To Compete or Collude: Bidding Incentives in Ethereum Block Building Auctions

Fei Wu
King's College London
London, UK
fei.wu@kcl.ac.uk

Thomas Thiery
Ethereum Foundation
Lisbon, Portugal
thomas.thiery@ethereum.org

Stefanos Leonardos
King's College London
London, UK
stefanos.leonardos@kcl.ac.uk

Carmine Ventre
King's College London
London, UK
carmine.ventre@kcl.ac.uk

Abstract

The block-building process on the Ethereum network has changed significantly with an upgrade of its consensus protocol. Network participants access blocks through block building auctions at a decentralized financial market, termed *builder market*, where *builders* vie for the right to build blocks and earn *Maximal Extractable Value* (MEV) rewards. This paper employs empirical game-theoretic analysis to examine builders' strategic bidding incentives in the Ethereum block building auctions, termed *MEV-Boost auctions*. We study various scenarios with different auction game settings and evaluate how critical elements such as network connectivity and access to MEV opportunities impact builders' strategic bidding incentives. Through our analyses, we highlight the challenge of creating a decentralized yet competitive builder market.

CCS Concepts

• **Applied computing → Online auctions**.

Keywords

Decentralized Finance, Ethereum Blockchain, Auction, Maximal Extractable Value, Empirical Game-Theoretic Analysis

ACM Reference Format:
Fei Wu, Thomas Thiery, Stefanos Leonardos, and Carmine Ventre. 2024. To Compete or Collude: Bidding Incentives in Ethereum Block Building Auctions. In *5th ACM International Conference on AI in Finance (ICAIF '24), November 14–17, 2024, Brooklyn, NY, USA*. ACM, New York, NY, USA, 9 pages. https://doi.org/10.1145/3677052.3698629

1 Introduction

With a total value of about 90b USD locked in Decentralized Finance (DeFi), permissionless blockchains like Ethereum [3] have proven highly successful in providing financial services. Unlike traditional financial systems, DeFi operations are transparent, with transactions broadcasted across peer-to-peer networks and recorded on the ledger through the consensus of network participants, hereafter referred to as validators, who process transactions.

Without central authorities to control activities in decentralized financial markets, self-interested validators often maximize their profits at the expense of user experience. These profits, known as *Maximal Extractable Value* (MEV), are derived from validators' authority over transaction inclusion, exclusion, and sequencing in the blocks [5]. MEV has emerged as a concept of significant importance and complexity, as potential problems caused by MEV constitute a severe bottleneck in the public adoption of decentralized financial markets and limit the positive impact of the underlying blockchain technology on our society. Specifically, while MEV can lead to increased revenue for validators, it also introduces centralizing risks by disproportionately favoring well-resourced validators with substantial capital and computational power, as extracting MEV demands significant resources.

To reduce computational demands and enhance decentralization for validators, Ethereum introduced *Proposer-Builder Separation* (PBS) [8]. It allows the proposer[1] to outsource the task of block construction and MEV extraction to specialized entities called *builders* at the *builder market*, thus eliminating the edges of well-resourced validators in MEV extraction. Under the current PBS framework, proposers can opt into a software named *MEV-Boost* [9] to access blocks through trusted intermediaries known as *relays* at the builder market, where builders compete in an auction by submitting their block with a bid to relays. This auction is termed *MEV-Boost auction* and operates as an English auction. The auctioneer, i.e., the proposer, terminates the auction by selecting the block with the highest bid from the relays and collects a significant portion of MEV in auction revenue. The winning builder, upon their block being selected, also stands to potentially realize profits.

The bids submitted by builders depend on the value of the blocks they produce, which originates from two primary sources: public transactions and private orderflow. The value of public transactions consists of transaction fees and MEV from user transactions broadcast in the network, pending in the *public mempool* accessible to all builders. The value of private orderflow consists of transaction fees and MEV from transaction *bundles privately* sent from orderflow providers. Orderflow providers detect MEV extraction opportunities [18, 22] and strategically bundle their transactions with user

ICAIF '24, November 14–17, 2024, Brooklyn, NY, USA

ACM ISBN 979-8-4007-1081-0/24/11
https://doi.org/10.1145/3677052.3698629

[1]A validator chosen at each consensus round to add a new block to the blockchain.

transactions to extract value from them. These bundles are sent privately, either exclusively to one builder or to multiple builders.

Ideally, builders compete with equal resources, and a large portion of MEV is captured through the auction mechanism by the proposer, rendering a competitive and decentralized builder market. However, to gain a competitive edge in auctions, builders invest in optimizing their connection to orderflow providers and relays, targeting lower *latency* and higher orderflow access. Often, builders forge agreements with orderflow providers for access to orderflow. Moreover, some large builders operate their orderflow providers and initiate a relay to improve latency further and gain exclusive access to high-value orderflow. While this strategy enhances their bidding performance, it also contributes to centralization in the builder market; currently, despite over 30 active builders, the market is dominated by a few large builders. While over 90% of Ethereum blocks being built through MEV-Boost auctions [27], there is little literature on builders' strategic bidding incentives in the auctions. While studies have highlighted the effects of latency and exclusive orderflow on builders' performance and market centralization [1, 2, 12–14], practical insights into how these advantages affect builders' incentives remain limited.

Another concern under the current PBS framework is the engagement in *timing games* by proposers [24], which involves strategically delaying the auction termination and the block proposal to obtain a higher bid, optimizing their profits. Such practices can undermine Ethereum's consensus. A viable countermeasure could involve relays enforcing stricter timeliness and rejecting further bids [23]. However, there is notably scarce literature on how such timing games countermeasure will affect builders' strategic bidding.

In this paper, we address these questions and thereby contribute to bridging the gap concerning the interaction between builders' strategic bidding incentives and the MEV-Boost auction mechanism. Given the inherent complexities of MEV-Boost auctions, conducting a direct theoretical analysis is challenging. Instead, we use the agent-based simulation framework and the empirically validated bidding strategies proposed by [31] (Section 3) to conduct an empirical game-theoretic analysis (EGTA). By leveraging empirical game-theoretic techniques, we efficiently reduce the dimensionality of the games and manage the computational demands, enabling extensive simulation and analysis of MEV-Boost auction (Section 4). We study builders' strategic bidding in auctions under varying scenarios and examine how advantages in latency and orderflow access as well as relay enforcement impact builders' incentives (Section 5&6). At a high level, this paper makes the following contributions:

(1) We find that, under ideal conditions of a builder market (similar latency and orderflow access), builders are incentivized to marginally increase their bids to outbid each other (i.e., *colluding*) rather than bidding their full valuation. Although this collusion potentially enhances decentralization in the market, with builders having an equal chance of winning the auction, the winning bids are much lower than the actual block values. As a result, the MEV-Boost auction mechanism does not efficiently capture MEV, and builders retain a large proportion of MEV for themselves.

