterbidding strategy. However, we cannot draw any definite conclusions, as this could also be caused by other behaviors: two builders may include similar transactions in their block and bid a similar fraction of the total block reward. This case would also result in a small bid difference even in the absence of reactive counterbidding. Nevertheless, this analysis suggests that there exists high potential for refinement of many builders bidding strategy as a large bid difference results in unnecessary tokens being passed on to the proposer. In the next paragraph we move to an alternate approach to identify reactive counterbidding, which is based on a more rigorous notion.

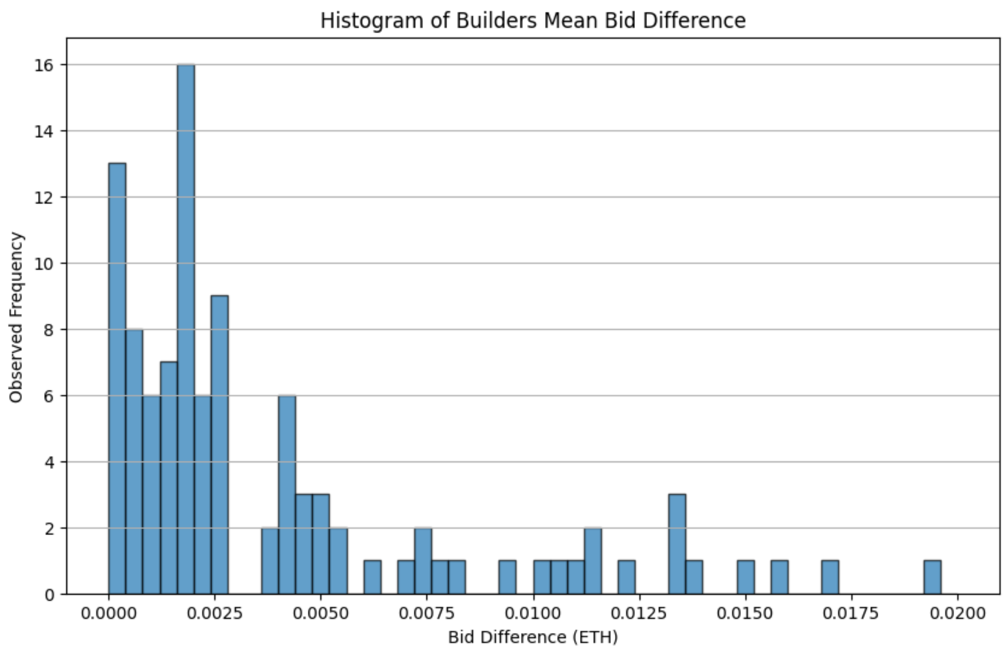

Figure 3.11: The above histogram illustrates the occurrences of builders mean weekly bid difference, where the bid difference refers to the increment between the winning bid value and the second highest bid value (highest loosing bid value). The x-axis indicates the mean bid difference in ETH and the y-axis the observed frequency of a given bid difference. As we are primarily interested in of small bid differences, the x-axis is cut off at $0.02ETH$.

In a bidding auction, builders may submit the same block multiple times with different bid values (see section 3.3.2). To identify builders using a reactive bidding strategy, we assume the notion that reactive bidders submit identical blocks with different bid values, as opposed to fraction-based bidders, which only submit one bid for each block hash. Thus, we calculate the number of duplicate block hash submissions of each builder per auction. Histogram 3.12 categorizes builders based on their mean number of duplicates. Although the majority of builders submit 0 duplicates (left-most bin), we indeed observe several builders submitting duplicate block hashes: 15% of all builders submit on average more than 4 duplicates. The highest builder (0x90000098...) on histogram 3.12 averages at over 80 duplicates per auction. Based on our initial notion, this analysis would thus indicate that a notable proportion of builders follow a reactive bidding strategy. However, despite this, when counting the total duplicate block submissions per builder per auction over all auctions, we determine that 72.5% of builders per auction submit zero duplicate block hashes. Thus, a large fraction of builders do not use a reactive bidding strategy and have a high potential for improving their bidding strategy.

Lastly, we tested our hypothesis, that bidders using a reactive bidding strategy, achieve a higher success frequency. This may either be due to (1) reactive bidding directly achieving a higher success frequency than other strategies (e.g. by bidding a smaller increment compared to the total block value) or, (2) builders using optimized bidding strategies may also tend to use more sophisticated block building strategies (resulting in a higher total block value). To evaluate this hypothesis, we conducted a correlation test between the weekly mean number of duplicate submissions of builders and the mean weekly success frequency of builders. The Spearman Rank Correlation Coefficient of 0.137 indicates a weak positive correlation between the number of duplicate submissions and the success frequency of builders. Thus, this test indicates that, as suspected, builders using reactive bidding strategies achieve a higher success frequency. However, note that the Spearman Rank Correlation Coefficient is close to zero and its significance might be limited. All values of the correlation test can be found in Table C.1 in the Appendix.

