Leveraging Machine Learning for Bidding Strategies in Miner Extractable Value Auctions

Christoffer Raun*
ETH Zurich
Switzerland
christoffer.raun@inf.ethz.ch

Benjamin Estermann*
ETH Zurich
Switzerland
estermann@ethz.ch

Liyi Zhou
Imperial College London
United Kingdom
liyi.zhou@imperial.ac.uk

Kaihua Qin
Imperial College London
United Kingdom
kaihua.qin@imperial.ac.uk

Roger Wattenhofer
ETH Zurich
Switzerland
wattenhofer@ethz.ch

Arthur Gervais
University College London
United Kingdom
a.gervais@ucl.ac.uk

Ye Wang
University of Macau
China
wangye@um.edu.mo

ABSTRACT

The emergence of blockchain technologies as central components of financial frameworks has amplified the extraction of market inefficiencies, such as arbitrage, through Miner Extractable Value (MEV) from Decentralized Finance (DeFi) smart contracts. Exploiting these opportunities often requires fee payment to miners and validators, colloquially termed as bribes. The recent development of centralized MEV relayers has led to these payments shifting from the public transaction pool to private channels, with the objective of mitigating information leakage and curtailing execution risk. This transition instigates highly competitive first-price auctions for MEV. However, effective bidding strategies for these auctions remain unclear.

This paper examines the bidding behavior of MEV bots using Flashbots' private channels, shedding light on the opaque dynamics of these auctions. We gather and analyze transaction data for the entire operational period of Flashbots, providing an extensive view of the current Ethereum MEV extraction landscape. Additionally, we engineer machine learning models that forecast winning bids whilst increasing profitability, capitalizing on our comprehensive transaction data analysis. Given our unique status as an adaptive entity, the findings reveal that our machine learning models can secure victory in more than 50% of Flashbots auctions, consequently yielding superior returns in comparison to current bidding strategies in arbitrage MEV auctions. Furthermore, the study highlights the relative advantages of adaptive constant bidding strategies in sandwich MEV auctions.

1 INTRODUCTION

DeFi has become a significant catalyst for recent blockchain adoption, transitioning trading activities from centralized intermediaries like custodians, banks, and brokers to transparent and immutable on-chain smart contracts. DeFi offers a broad range of financial products, such as borrowing and lending [20], exchanges [4], and leveraged trading [1].

Just as in traditional finance, High Frequency Trading (HFT) opportunities such as arbitrage and liquidation have emerged within DeFi. This evolving DeFi landscape is primarily driven by the competition for MEV among market participants [9]. MEV represents the potential profit achievable by exploiting financial opportunities within a blockchain through transactions [25]. Because blockchain operations are entirely transparent, all transactions and smart contracts are publicly accessible [7]. Additionally, due to the asynchronous nature of a blockchain's Peer-to-Peer (P2P) network, there's an inherent time delay between the initiation (e.g., the signing of a transaction) and execution (i.e., the mining) of a transaction [9]. This allows HFT traders to monitor transactions on the P2P network and infer their competitors' actions and reactions, turning HFT within DeFi into a highly competitive game [38].

Miners, motivated by financial gain, prioritize transactions that offer the highest bribe per unit of computation (also known as a bid). Consequently, DeFi users engage "MEV auctions" to gain an advantage through front-running [9]. The growing demand for MEV has led to increased fees, inefficient transactions, and potential security risks, posing significant challenges to the DeFi ecosystem [25]. Centralized relay services like Flashbots [2] have established network peering

[*] Both authors contributed equally to the paper

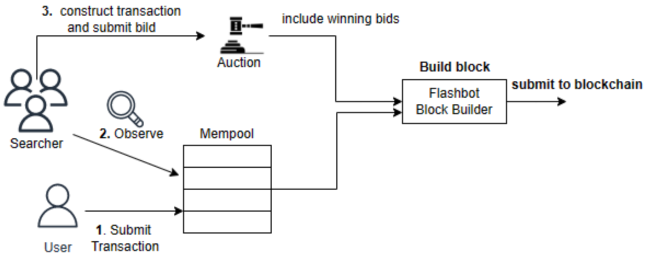

Figure 1: The diagram outlines the transaction submission and the MEV auction process, as well as the resultant MEV extraction.

agreements with miners, allowing DeFi users to directly send transactions to miners. Unlike P2P network broadcasting, MEV extractors using a relay service cannot observe their competitors' bids during the auction, leading to the adoption of sealed-bidding strategies (cf. Fig 1). Flashbots, the largest relay service, has secured cooperation with over 80% of Ethereum miners [7], which results in non-transparent markets of MEV auctions among all system participants. The current state of high-frequency bribery auction marketplaces remains somewhat opaque, and the distribution of MEV among different system stakeholders is not well understood. This lack of clarity hinders efforts to improve the blockchain system to better avoid the security challenges brought by the increasing MEV.

To gain a deeper understanding of the strategies employed by MEV extractors and to elucidate how their bidding strategies are influenced by market dynamics, this paper presents a comprehensive analysis of the bidding behaviors of MEV bots on Ethereum that submit transactions via Flashbots' private channels. Specifically, we closely examine the types of transactions submitted, the timing of these transactions, and the bribes offered to miners, also referred to as validators.