(2) We find that latency improvements enable builders to bid strategically and maximize their profit. In contrast to the symmetric (idealized) scenario studied above, we also find that the disparity in latency between builders can serve as a critical element that increases the efficiency of the current MEV-Boost auction mechanism. We further show that significant advantages in private orderflow access incentivize builders to refuse to collude and bid truthfully to increase their market shares (thus dominate the market).

(3) We demonstrate the timeliness enforced by relay impacts builders' bidding incentives when builders have different latencies. We show that this enforcement contributes to enhancing the auction efficiency by forcing competitive players with a latency advantage to bid their full valuation.

2 Background and Related Works

Proof-of-Stake Ethereum: The Ethereum blockchain switched on its Proof-of-Stake (PoS) consensus protocols in September 2022 [7]. Network participants of PoS Ethereum are called validators, and each validator must deposit 32 ETH minimum as collateral. In PoS Ethereum, time progresses in 12-second *slots* [4]. At each slot, a validator is randomly chosen to be the proposer to propose a block at the start of the slot, as specified by the protocol [6].

PBS and MEV-Boost: Currently, Ethereum does not have an in-protocol PBS design due to the challenge of ensuring trustless interactions where proposers cannot steal MEV rewards and builders must deliver promised bid value. MEV-Boost [9] is an out-of-protocol implementation of PBS and is deployed with PoS Ethereum. It introduces relays, third-party trusted intermediaries, to resolve the reliance issue between proposers and builders. The proposer can either build the block themselves (i.e., *local production*) or have it built by builders via MEV-Boost. The MEV-Boost auction for the block of slot n typically starts around the beginning of slot $n-1$ and terminates at the end of slot $n-1$ when the proposer selects and signs the winning block header. Within each auction cycle, builders submit blocks alongside their bids to relays. Relays are tasked with validating blocks from builders, selecting the winning block with the highest bid, and only forwarding its block header to the proposer. If the proposer receives multiple block headers with the same bid value from different relays, the winning block selection can be viewed as a random selection using the hash value as a tiebreaker.[2] Upon the proposer returning the signed header to the relay, the relay publishes the full block to the network.

Related Works. To the best of our knowledge, the study on builders' strategic bidding behaviors in MEV-Boost auctions starts from [17], where the authors identified bid erosion and bid shielding. Subsequent analyses by [25] observed varied bidding behaviors among builders in MEV-Boost auctions, confirming diverse strategies. [21] demonstrated that a latency disadvantage adversely affects builders' bidding performance. [1, 2, 12–14] shows private orderflow contributes to better bidding performance and its centralizing effect on the builder market. [16] discussed Ethereum's evolving orderflow landscape, emphasizing the challenge of balancing increased competition with market decentralization. [31] introduced a game-theoretic model for MEV-Boost auctions alongside four bidding strategies, and conducted simulations to assess the effects of latency and orderflow access on bidding performance. [32] further

[2]https://github.com/flashbots/mev-boost/blob/74a8ecb36ee36952d7b622570f6a7b7195ddc599/server/service.go#L449.

empirically assessed the competitiveness and efficiency of MEV-Boost auctions. [20] demonstrated a positive correlation between a builder's market share and the diversity of their orderflow, as well as a positive correlation between builder profitability and their access to exclusive private orderflow providers. [23, 24] analyzed the impact of proposer timing games on Ethereum consensus and suggested relay enforcement as potential mitigation. Building upon these works, our work is the first to examine builders' strategic bidding incentives in MEV-Boost auctions and analyze how latency improvements, orderflow access advantage, and relay enforcement influence these incentives.

3 MEV-Boost Auction Game Model

We employ the MEV-Boost auction model and bidding strategies proposed by [31]. To make our paper self-contained, we briefly introduce here the necessary background.

We consider a set of $N = \{1, \dots, n\}$ builders competing in the MEV-Boost auction game. Each builder, indexed by i, employs a bidding strategy s_i which can be described as a function $\beta_{s_i} : X \to \mathbb{R}_+$ so that the bid of player i at time t is $\beta_{s_i}(x_{i,t})$, where $x_{i,t} \in X$ represents a vector of input variables at time $t \geq 0$. Whenever redundant, we will omit the dependence on s_i and $x_{i,t}$ and simply write $\beta_{i,t}$. These inputs are discussed next, where we use the terms *player* and *builder* interchangeably.

Public signal $P(t)$ represents the maximum extractable value from public transactions broadcast in the mempool at time t, accessible to all builders. New pending transactions are submitted to the mempool as the auction advances. This process is modeled by a compound Poisson process, where the number of transactions $N(t)$ up to time t follows a Poisson distribution with rate λ_p and each transaction's value, V_j, is randomly drawn from a log-normal distribution. The public signal, $P(t)$, is the cumulative sum of values of $N(t)$ transactions, given by the equation:

$$P(t) := \sum_{j=1}^{N(t)} V_j,$$

where $N(t) \sim \text{Poisson}(\lambda_p \cdot t)$ and $V_j \sim \text{Log-normal}(\xi_1, \omega_1)$.

Private signal $E_i(t)$ represents the *private orderflow* secured from orderflow providers. Builders often receive similar orderflow because some orderflow are commonly shared among them. To account for the exclusiveness and correlation of orderflow among players, we introduce an *orderflow access probability*, $\pi_i \in [0, 1]$, for each player $i \in N$, to represent that player's probability of accessing each orderflow. The probabilities, $(\pi_i)_{i \in N}$, remain constant throughout the auction interval. Similar to the public signal, the number of private orderflow $N_i(t)$ accessed by player i up to time t follows a Poisson distribution with rate λ_e but is also influenced by π_i. Each orderflow's value O_j is randomly drawn from a log-normal distribution. The private signal, $E_i(t)$, of player i is

$$E_i(t) := \sum_{j=1}^{N_i(t)} O_j,$$

where $N_i(t) \sim \text{Poisson}(\lambda_e \cdot t \cdot \pi_i)$ and $O_j \sim \text{Log-normal}(\xi_2, \omega_2)$. $E(t)$ denotes the total value of private orderflow, during the slot at time t. Thus, the *aggregated signal*, $L_i(t)$, of player i and the *total signal* in the auction, $L(t)$, at time t can be given by combining the public signal and the private signal:

$$L_i(t) := P(t) + E_i(t) \text{ and } L(t) := P(t) + E(t).$$

Given the positive correlations between bid arrival times and bid values [28], we assume that all MEV opportunities are persistent throughout the auction and have uniform values for all builders who have access to them.