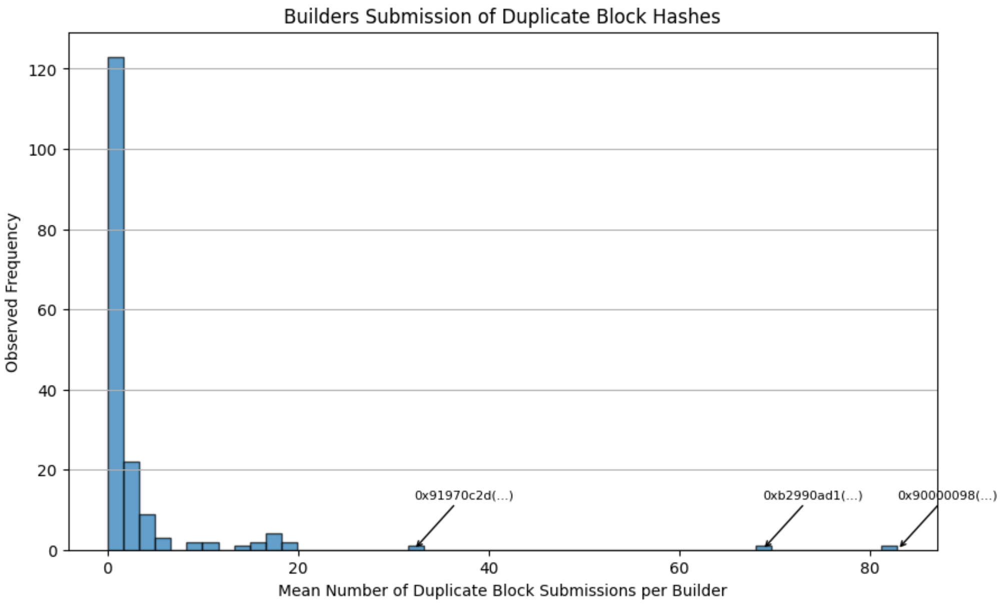

Figure 3.12: The above histogram illustrates the duplicate block hash submissions of builders. We calculated the mean duplicate submissions per auction per builder. The x-axis indicates the mean number of duplicate submissions and the y-axis the corresponding observed frequency.

Chapter 4

Conclusion and Future Research

The purpose of this thesis was to investigate player interactions and behaviors within the PBS ecosystem. Initially, we provided an overview of the stakeholders that occupy the market. Our observations revealed that both the relay and builder segment of the market are of an oligopolistic structure, primarily dominated by a handful of players (Section 3.1).

Next, we gave a comprehensive overview of how the block rewards and MEV profits are shared between searchers, builders and proposers (for blocks built by Flashbots builders). Our findings revealed that proposers consistently hold the largest profit portion (75-90%), followed by searchers (5-20%) and builders (typically $<5\%$). However, when breaking the profits into individual players, the Flashbots builders on average achieve the highest profit. This may indicate a disproportionate profit margin favoring builders over searchers and proposers, however, this would have to be confirmed by analyzing a larger number of builders. Future research could dive deeper into this aspect. Moreover, we revealed a high profit variability between individual searchers and proposers, and as a result, we cannot draw any decisive conclusions regarding which stakeholder category is most profitable (Section 3.2).

In the last section of this thesis, we introduced a model for builder auctions that arise when builders compete for their block to be selected by a proposer. Firstly, we assess several auction heuristics (bidding frequency, auction participants, success frequency) in practice, and secondly, evaluate the presence of two concrete bidding strategies: fraction-based bidding and reactive counterbidding. Our finding suggest that while the majority of builders do not seem to use a reactive bidding strategy, a notable fraction of bidders (15%) do. Furthermore, our analysis revealed that many builders employ a very simple and straightforward bidding strategy and there is high potential for refinement of these strategies. In future research one could build upon these results by either (1) formally evaluating an ideal bidding strategy in this game model or (2) practically testing different bidding approaches (Section 3.3).

Before concluding this thesis I would like to mention two final thoughts which became evident to me through the course of this study. Firstly, this thesis continuously revealed that the concept of PBS is still in its infancy (volatility between players in the network, observed instabilities during major changes, concentrated markets, etc.). The PBS ecosystem is yet to undergo significant evolution before achieving maturity. Secondly, PBS sheds light on some philosophical questions tied to the Ethereum blockchain: there is a fundamental clash between the vision of decentralization in Ethereum and the need for pragmatic solutions to market forces. This tension manifests in various ways, the following are a few examples: (1) PBS strives to democratize MEV access for all, however, currently less than 10 organizations operate the critical relay segment - far from what one can call a decentralized network. (2) Further, private transaction channels, on one hand, may be viewed as a centralizing force, but they also serve to protect users from malicious entities. (3) Or lastly, one can consider the Lido network, which controls over 30% of the staked ETH and dominates the proposer segment [19], however, at the same time Lido strives for decentralization within its own network of users.

In the future, it will be extremely interesting to observe what smart solutions will be employed to address these existing challenges and what path will be chosen balancing the vision of decentralization and on-ground pragmatics.

Bibliography

[1] P. Daian, S. Goldfeder, T. Kell, Y. Li, X. Zhao, I. Bentov, L. Breidenbach, and A. Juels, "Flash boys 2.0: Frontrunning in decentralized exchanges, miner extractable value, and consensus instability," in *2020 IEEE Symposium on Security and Privacy (SP)*, 2019.

[2] A. Wahrstätter1, L. Zhou, K. Qin, D. Svetinovic, and A. Gervais, "Time to bribe: Measuring block construction markets," 2023.

[3] L. Heimbach, L. Kiffer, C. F. Torres, and R. Wattenhofer, "Ethereum's proposer-builder separation: Promises and realities," 2023.