The ***latency***, $\Delta_i > 0$ of each player depends mostly on that player's network connectivity and geographic location. It quantifies the delay in the relay's acceptance of bids relative to the player's access to a signal update and their subsequent bid submission. It is assumed to be known and constant during the auction and to only affect the player's bidding action.

The ***profit margin***, $pm_i > 0$, quantifies player i's risk tolerance and profit expectations. The ***valuation*** of player i, $v_i(t)$, represents the highest bid value that player i can place at time t while ensuring a positive profit, and is defined as: $v_i(t) := L_i(t) - pm_i$.

The ***current highest bid***, denoted by $\max_{j \in N}\{\beta_{j,k}\}_{k \leq t}$, represents the highest bid among all bids submitted by all builders up to time t. This information is known to all players.

The auction interval is defined as $[0, T]$, where T denotes the time when the proposer selects the highest winning bid. Instead of T being exactly equal to 12 seconds as expected, the winning bid is typically selected around $T = 12$ seconds due to factors such as latency or timing games. Thus, T is randomly drawn from a normal distribution with mean 12 and standard deviation σ.

3.1 Strategies and payoffs

We consider a strategy space S with three strategies, $S = \{s_0, s_1, s_2\}$, where s_0 represents the *truthful* strategy, s_1 represents the *adaptive* strategy, and s_2 represents the *last-minute* strategy. We will qualify players by their strategy; so, e.g., "truthful players" are those playing the truthful strategy. These strategies are defined as follows.

Truthful players consistently bid their valuation. Adaptive players either incrementally exceed the current highest bid by a marginal value $\delta > 0$, or bid their valuation when outbidding is unfeasible. Last-minute players hold their bids initially and reveal their valuation starting from the expected auction termination (the 12-second slot boundary). The bidding behaviours of these strategies are summarised in Table 1, where we use $v_i(t)_+$ to denote the positive part of $v_i(t)$, i.e., $v_i(t)_+ := \max\{v_i(t), 0\}$.

Together with the assumption that $pm_i > 0$ for all builders $i \in N$, the above definitions implicitly assume that builders are not willing to win the auction at a negative profit. However, in current practice, builders are willing to win the auction by subsidizing [32]. The main reason that we introduce this assumption is that we consider static (non-repeated) auction games.

We consider 10 players in the auction, since currently, the top 10 builders build 98.65% of the total blocks built via MEV-Boost. Each player $i \in N = \{1, \dots, 10\}$, selects a pure strategy $s_i \in S$, and bids according to their chosen strategy throughout the auction interval. The collection of strategies selected by all players forms a strategy profile $s = (s_1, s_2, ..., s_{10})$. The payoff $u_i(s)$ of player i is given by

$$u_i(s_i, s_{-i}) := \begin{cases} L_i(t_w) - \beta_{s_i}(x_{i,t_w}) & \text{if } \beta_{s_i}(x_{i,t_w}) = \max_{j \in N}\{\beta_{j,k}\}_{k \leq T}, \\ 0 & \text{otherwise.} \end{cases}$$

Table 1: Bidding strategies.

Strategy	Bid value at time $t \leq T$
Truthful	$\beta_{s_0}(x_{i,t}) = v_i(t)_+$
Adaptive	$\beta_{s_1}(x_{i,t}) = \min\{v_i(t), \max_{j\in N}\{\beta_{j,k}\}_{k\leq t}\} + \delta\}_+$
Last-minute	$\beta_{s_2}(x_{i,t}) = v_i(t)_+ \times \mathbf{1}\{t \geq \theta\}$, where $\theta = 12 - \Delta_i$.

where t_w denotes the submission time of the winning bid and $s = (s_i, s_{-i})$ as is standard.

3.2 Model calibration: orderflow estimation

The model is implemented with Agent-Based Modeling techniques. For technical reasons, we assume that time evolves at discrete time steps of 10ms increments. The state of the auction is updated after all the players take their bidding actions simultaneously.

To inform the settings of our model, we use the mempool data [10] and the on-chain data maintained by Flashbots and Dune for the period February 23 to April 7 2024 [11].[3] On average, each block contains 143 public transactions, with each transaction valued at approximately 0.00021 ETH, which collectively accounts for nearly 40% of the total block value. Additionally, builders earn an average profit of 0.0066 ETH from winning an auction. In the simulation, we assume that this profit margin is symmetric across all players.

It is worth noting that the available data sources exhibit a certain degree of bias concerning the private orderflow. The on-chain data *only* reveals the private orderflow included by the winning builder. Since that builder wins the auction, we presume that their private orderflow access surpasses that of other competing builders in the auction. However, the actual access to private orderflow by a builder remains undisclosed, irrespectively of their success in the auction. This lack of information stems from the data not being recorded on-chain and not being available from any Relay Data API.

Consequently, we assume that private orderflow access among players is randomly distributed on a domain which we can estimate from on-chain data. This approach guarantees that regardless of the auction's winner, the volume of private orderflow included in the winning block aligns with expectations set by on-chain metrics. Specifically, the average private orderflow of the top 10 builders contributes between 11.2% and 14.0% of the total transactions in their respective winning blocks. We subsequently infer that a 14.0% inclusion of private orderflow represents the maximal volume achievable by players, i.e., the total transaction number in the private mempool, corresponding to a private orderflow access probability of 100%. An 11.2% inclusion denotes the minimal threshold, equating to an 80% access probability. Thus, we delineate the distribution of π_i to be a uniform distribution spanning the interval $[0.8, 1.0]$.