[4] Ethereum: The merge. Available online. Ethereum Foundation. Accessed on June 2, 2023. [Online]. Available: https://ethereum.org/en/roadmap/merge/

[5] The state of the merge: An update on ethereum's merge to proof of stake in 2022. Available online. Consensys. Accessed on November 4, 2023. [Online]. Available: https://consensys.io/blog/the-state-of-the-merge-an-update-on-ethereums-merge-to-proof-of-stake-in-2022

[6] Ethereum developers documentation. Available online. Ethereum Foundation. Accessed on June 2, 2023. [Online]. Available: https://ethereum.org/en/developers/docs/

[7] Sandwich overview | eigenphi. Available online. EigenPhi. Accessed on June 2, 2023. [Online]. Available: https://eigenphi.io/mev/ethereum/sandwich

[8] Proposer/block builder separation-friendly fee market designs. Available online. Ethereum Research. Accessed on June 2, 2023. [Online]. Available: https://ethresear.ch/t/proposer-block-builder-separation-friendly-fee-market-designs/9725

[9] S. Yang, F. Zhang, K. Huang, X. Chen, Y. Yang, and F. Zhu, "Sok: Mev countermeasures: Theory and practice," 2022.

[10] Mev-boost architecutre. Available online. Flashbots. Accessed on June 2, 2023. [Online]. Available: https://github.com/flashbots/mev-boost

[11] Flashbots docs. Available online. Flashbots. Accessed on June 2, 2023. [Online]. Available: https://docs.flashbots.net

[12] Office of foreign assets control. Available online. U.S. DEPARTMENT OF THE TREASURY. Accessed on November 4, 2023. [Online]. Available: https://ofac.treasury.gov

[13] Flashbots. Relay api. Available online. Accessed on June 15, 2023. [Online]. Available: https://flashbots.github.io/relay-specs/

[14] Flashbots. Mev-boost capella upgrades. Available online. Accessed on July 28, 2023. [Online]. Available: https://flashbots.notion.site/MEV-Boost-Capella-Upgrades-00cea01704794f6eb4f792c55b69c441

[15] Binance. Binance api endpoint. Available online. Accessed on July 29, 2023. [Online]. Available: https://api.binance.com/api/v3/klines

[16] Maximal extractable value (mev). Available online. Ethereum. Accessed on November 11, 2023. [Online]. Available: https://ethereum.org/en/developers/docs/mev/

[17] T. Gupta, M. M. Pai, and M. Resnick, "The centralizing effects of private order flow on proposer-builder separation," 2023.

[18] Etherscan. Ether daily price (usd) chart. Available online. Accessed on November 4, 2023. [Online]. Available: https://etherscan.io/chart/etherprice

[19] Lido. Available online. Accessed on September 25, 2023. [Online]. Available: https://lido.fi/ethereum

[20] Stakefish. Available online. Accessed on September 25, 2023. [Online]. Available: https://stake.fish

[21] Flashbots. Relay api notion documentation. Available online. Accessed on October 21, 2023. [Online]. Available: https://flashbots.notion.site/Relay-API-Documentation-5fb0819366954962bc02e81cb33840f5#854339c909a042d0bbca6e8f8069674e

[22] E. Network. Eden relay overview. Available online. Accessed on October 26, 2023. [Online]. Available: https://docs.edennetwork.io/eden-relay/overview/

Appendix A

Market Dynamics

Relay	**URL endpoint**
Flashbots	https://boost-relay.flashbots.net
Bloxroute (ethical)	https://bloxroute.ethical.blxrbdn.com
Bloxroute (max profit)	https://bloxroute.max-profit.blxrbdn.com
Bloxroute (regulated)	https://bloxroute.regulated.blxrbdn.com
Eden	https://relay.edennetwork.io
Ultrasound	https://relay.ultrasound.money
Agnostic Gnosis	https://agnostic-relay.net
Manifold	https://mainnet-relay.securerpc.com
Aestus	https://mainnet.aestus.live
Blocknative	https://builder-relay-mainnet.blocknative.com

Table A.1: List of relay endpoints queried.

Builder, Website and Source	Builder Pubkey Addresses
Flashbots https://boost-relay.flashbots.net https://docs.flashbots.net/flashbots-mev-boost/block-builders	0xa1dead01e65f0a0eee7b5170223f20c8f0cbf122eac3324d61afbdb33a8885ff8cab2ef514ac2c7698ae0d6289ef27fc 0x81beef03aafd3dd33ffd7deb337407142c80fea2690e5b3190cfc01bde5753f28982a7857c96172a75a234cb7bcb994f 0x81babeec8c9f2bb9c329fd8a3b176032fe0ab5f3b92a3f44d4575a231c7bd9c31d10b6328ef68ed1e8c02a3dbc8e80f9 0xa1defa73d675983a6972e8686360022c1ebc73395067dd1908f7ac76a526a19ac75e4f03ccab6788c54fdb81ff84fc1b
BloXroute (Max profit) https://docs.bloxroute.com https://docs.bloxroute.com/apis/mev-solution/list-of-bloxroute-builders	0x8b8edce58fafe098763e4fabdeb318d347f9238845f22c507e813186ea7d44adecd3028f9288048f9ad3bc7c7c735fba 0x94aa4ee318f39b56547a253700917982f4b737a49fc3f99ce08fa715e488e673d88a60f7d2cf9145a05127f17dcb7c67 0x976e63c505050e25b70b39238990c78ddf0948685eb8c5687d17ba5089541f37dd3c45999f2db449eac298b1d4856013 0xafc9274fe595e8cff421ab9e73b031f0dff707ea1852e2233ff070ef18e3876e25c44a9831c4b5f802653d4678ccc31f
BloXroute (Regulated) https://docs.bloxroute.com https://docs.bloxroute.com/apis/mev-solution/list-of-bloxroute-builders	0xb9b50821ec5f01bb19ec75e0f22264fa9369436544b65c7cf653109dd26ef1f65c4fcaf1b1bcd2a7278afc34455d3da6 0x80c7311597316f871363f8395b6a8d056071d90d8eb27defd14759e8522786061b13728623452740ba05055f5ba9d3d5 0x965a05a1ba338f4bbbb97407d70659f4cea2146d83ac5da6c2f3de824713c927dcba706f35322d65764912e7756103e2
BloXroute (Ethical) https://docs.bloxroute.com https://docs.bloxroute.com/apis/mev-solution/list-of-bloxroute-builders	0x95701d3f0c49d7501b7494a7a4a08ce66aa9cc1f139dbd3eec409b9893ea213e01681e6b76f031122c6663b7d72a331b 0xb086acdd8da6a11c973b4b26d8c955addbae4506c78defbeb5d4e00c1266b802ff86ec7457c4c3c7c573fa1e64f7e9e0 0x82801ab0556f7df1fb9bb3a61ca84beea8285a8dc3c455a7ea16a8b2993fe06058e0e7d275b28ea5d9f2ae995aa72605
Eden https://docs.edennetwork.io Pubkey addresses from paper [3]	0xa5eec32c40cc3737d643c24982c7f097354150aac1612d4089e2e8af44dbeefaec08a11c76bd57e7d58697ad8b2bbef5 0x8e39849ceabc8710de49b2ca7053813de18b1c12d9ee22149dac4b90b634dd7e6d1e7d3c2b4df806ce32c6228eb70a8b 0x91970c2db7c12510acb2e9c45844f7de602f83a7f31064f7ca04a807b607d7aebfc0abda73c036a92e5c3e56ebca04b7 0xa412007971217a42ca2ced9a90e7ca0ddfc922a1482ee6adf812c4a307e5fb7d6e668a7c86e53663ddd53c689aa3d350
Titan Builder https://docs.titanbuilder.xyz https://docs.titanbuilder.xyz/builder-public-keys	0xb67eaa5efcfa1d17319c344e1e5167811afbfe7922a2cf01c9a361f465597a5dc3a5472bd98843bac88d2541a78eab08 0x94a076b27f294dc44b9fd44d8e2b063fb129bc85ed047da1cefb82d16e1a13e6b50de31a86f5b233d1e6bbaca3c69173 0xb26f96664274e15fb6fcda862302e47de7e0e2a6687f8349327a9846043e42596ec44af676126e2cacbdd181f548e681 0x95c8cc31f8d4e54eddb0603b8f12d59d466f656f374bde2073e321bdd16082d420e3eef4d62467a7ea6b83818381f742 0xabf1ad5ec0512cb1adabe457882fa550b4935f1f7df9658e46af882049ec16da698c323af8c98c3f1f9570ebc4042a83
rsync-builder https://rsync-builder.xyz/docs https://rsync-builder.xyz	0x978a35c39c41aadbe35ea29712bccffb117cc6ebcad4d86ea463d712af1dc80131d0c650dc29ba29ef27c881f43bd587 0x83d3495a2951065cf19c4d282afca0a635a39f6504bd76282ed0138fe28680ec60fa3fd149e6d27a94a7d90e7b1fb640 0x945fc51bf63613257792926c9155d7ae32db73155dc13bdfe61cd476f1fd2297b66601e8721b723cef11e4e6682e9d87 0x8e6df6e0a9ca3fd89db2aa2f3daf77722dc4fbcd15e285ed7d9560fdf07b7d69ba504add4cc12ac999b8094ff30ed06c 0x8aab0ed724d2c7f94af139bd2249ab511f08474ac69e761e56918403c81c358f5f8a6d61c62a86dc4cd7bcad935f49d9 0xacfdcf458829f4693168a57d0659253069d687682bc64ec130d935ecb6e05ccfb80c138bd3cf53546c86715696612ec8
Blocknative https://www.blocknative.com/block-building Pubkey addresses from paper [3]	0x9000009807ed12c1f08bf4e81c6da3ba8e3fc3d953898ce0102433094e5f22f21102ec057841fcb81978ed1ea0fa8246 0x8000008a03ebae7d8ab2f66659bd719a698b2e74097d1e423df85e0d58571140527c15052a36c19878018aaebe8a6fea 0xa66f3abc04df65c16eb32151f2a92cb7921efdba4c25ab61b969a2af24b61508783ceb48175ef252ec9f82c6cdf8d8fd 0xa00000a975dffbd1ef61953ac6c90b52b70eb0188eb9d030774346c9248f81e875f7e8bc56c4bbbda297a9543cfa051d
Builder 0x69 https://docs.builder0x69.io Pubkey addresses from paper [3]	0xb194b2b8ec91a71c18f8483825234679299d146495a08db3bf3fb955e1d85a5fca77e88de93a74f4e32320fc922d3027 0x8bc8d110f8b5207e7edc407e8fa033937ddfe8d2c6f18c12a6171400eb6e04d49238ba2b0a95e633d15558e6a706fbe4 0xa971c4ee4ac5d47e0fb9e16be05981bfe51458f14c06b7a020304099c23d2d9952d4254cc50f291c385d15e7cae0cf9d 0xa4fb63c2ceeee73d1f1711fadf1c5357ac98cecb999d053be613f469a48f7416999a4da35dd60a7824478661399e6772 0xb8fceec09779ff758918a849bfe8ab43cea79f6a98320af0af5b030f6a7850fcc5883cb965d02efb10eed1ffa987e899
Manifold https://kb.manifoldfinance.com	0xa25f5d5bd4f1956971bbd6e5a19e59c9b1422ca253587bbbb644645bd2067cc08fb854a231061f8c91f110254664e943
beaverbuild https://beaverbuild.org Pubkey addresses from paper [3]	0x96a59d355b1f65e270b29981dd113625732539e955a1beeecbc471dd0196c4804574ff871d47ed34ff6d921061e9fc27 0xb5d883565500910f3f10f0a2e3a031139d972117a3b67da191ff93ba00ba26502d9b65385b5bca5e7c587273e40f2319 0x8dde59a0d40b9a77b901fc40bee1116acf643b2b60656ace951a5073fe317f57a086acf1eac7502ea32edcca1a900521 0xaec4ec48c2ec03c418c599622980184e926f0de3c9ceab15fc059d617fa0eafe7a0c62126a4657faf596a1b211eec347

Table A.2: The table above contains a list of well-known builders included in this analysis and their respective pubkey addresses. Note that identifying pubkey addresses of certain organizations can be challenging, and this table represents a non-exhaustive list.