4 Empirical Games

In this section, we introduce the definitions of the empirical games, the representations of payoffs in these games, and the tools that we utilize to solve them. Explicitly solving the above game is computationally hard. The reason is that as players can independently decide their bidding strategy, resulting in 3^{10} distinct strategy profiles. To tackle this issue, we follow an empirical game-theoretic approach where we exploit certain symmetries between players to reduce the game size. Specifically, each player is characterized by their latency and their private orderflow access probability (distribution). Accordingly, we analyze three variants in which either one or both of these attributes are uniform between some players.

Equal latency and private orderflow access distribution. We begin by analyzing games where all builders have the same latency and prior distribution on private orderflow access. Specifically, for every auction simulation, each player's probability π_i of accessing private orderflow is drawn from the same prior distribution, namely uniform on $[0.8, 1]$.

In this case, we can reduce the size of the underlying game by replacing it with an *anonymous game*, where a player's payoff is invariant to permutations of other players [30]. In other words, a player's payoff depends only on the number of other players playing each strategy. Therefore, we can represent a strategy profile by a vector of the number of players playing each strategy, which allows us to reduce the number of strategy profiles to $\binom{10+|S|-1}{10} = 66$.

To store the payoff information, we want to use the *heuristic payoff table* (HPT) [29], where *payoffs of each strategy* are stored as a function only of the number of players using it. However, while the private orderflow access probabilities are equal in expectation (drawn uniformly from the same prior distribution), every realization can be different. This implies that players can have different payoffs if they interchange their strategies which, in turn, implies that the payoffs of each strategy are not unique. Thus, HPT cannot be directly applied since this game is, in fact, asymmetric [26].

To overcome this issue, we let the payoff of each strategy be the *average payoff of the players* using it in each strategy profile. Moreover, we further reduce the impact of randomness introduced by each player's access to private orderflow, by letting each player's payoff be the average profit out of 1,000 auction simulations for each strategy profile. This leverages the fact that the players with the same latency using the same strategy tend to have the same payoff (in expectation) due to their private orderflow access being drawn from the same prior distribution.

Formally, let the HPT, $\mathcal{H} = (\mathcal{N}, \mathcal{U})$, where $\mathcal{N}$ is a matrix of profile representations of dimension $\binom{10+|S|-1}{10} \times |S|$, and $\mathcal{U}$ is a matrix of payoffs of the same dimension. Entry $\mathcal{N}_{k,j}$ in $\mathcal{N}$ describes the number of players choosing strategy s_j, $j \in \{0, 1, 2\}$ in strategy profile s^k, and entry $\mathcal{U}_{k,j}$ in $\mathcal{U}$ describes the average payoff of players choosing strategy s_j in profile s^k. $\mathcal{U}_{k,j}$ can be given by

$$\mathcal{U}_{k,j} = \begin{cases} \frac{1}{\mathcal{N}_{k,j}} \sum_{i:s_i=s_j} u_i\left(s^k\right) & \text{if } \mathcal{N}_{k,j} > 0, \\ 0 & \text{otherwise.} \end{cases}$$

Different latencies. While the previous case captures an idealized scenario, in practice, builders experience different latencies due to variations in their connectivity to orderflow providers and relays. In

[3] March 13 data is excluded due to an error caused by the Ethereum EIP-4844 upgrade.

this part, we consider games where players have the same distribution of private orderflow access but different latencies. Specifically, we examine scenarios involving 5 players with low latency and 5 players with high latency.

This setup, which resembles current practice, allows us to consider the auction game as a *role-symmetric game* [30], in which players are divided into two *roles* based on their latency: r_l for the 5 low-latency players and r_h for the 5 high-latency players. Within each role, the payoff of each strategy is represented by the average payoff of the players within that role adopting that strategy.

Formally, let the role, r_i of each player i be $\in \{r_l, r_h\}$, indicating whether player i belongs to the low-latency group (r_l) or the high-latency group (r_h). We extend the HPT $\mathcal{H} = (\mathcal{N}^l \times \mathcal{N}^h, \mathcal{U}^l \times \mathcal{U}^h)$. $\mathcal{N}^l \times \mathcal{N}^h$ is a matrix of strategy profile representations of the dimension of $\binom{5+|S|-1}{5}^2 \times 2|S|$, where $\mathcal{N}^l$ is a counts matrix for r_l and $\mathcal{N}^h$ is a counts matrix for r_h. $\mathcal{U}^l \times \mathcal{U}^h$ is a matrix of payoffs of the same dimension. Entry $\mathcal{N}^r_{k,j}$ in $\mathcal{N}^r, r \in \{l, h\}$, describes the number of players choosing strategy $s_j, j \in \{0, 1, 2\}$ within role r in the strategy profile s^k, and entry $\mathcal{U}^r_{k,j}$ in $\mathcal{U}^r, r \in \{l, h\}$ describes the average payoff of players within the role r choosing the strategy s_j in the profile s^k. $\mathcal{U}^r_{k,j}$ can be given by

$$\mathcal{U}^r_{k,j} = \begin{cases} \frac{1}{\mathcal{N}^r_{k,j}} \sum_{i:s_i=s_j, r_i=r} u_i\left(s^k\right) & \text{if } \mathcal{N}^r_{k,j} > 0, \\ 0 & \text{otherwise.} \end{cases}$$

Different private orderflow access distributions. In fact, the disparity in private orderflow access between builders is significant. We finally consider games where players have the same latency but different private orderflow access probabilities. Similarly, we examine scenarios involving 5 players with high private orderflow access probability and 5 players with low private access probability and consider the games as role-symmetric games.

Varied/fixed auction interval. Additionally, for all three empirical games above, we consider two distinct scenarios and study the players' incentives under these scenarios: 1) the auction interval varying around 12 seconds ($\sigma = 0.1$), where last-minute players face a 50% chance of successfully revealing their bids before the auction closes (Section 5), and 2) the auction interval being fixed to 12 seconds ($\sigma = 0$), i.e., relay enforcement of rejecting bids after the beginning of the slot, where last-minute players typically bid at the very end of the auction interval (Section 6).

4.1 α-Rank

To solve the above games, we employ the *α-Rank* algorithm [19]. The α-Rank algorithm describes a stochastic evolutionary process to model a selection-mutation process of a set of populations. It provides a dynamic solution to the game by understanding agents' behaviors and predicting their convergence. The solution (equilibrium) is presented by a ranking of strategy profiles with their stationary probabilities within the unique stationary distribution of the α-Rank Markov Chain, also known as the α-Rank score. This distribution indicates the average time the system spends in each strategy profile. To convey the equilibrium clearly, we present the results as the average number of players using each strategy across all the profiles weighted by their stationary probabilities, which captures the expected frequency of each strategy being used by players in the long run.