Appendix B

Distribution of Profits

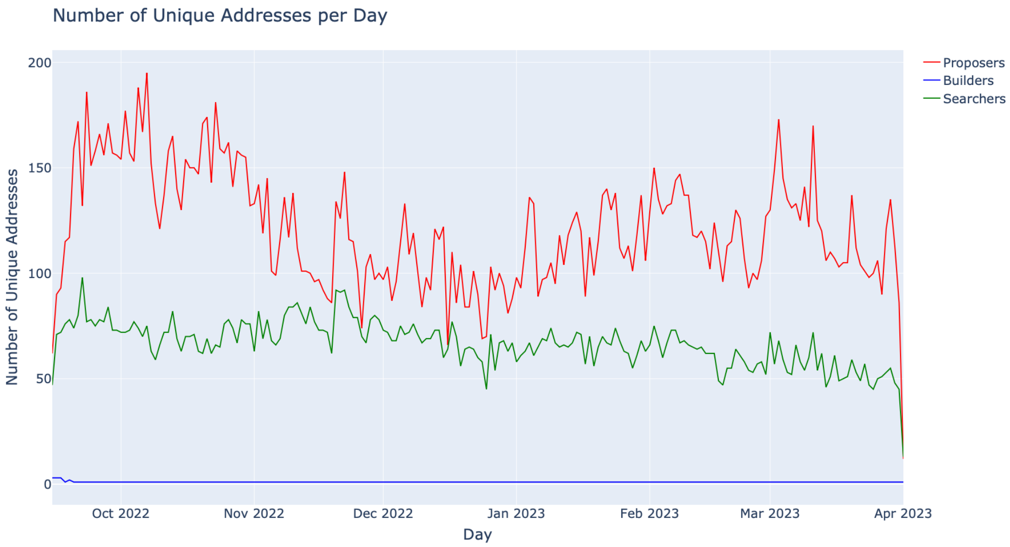

Figure B.1: Number of unique addresses profit is shared between (for searchers, builders and proposers).

Variables	**Spearman Rank Correlation Coefficient**	**p-value**
(1) Searcher Profit (2) ETH to USD price changes	-0.009	0.962
(1) Builder Profit (2) ETH to USD price changes	-0.264	0.166
(1) Proposer Profit (2) ETH to USD price changes	-0.139	0.471

Table B.1: The above table displays the results of correlation tests between the searcher, builder and proposer profits compared to price change of Ethereum (ETH to USD price changes). We used the Spearman Rank Coefficient.

Appendix C

Builder Auctions

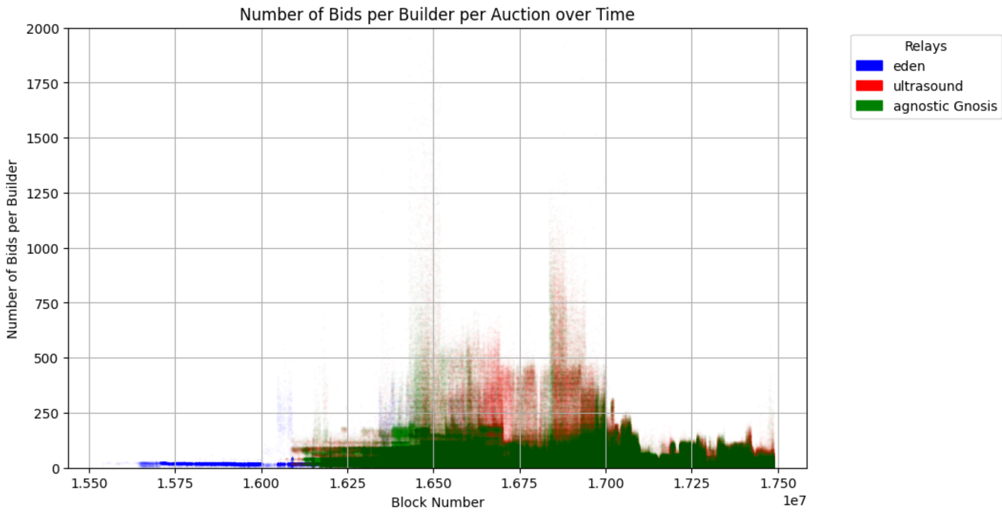

Figure C.1: The above figure illustrates the number of bids builder builder per auction (i.e. block number) over time. The x-axis indicates the elapsed time and the y-axis the number of bids. Each dot represents the number of bids of a single builder during one auction (i.e. for each auction we plot one dot for each unique builder). The scatter dots are colored by the relay which the block was proposed through.