The probability $\rho^i_{\sigma,\tau}(s_{-i})$ that player i in a population with strategy σ switches to strategy τ, given the strategy choices s_{-i} of other players, is given by

$$\rho^i_{\sigma,\tau}(s_{-i}) = \begin{cases} \dfrac{1 - e^{-\alpha(u_i(\tau,s_{-i}) - u_i(\sigma,s_{-i}))}}{1 - e^{-m\alpha(u_i(\tau,s_{-i}) - u_i(\sigma,s_{-i}))}} & \text{if } u_i(\tau, s_{-i}) \neq u_i(\sigma, s_{-i}), \\ \dfrac{1}{m} & \text{if } u_i(\tau, s_{-i}) = u_i(\sigma, s_{-i}), \end{cases}$$

where m is the number of players in the population, and $u_i(\tau, s_{-i})$ and $u_i(\sigma, s_{-i})$ are the payoffs to player i when choosing strategies τ and σ, respectively. To ensure that even a tiny payoff difference will contribute to a strategy switch, the *ranking-intensity* α, must be sufficiently large. Instead of starting from a small value and increasing it exponentially as suggested in [19], which is computationally time-consuming, we use the method of calculating an estimated lower-bound value of α provided by OpenSpiel [15]. The lower-bound value of α is given by

$$\alpha > \frac{2}{\min_{u_i(\tau)>u_i(\sigma)}(u_i(\tau) - u_i(\sigma))}.$$

5 Empirical Game-Theoretic Analysis

In this section, we present our experimental results using the three game variants defined above. We study builders' incentives for choosing strategies under varying conditions, investigate the impact of latency improvements and orderflow access advantage on builders' incentives, and analyze the state of the MEV-Boost auction in equilibrium.

5.1 Collusion in the symmetric game

In the games where all 10 players share identical latency and the prior distribution of private orderflow access probability, the strategy profile where all 10 players adopt the adaptive strategy, demonstrates its dominance with a stationary probability of 0.99997. The result shows that, once the system reaches this state, transitioning to other states is exceedingly unlikely due to the evolutionary stability of this profile and the success of the strategies involved.

Specifically, in the auction simulation, when all the players employ the adaptive strategy, they increase their bids incrementally with the marginal value δ simultaneously, contributing to uniform bid values among all builders at the end. Consequently, the auction terminates with a relatively low winning bid value and a winner randomly selected, meaning all players have an equal chance to win. This scenario implies that a significant portion of block value accrues to the winning builder rather than to the proposer, highlighting inefficiencies in the auction mechanism to capture MEV.

To evaluate the auction's capability of capturing MEV, we define auction *efficiency* as the ratio of the winning bid to the total signal value. In the simulation, we use $\delta = 0.0001$ ETH, which results in a median auction efficiency of 46.65%. It is worth noting that further reduction in the marginal value δ can lead to even lower winning bid values and efficiency.[4] Figure 1 illustrates the correlation between auction efficiency and the value of δ.

[4]The varying δ values have negligible impact on the stationary distribution.

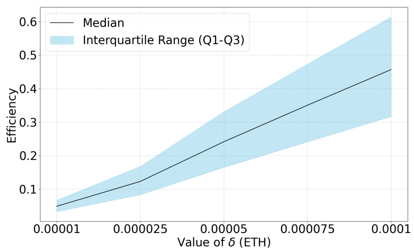

Figure 1: Auction efficiencies under varying δ values when all 10 builders use adaptive strategy and share an identical latency of 10ms.

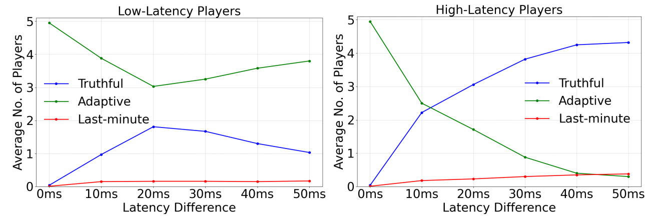

Figure 2: Average usage of each strategy by low-latency players (left) and high-latency players (right) across all profiles under varying latency differences as computed by α-Rank.

These findings indicate that, in scenarios where all builders experience identical latency and have comparable private orderflow access, given the defined strategy space, builders are disinclined to bid their full valuation (adopt the truthful strategy). Instead, they are incentivized to maximize their profitability by together increasing their bids incrementally with a small margin, sharing an equal chance of winning the auction. This equal chance to win might foster a more decentralized builder market, with builders potentially sharing the market equally. However, such behavior could be perceived as a form of collusion, which, while enhancing decentralization, reduces competitive bidding and auction efficiency.

While the builders are incentivized to maximize their profit by colluding, to prevent the proposer from falling back to local production, they need to ensure that their final bid matches or exceeds the maximum extractable value from the public mempool, i.e., public signal—approximately 40% of the total signal value. In our simulation, setting $\delta = 0.0001$ ETH results in a bid value corresponding to 46.65% (median) of the total block value, demonstrating that the simulated auctions remain effective under these settings.

5.2 Impact of latency

In the role-symmetric games where all players have the same distribution of private orderflow access but experience two types of latencies, the equilibrium of the previous (symmetric) situation can be disrupted: if all players employ the adaptive strategy, the market will not be equally shared anymore, as high-latency players have a lower chance of winning than low-latency players. It is worth mentioning that the adaptive strategy is more sensitive to latency variations, and a higher latency typically results in a slower reaction. Consequently, high-latency players may be incentivized to switch to the truthful or last-minute strategy, as the bidding behaviors defined by these two strategies are less affected by latency.

The primary factor influencing players' strategy choices in this case is the difference in latency between the low-latency and high-latency players. To analyze these effects, we consider scenarios under varying latency differences between low- and high-latency players, from 0ms (previous symmetric scenario) to 50ms, and analyze each latency difference scenario as a separate game. Specifically, low-latency players maintain a fixed latency of 10ms, while the latency of high-latency players starts at 10ms and increases by 10ms increments. We then analyze the equilibria of these games to understand the impact of latency on strategy choices.