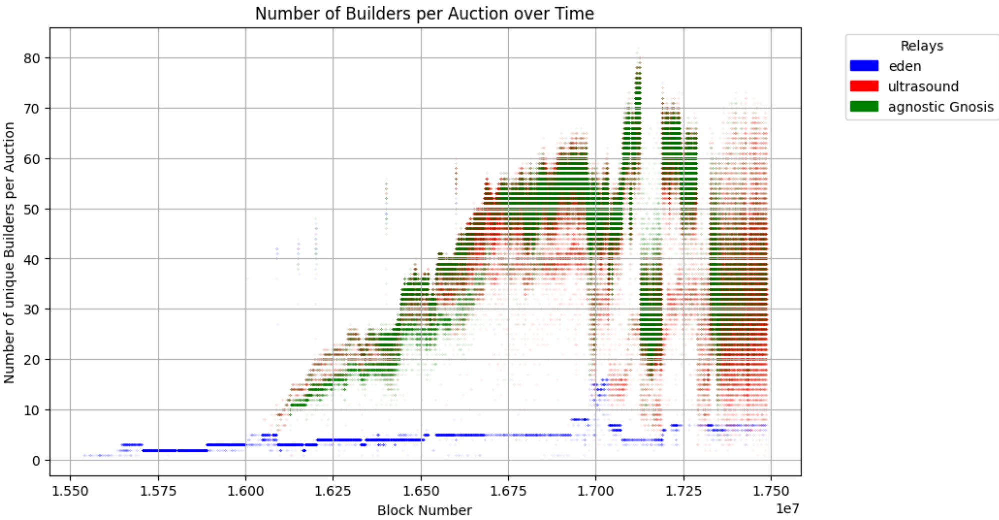

Figure C.2: The above figure illustrates the number of builders per auction (i.e. block number) over time. The x-axis indicates the elapsed time and the y-axis the number of unique builders. Thus, each dot represents, the number of unique builders of a single auction and it is colored by the relay which the block was proposed through.

Variables	**Dataset**	**Spearman Rank Correlation Coefficient**	**p-value**	**95% Confidence Interval**
(1) Nr of Bids / Auction (2) Mean Nr. of Bids / Builder / Auction	All Blocks	0.785	0.0	(0.783, 0.787)
(1) Nr of Bids / Auction (2) Nr. of Builders / Auction	All Blocks	0.568	0.0	(0.566, 0.570)
(1) Nr. of Bids / Auction (2) Block Reward	All Relays	-0.140	0.0	(-0.145, -0.136)
	Eden	-0.189	0.0	(-0.202, -0.177)
	Ultrasound	-0.122	0.0	(-0.128, -0.116)
	Agnostic Gnosis	-0.079	0.0	(-0.087, -0.071)
(1) Nr of Bidders / Auction (2) Block Reward	All Relays	-0.070	0.0	(-0.076, -0.066)
	Eden	-0.117	0.0	(-0.129, -0.103)
	Ultrasound	-0.038	0.0	(-0.044, -0.031)
	Agnostic Gnosis	0.016	0.0004	(0.007, 0.024)
(1) Mean Nr of Bids / Auction / Winning Builder (2) Success Frequency / Builder	All Builders	0.266	$9.72e^{-19}$	(0.209, 0.284)
(1) Nr of Duplicate Blocks / Auction / Builder (2) Success Frequency / Builder	All Builders	0.137	$3.9e^{-7}$	(0.086, 0.190)

Table C.1: Correlation tests of various data points of builder auctions. We tested for a correlation between variable x and variable y using the Spearman Rank Coefficient. Bootstrapping was used to determine the 95% confidence interval.

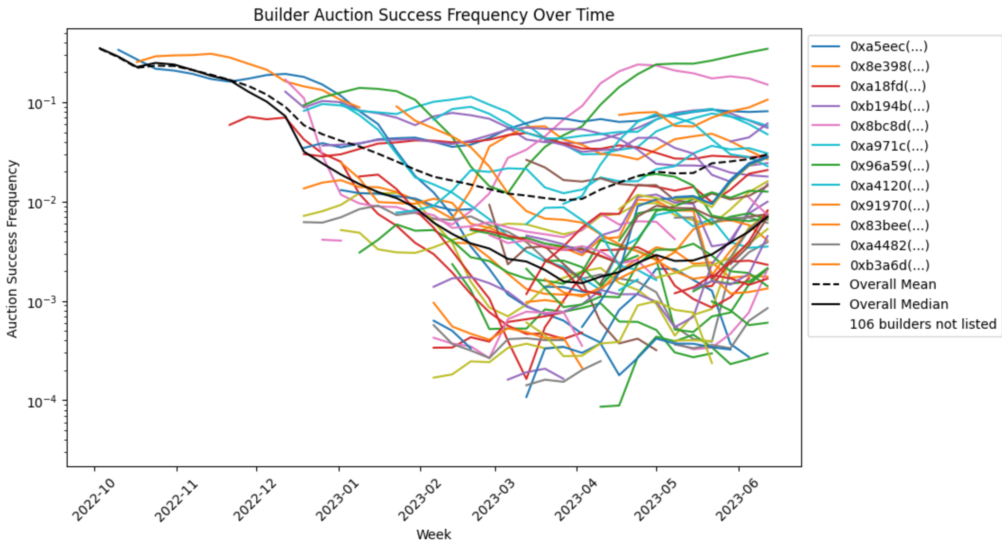

Figure C.3: The above figure illustrates the auction success frequency of builders over time. The x-axis indicates the elapsed time and the y-axis is a log scale of the success frequency (*number of auctions won / number of auctions participated*). Each line represents the 4-week moving average of the weekly success frequency of a builder. Only the top 12 builders are listed in the legend, in addition to the weekly mean and median.