Our findings show that when the latency difference is only 10ms, the equilibrium shows that players are still more incentivized to adopt the adaptive strategy and collude. As suggested by our simulations, despite a lower win rate of 4.45%, high-latency players continue to collude with low-latency players, attracted by the potential to capture approximately 50% of the block value upon winning. However, as the latency difference exceeds 10ms, high-latency players' win rates decline significantly, prompting a strategic shift towards the truthful strategy for improved performance.

Figure 2 presents the equilibria computed by α-Rank under varying latency differences. As the high-latency players' incentives for switching to the truthful strategy become increasingly strong, when the latency difference becomes 20ms, the equilibrium shifts, with low-latency players being incentivized to adopt the adaptive strategy while high-latency players being incentivized to adopt the truthful strategy. This shift occurs because the effectiveness of the adaptive strategy diminishes significantly for high-latency players as their latency increases, making the truthful strategy more appealing. In response to high-latency players' shift to the truthful strategy, we also observe a reaction from low-latency players when the latency difference is not greater than 20ms. As the latency difference increases, low-latency players have a stronger incentive to maintain the adaptive strategy, while high-latency players are better suited to the truthful strategy. Furthermore, when the latency for high-latency players is particularly high, even with a 50% risk of missing the submission window, the last-minute strategy proves more effective than the adaptive strategy. This is attributed to the decreasing effectiveness of the adaptive strategy as latency increases, combined with the occasional effectiveness of the last-minute strategy against the adaptive strategy employed by low-latency players.

The results underscore the significant impact of latency improvements on builders' incentives for strategic bidding. Although both low-latency and high-latency builders have comparable access to private orderflow in the game settings, latency advantages facilitate faster access to transactions and quicker bid updates. This capability proves crucial near the auction termination because if a late transaction occurs, low-latency builders can include it and update their bids before the auction closes, unlike high-latency builders who, despite having access to the same transaction, cannot update their bids in time. Thus, builders who benefit from a latency advantage are incentivized to adopt the adaptive strategy, thereby maximizing their profits by marginally outbidding. Conversely, builders with a latency disadvantage are compelled to bid truthfully to enhance performance and offset their latency disadvantage.

Furthermore, the results indicate that the latency difference between builders under the current PBS framework serves as a crucial element that makes MEV-Boost auction efficient. Although high-latency builders may win infrequently, their truthful bidding behavior still pressures low-latency adaptive players to place higher bids, thereby enhancing the auction efficiency.

5.3 Impact of orderflow access

We next proceed to analyze builders' incentives when their orderflow access probabilities are different (i.e., not from the same prior distribution). To isolate the effects of latency, we mirror our previous approach in the study of latency effects, and examine scenarios involving 5 low- and 5 high-orderflow players, all with the same latency of 10ms, as described in the third scenario in Section 4.

Surprisingly, despite varying differences in the orderflow access probability between low- and high-orderflow players, the equilibrium outcomes are consistent with the previous symmetric situation: all players play the adaptive strategy. This consistency arises because, under the current simulation settings, the ultimate bid value approximates the public signal value, as discussed in Section 5.1. As a result, both high-orderflow and low-orderflow players remain competitive at the ultimate bid value, rendering the differences in orderflow access inconsequential.

Nevertheless, there exists a scenario where the high-orderflow players are incentivized to adopt the truthful strategy and dominate the market without sharing it with the low-orderflow players. This is the case, when builders can dynamically adjust their profit margin based on their private orderflow volume. To study this effect, instead of having a fixed profit margin value symmetric for all players, we set the profit margin of each player to be equal to 50% of their private orderflow volume, i.e., 50% of their private signal. Similarly, we analyze each probability difference scenario as a separate game. We set the orderflow access probability of all players to 50% and increase the high-orderflow players' access probability in 10% increments.

Figure 3 presents the equilibria for the above scenarios. As we see, high-orderflow players are more incentivized to adopt the adaptive strategy and collude with low-orderflow players when their orderflow access and, hence, their profit margin, are low (difference below 30%). However, as their orderflow access probability and, thereby their profit margin, increase, they are increasingly incentivized to adopt the truthful strategy. Thus, they increasingly refuse to collude with low-orderflow players and capture a higher win rate. When the probability difference exceeds 30%, meaning high-orderflow players' access probability surpasses 80%, we observe a significant shift in equilibrium.

For low-orderflow players, their incentive to maintain the adaptive strategy remains strong. Given their limited access to orderflow, adopting the truthful strategy is suboptimal because high-orderflow players can easily outbid them. Consequently, their best chance to win is by playing the adaptive strategy and colluding with high-orderflow players. However, as high-orderflow players switch to the truthful strategy and refuse to collude, low-orderflow players adopting the adaptive strategy are unable to outbid the high-orderflow players, meaning that the adaptive strategy and the truthful strategy

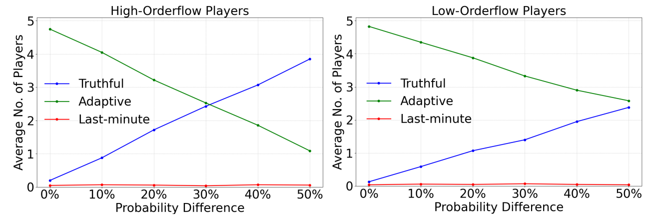

Figure 3: Average usage of each strategy by high-orderflow players (left) and low-orderflow players (right) across all profiles under varying probability differences.

are equally effective. Thus, we observe the usage of these two strategies by low-orderflow players tends to converge as the probability difference increases.

6 Impact of Relay Enforcement

In this section, we investigate how the relay enforcement of rejecting bids after the start of the slot affects builders' incentives, by fixing the simulated auction interval to match exactly the 12-second slot duration (i.e., $\sigma = 0$). We assume that the relay conducts this enforcement honestly. This enforcement effectively terminates the auction at the start of the slot, ensuring no further bids are accepted and eliminating the proposer's incentive for delaying bid selection. It is worth noting that the proposer is not incentivized to select a winning bid earlier, as they might miss a higher bid. Therefore, the builders know that the auction will be terminated at a fixed time point, i.e., the beginning of the slot.

As it turns out, the lack of ambiguity in auction termination significantly affects builders' incentives. Under non-random termination, the last-minute players, who no longer face a 50% chance of out-of-time revelation, will ultimately bid their valuation at the end of the auction, similar to truthful players. This strategy is particularly effective against the adaptive strategy, as it reveals the valuation at the final moment, thereby denying adaptive players any opportunity to react.

Next, we revisit the three game settings under the assumption that the auction terminates at exactly 12 seconds (and this is known to players). In the first symmetric situation, we find that the profile where all players adopt the adaptive strategy continues to dominate with a stationary probability of 0.99506. In the games where players are divided into high-orderflow players and low-orderflow players (Section 5.3), the relay enforcement also has a limited impact, as high-orderflow players are able to dominate the market with both truthful and last-minute strategies.

However, the equilibria of the games where players experience different latencies, in which low-latency players are incentivized to adopt the adaptive strategy and high-latency players are incentivized to adopt the truthful strategy, are disrupted. When the auction interval is deterministic, high-latency players use the last-minute strategy, which undermines the effectiveness of the adaptive strategy employed by low-latency players. This forces low-latency players to switch to either the truthful or last-minute strategies to remain competitive.

Figure 4 displays the experimental results. Similar to the scenarios with varying auction intervals, all players are still more

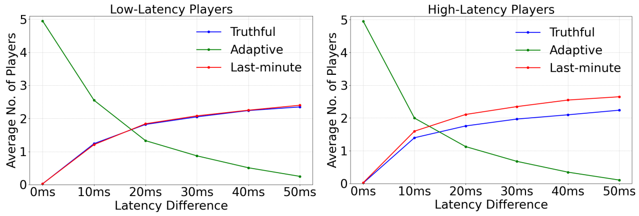

Figure 4: Average usage of each strategy by low-latency players (left) and high-latency players (right) across all profiles with relay enforcement under varying latency differences.

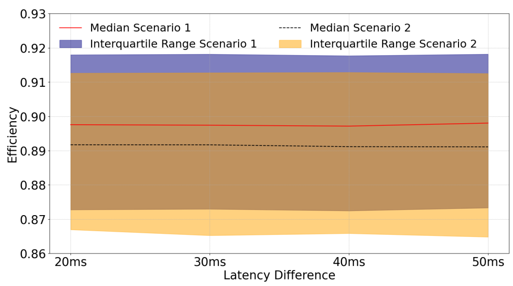

Figure 5: Auction efficiency of the most stable strategy profiles under relay enforcement (Scenario 1) and without relay enforcement (Scenario 2).

incentivized to collude by adopting the adaptive strategy when the latency difference does not exceed 10ms. However, as players' latency increases, the adaptive strategy becomes less effective for high-latency players, who then favor the last-minute strategy due to its potential to disrupt the adaptive bidding of low-latency opponents. This strategic shift prompts even the low-latency players to adopt either the truthful or last-minute strategies, as the effectiveness of the adaptive strategy diminishes in response to increasing latency differences and the heightened incentive for high-latency players to utilize the last-minute strategy.

While the relay enforcement has limited effect on builders' incentives for collusion under ideal conditions, it contributes to offsetting the negative impact on the auction efficiency caused by the latency asymmetries between builders, by forcing low-latency builders to abandon marginally outbidding and bid their full valuation. This is due to the increased effectiveness of the last-minute strategy which allows the high-latency builders to compete and curb inequalities. In turn, this has a noticeable effect on enhancing auction efficiency.

To study this effect, we compare the auction efficiency between the most stable strategy profiles under two scenarios: with enforcement (Scenario 1), where low-latency players adopt the naive strategy while high-latency players employ the last-minute strategy, and without enforcement (Scenario 2), where low-latency players utilize the adaptive strategy while high-latency players adopt the naive strategy. Figure 5 presents the simulation results. We show that, under varying latency differences between the low- and high-latency players, the auction efficiency is enhanced with relay enforcement.

7 Limitations

The results and analyses presented in this paper are inherently based on and constrained by the model, calibration methods, strategy space, and game definitions. To ensure tractable equilibria in the games with available computational resources, we analyzed the games as *only* one auction with limited information asymmetries (e.g., private orderflow and profit margin) among builders. Due to this constraint and the lack of data, we applied the estimation metrics and assumed symmetry among builders for these variables (Section 3.2).

The strategies were formulated simply: consistently bidding the full valuation (truthful strategy), marginally outbidding when feasible (adaptive strategy), and a last-minute strategy specialized for scenarios under relay enforcement. Whilst these strategies are considered naturally adopted by agents and empirically validated [31], it is conceivable to imagine that builders' strategic behaviors and MEV-Boost auction dynamics are far more complex. Our results provide limited insights if more asymmetric information, a richer strategy space, and consecutive auctions are considered. Despite these limitations, our analyses shed light on builders' strategic bidding incentives in MEV-Boost auctions.

8 Concluding Discussion

In this paper, we explore builders' incentives for strategic bidding in MEV-Boost auctions through empirical game-theoretic analysis. Our findings indicated that, with our modeling choices, under ideal conditions of a builder market that leads to decentralization, builders are incentivized to collude by marginally outbidding each other rather than competing by bidding their true valuation, resulting in a low auction efficiency for MEV capture. We demonstrated that builders can marginally outbid to maximize their profits with a latency advantage and can refuse collusion to dominate the market with greater orderflow access. We show that relay enforcement as a mitigation of timing games impacts builders' bidding incentives.

While the current PBS design enhances validator decentralization, it centralizes the builder market by shifting trustless interaction challenges from proposers and builders to builders and orderflow providers. Consequently, orderflow providers engage in private deals with builders, contributing to disproportionate orderflow distribution among builders. However, even if a trustless mechanism were available to distribute orderflow equally and privately among all builders, our results indicate that the current MEV-Boost auction mechanism would not efficiently capture MEV under such conditions as builders are incentivized to collude. The above results contribute to the ongoing discussion about the challenge of creating a decentralized yet competitive and efficient market and provide evidence that current market structures and mechanisms may require fundamental changes to address centralization and efficiency concerns. We highlight the importance of further research to validate the robustness of our findings, particularly by incorporating a more comprehensive strategy space and examining consecutive auction games, as well as exploring block building auction mechanisms that discourage collusion and promote genuine competition among builders.

References

[1] Data Always. 2023. The Influence of CeFi-DeFi Arbitrage on Order-Flow Auction Bid Profiles. https://ethresear.ch/t/the-influence-of-cefi-defi-arbitrage-on-order-flow-auction-bid-profiles/17258. Accessed: 2024-04-29.