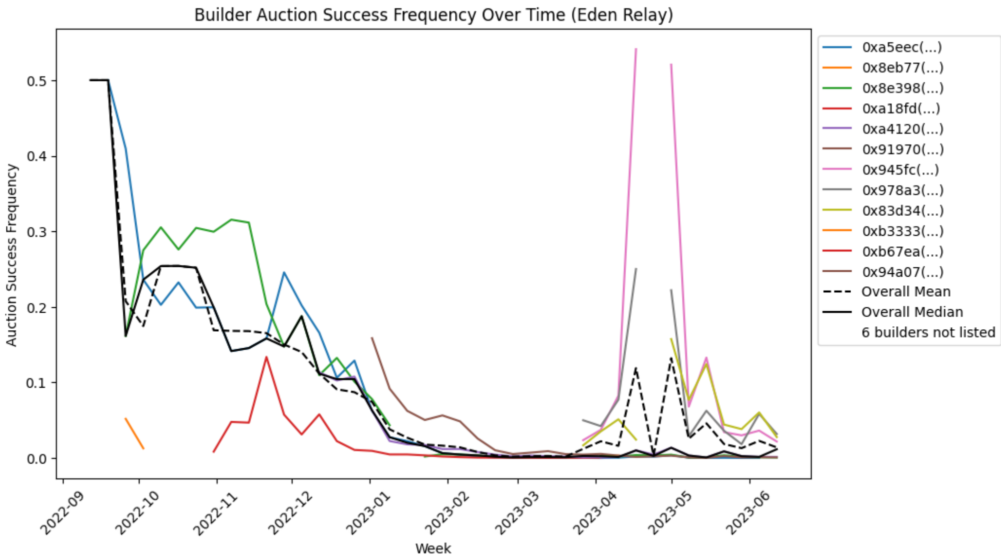

Figure C.4: The above figure illustrates builders' bidding success frequency of auctions where the winning block was proposed through the relay Eden. The x-axis indicates the elapsed time and the y-axis the success frequency (*number of auctions won / number of auctions participated*). Each line represents the 4-week moving average of the weekly success frequency of a builder. Only the top 12 builders are listed in the legend, in addition to the weekly mean and median.

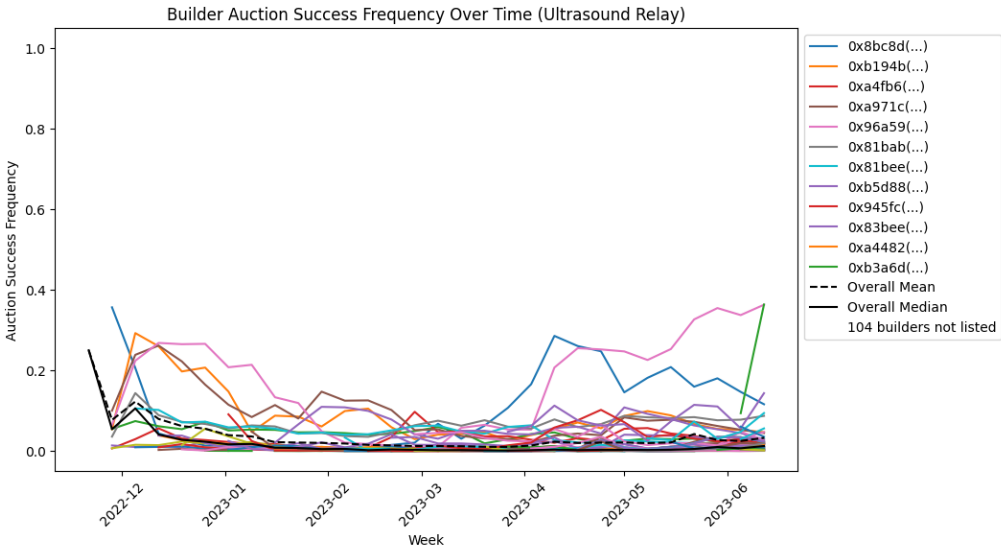

Figure C.5: The above figure illustrates builders' bidding success frequency of auctions where the winning block was proposed through the relay Ultrasound. The x-axis indicates the elapsed time and the y-axis the success frequency (*number of auctions won / number of auctions participated*). Each line represents the 4-week moving average of the weekly success frequency of a builder. Only the top 12 builders are listed in the legend, in addition to the weekly mean and median.