[2] Maryam Bahrani, Pranav Garimidi, and Tim Roughgarden. 2024. Centralization in Block Building and Proposer-Builder Separation. arXiv:2401.12120 [cs.GT] https://arxiv.org/abs/2401.12120

[3] Vitalik Buterin et al. 2014. A next-generation smart contract and decentralized application platform. *white paper* 3, 37 (2014), 2–1.

[4] Vitalik Buterin et al. 2020. Combining GHOST and casper. *arXiv preprint arXiv:2003.03052* (2020).

[5] Philip Daian et al. 2019. Flash Boys 2.0: Frontrunning, Transaction Reordering, and Consensus Instability in Decentralized Exchanges. arXiv:1904.05234 [cs.CR] https://arxiv.org/abs/1904.05234

[6] Ethereum. 2022. Honest Validator - Ethereum Consensus Specifications. https://github.com/ethereum/consensus-specs/tree/dev/specs. Accessed: 2024-04-13.

[7] Ethereum. 2022. The Merge. https://ethereum.org/en/roadmap/merge/. Accessed: 2024-04-29.

[8] Ethereum. 2024. Proposer-Builder Separation. https://ethereum.org/en/roadmap/pbs/. Accessed: 2024-04-29.

[9] Flashbots. 2022. MEV-Boost. https://docs.flashbots.net/flashbots-mev-boost/. Accessed: 2024-04-29.

[10] Flashbots. 2024. Mempool Dumpster. https://github.com/flashbots/mempool-dumpster Accessed: 2024-04-08.

[11] Flashbots and Dune. 2024. Mempool Hygrometer. https://dune.com/flashbots/mempool-hygrometer Accessed: 2024-04-08.

[12] Tivas Gupta, Mallesh M. Pai, and Max Resnick. 2023. The Centralizing Effects of Private Order Flow on Proposer-Builder Separation. In *5th Conference on Advances in Financial Technologies (AFT 2023) (Leibniz International Proceedings in Informatics (LIPIcs), Vol. 282)*. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl, Germany, 20:1–20:15. https://doi.org/10.4230/LIPIcs.AFT.2023.20

[13] Quintus Kilbourn. 2022. Order flow, auctions and centralisation I - a warning. https://writings.flashbots.net/order-flow-auctions-and-centralisation. Accessed: 2024-04-23.

[14] Quintus Kilbourn. 2023. Order flow, auctions and centralisation II - order flow auctions. https://writings.flashbots.net/order-flow-auctions-and-centralisation-II. Accessed: 2024-04-23.

[15] Marc Lanctot et al. 2019. OpenSpiel: A Framework for Reinforcement Learning in Games. *CoRR* abs/1908.09453 (2019). arXiv:1908.09453 [cs.LG] http://arxiv.org/abs/1908.09453

[16] Angela Lu. 2024. Illuminating The Order Flow Landscape. https://writings.flashbots.net/illuminate-the-order-flow. Accessed: 2024-04-23.

[17] Michael Neuder. 2023. Bid cancellations considered harmful. https://ethresear.ch/t/bid-cancellations-considered-harmful/15500. Accessed: 2024-04-23.

[18] Alexandre Obadia, Alejo Salles, Lakshman Sankar, Tarun Chitra, Vaibhav Chellani, and Philip Daian. 2021. Unity is Strength: A Formalization of Cross-Domain Maximal Extractable Value. arXiv:2112.01472 [cs.CR]

[19] Shayegan Omidshafiei, Christos Papadimitriou, Georgios Piliouras, Karl Tuyls, Mark Rowland, Jean-Baptiste Lespiau, Wojciech M Czarnecki, Marc Lanctot, Julien Perolat, and Remi Munos. 2019. α-rank: Multi-agent evaluation by evolution. *Scientific reports* 9, 1 (2019), 9937.

[20] Burak Öz, Danning Sui, Thomas Thiery, and Florian Matthes. 2024. Who Wins Ethereum Block Building Auctions and Why? *arXiv preprint arXiv:2407.13931* (2024).

[21] Mallesh Pai and Max Resnick. 2023. Structural Advantages for Integrated Builders in MEV-Boost. arXiv:2311.09083 [econ.TH]

[22] 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)*. IEEE, 198–214.

[23] Caspar Schwarz-Schilling and Michael Nueder. 2023. Timing Games: Implications and Possible Mitigations. https://ethresear.ch/t/timing-games-implications-and-possible-mitigations/17612. Accessed: 2024-04-23.

[24] 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 preprint arXiv:2305.09032* (2023).

[25] Thomas Thiery. 2023. Empirical analysis of Builders' Behavioral Profiles (BBPs). https://ethresear.ch/t/empirical-analysis-of-builders-behavioral-profiles-bbps/16327. Accessed: 2024-04-23.

[26] Karl Tuyls et al. 2020. Bounds and dynamics for empirical game theoretic analysis. *Autonomous Agents and Multi-Agent Systems* 34 (2020), 1–30.

[27] Anton Wahrstätter. 2023. mevboost.pics. https://github.com/Nerolation/mevboost.pics. Accessed: 2024-04-20.

[28] Anton Wahrstätter, Liyi Zhou, Kaihua Qin, Davor Svetinovic, and Arthur Gervais. 2023. Time to Bribe: Measuring Block Construction Market. *arXiv preprint arXiv:2305.16468* (2023).

[29] William Walsh, Rajarshi Das, Gerald Tesauro, and Jeffrey Kephart. 2002. Analyzing Complex Strategic Interactions in Multi-Agent Systems. *AAAI-02 Workshop on Game-Theoretic and Decision-Theoretic Agents* (2002).

[30] Michael P. Wellman, Karl Tuyls, and Amy Greenwald. 2024. Empirical Game-Theoretic Analysis: A Survey. arXiv:2403.04018 [cs.GT]

[31] Fei Wu, Thomas Thiery, Stefanos Leonardos, and Carmine Ventre. 2024. Strategic Bidding Wars in On-chain Auctions. In *2024 IEEE International Conference on Blockchain and Cryptocurrency (ICBC)*. 503–511. https://doi.org/10.1109/ICBC59979.2024.10634354

[32] Sen Yang, Kartik Nayak, and Fan Zhang. 2024. Decentralization of Ethereum's Builder Market. arXiv:2405.01329 [cs.CR]