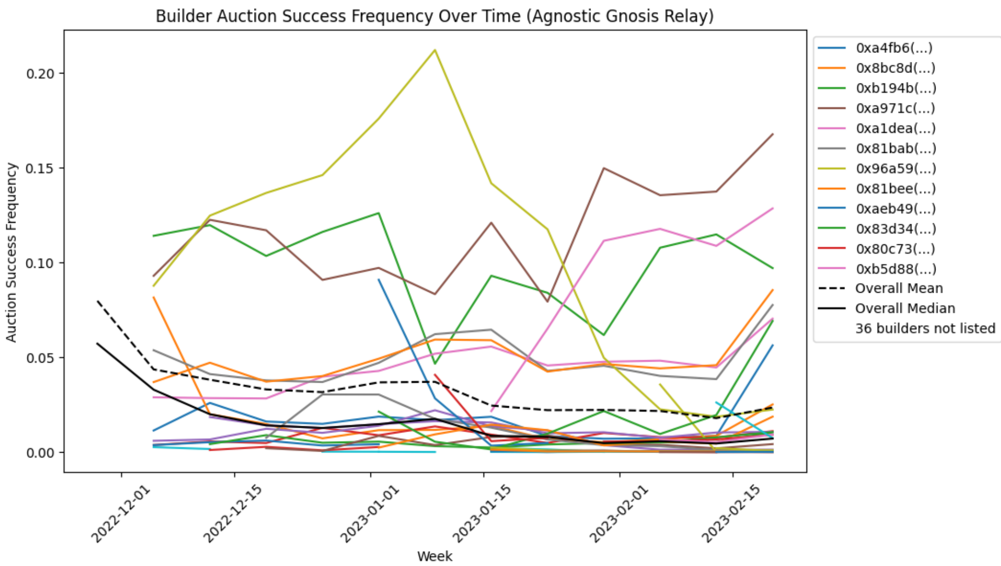

Figure C.6: The above figure illustrates builders' bidding success frequency of auctions where the winning block was proposed through the relay Agnostic Gnosis. The x-axis indicates the elapsed time and the y-axis the success frequency (*number of auctions won / number of auctions participated*). Each line represents the 4-week moving average of the weekly success frequency of a builder. Only the top 12 builders are listed in the legend, in addition to the weekly mean and median.

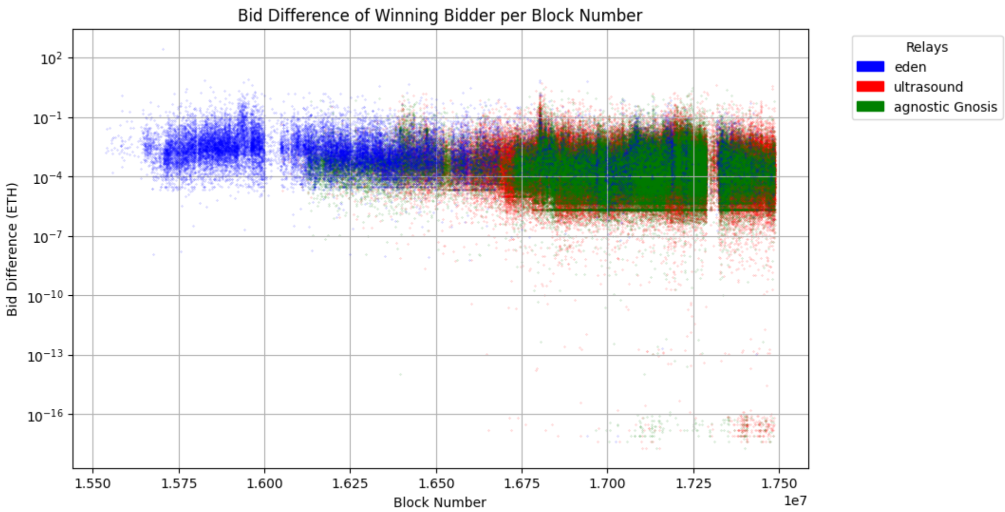

Figure C.7: The above scatter plot illustrates the bid difference of all auctions, where the bid difference refers to the increment between the winning bid value and the second highest bid value (highest loosing bid value). The x-axis indicates the block number and the y-axis the bid difference in ETH. Each dot represents the bid difference of a single auction, colour by the relay through which the block was proposed.

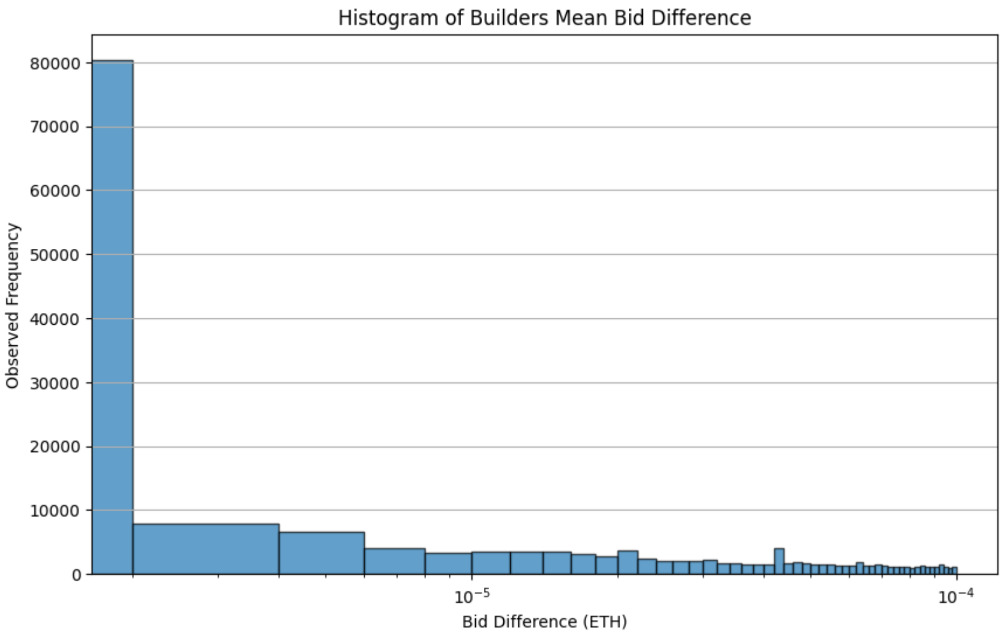

Figure C.8: The above histogram illustrates the bid difference of all auctions, where the bid difference refers to the increment between the winning bid value and the second highest bid value (highest loosing bid value). The x-axis indicates the bid difference between 0-0.0001 ETH (log-scale) split into 50 bins, and the y-axis the observed frequency. The minority of bid differences with a value ≤ 0.0001 ETH are not shown in this scale, to improve the readability of the histogram.