Moreover, we explore the feasibility of using machine learning models to predict winning bids in Flashbots auctions with the goal to increase MEV bot profits. These models consider a range of factors that impact bidding practices, including gas prices, timing, and transaction type. Through our machine learning model, we aim to illuminate the complex, rapidly evolving landscape of MEV auctions and provide actionable insights that can aid traders in more effectively and profitably navigating this space. Simultaneously, the interpretation of our model highlights the factors determining MEV generation and distribution among various stakeholders. This investigation thus provides valuable insights for market designers in the DeFi space for enhancing blockchain system stability. This paper makes the following key contributions:

(1) We conduct a thorough analysis of Ethereum MEV bots' bidding practices, particularly in relation to cyclic arbitrage and sandwich attacks, submitted via Flashbots' private channels. Our analysis spans the period from the inception of Flashbots until March 2023, effectively capturing the evolution of MEV auctions over time, which includes individual strategies and overarching trends.
(2) We develop machine learning models designed to predict winning bids in MEV auctions. We compare the performance of our machine learning models against historical data and baseline strategies to evaluate the models' effectiveness. The results show our models are able to win 50% of all arbitrage MEV auctions.
(3) We interpret the factors that determine winning strategies in MEV auctions using our machine learning models. Factors such as gas prices, timing, and transaction type are considered. Our results reveal that gas consumption and timing are the most significant factors. Moreover, we underline the limitations that linear regression models encounter when predicting the bribe ratio in sandwich MEV auctions. This finding aligns with the fact that most of the top bots implement adaptive, rather than constant, bidding strategies.

This study provides crucial insights into existing MEV bot bidding strategies in Flashbots auctions on Ethereum, while also exploring the potential of machine learning models within these auctions. Our investigation underscores the importance of understanding MEV bot behavior and introduces a framework for formulating effective bidding strategies, which could potentially mitigate the negative impacts of MEV extraction within the DeFi ecosystem.

2 BACKGROUND

2.1 Ethereum

Ethereum is a blockchain platform that enables users to conduct transactions and execute smart contracts without a central authority [12]. Initially operating under a Proof of Work (PoW) consensus mechanism, miners maintained the network by providing computational evidence of their contribution to the system. In return, they received block rewards and transaction fees in Ether (ETH). However, Ethereum has recently transitioned to a Proof of Stake (PoS) consensus algorithm, with the network maintained by validators that verify and append transactions to the blockchain.

A significant development in Ethereum's network architecture is the introduction of Proposer Builder Separation (PBS). PBS distributes the tasks of transaction bundling, sequencing, and block creation among different participants, thereby enhancing the network's decentralization. By dividing the roles of block builders and proposers, PBS aims to balance financial incentives and facilitate a fairer distribution of rewards within the network.

Smart contracts on the Ethereum platform are executed by the Ethereum Virtual Machine (EVM). Developers utilize this technology to build decentralized applications (DApps) and tokens, such as ERC-20 tokens, which can operate autonomously without intermediaries. Tokens are a type of virtual currency that can represent any asset or scarce item, from votes in a decentralized system to collectibles in a metaverse game.

Decentralized Exchanges (DEXs) represent a crucial application of Ethereum's smart contracts. These platforms enable peer-to-peer token trading, guided by straightforward rules defined in the smart contract, all without a central authority. One of the most prevalent DEX design paradigms is the Automated Market Maker (AMM), which employs a mathematical formula to determine the on-chain price for a pair of tokens. UniSwap v2 [14], UniSwap v3 [3], and SushiSwap (a UniSwap v2 fork) are among the most popular AMM DEX protocols currently employed on Ethereum.

2.2 MEV

MEV refers to the potential profit that can be extracted from transactions included within a blockchain block. In the current Ethereum ecosystem, numerous MEV bots scout and source these MEV opportunities. They identify lucrative transactions, package them, and submit them to the network.

MEV extraction involves the manipulation of transaction ordering, inclusion, and exclusion during the process of block creation, with the objective of acquiring profits exceeding the standard block rewards. The ordering of transactions is a critical aspect of MEV extraction and primarily involves two types: front-running and back-running. Front-running transpires when an MEV seeker raises the gas price to ensure their transaction is processed ahead of a targeted transaction. Conversely, back-running involves placing a transaction immediately after the target's transaction to capitalize on it. The deliberate ordering of transactions facilitates various forms of MEV extraction, such as sandwich attacks [32, 38] and cyclic arbitrage [31]. A sandwich attack involves positioning two transactions around a victim's transaction to extract value, while cyclic arbitrage capitalizes on price discrepancies across different DEXs through a series of token swaps.

MEV bots constantly scan the network for profitable MEV opportunities. The process of MEV extraction often leads to multiple extractors competing for the same opportunities, culminating in a Priority Gas Auction (PGA) [9, 25]. Such competition can result in high gas fees, network congestion, and increased block space usage. Additionally, MEV presents potential threats to the underlying consensus security of Ethereum, as demonstrated in recent research [36].

Flashbots. Flashbots [29] is a commercial entity that provides a centralized venue to perform a blinded gas price auction. Specifically, Flashbots auctions [28] provide a private communication channel between Ethereum users and validators for users to submit their transactions to an Ethereum block while promising, but not being able to prove privacy. The auction system started as a patch of the go-ethereum client (mev-geth) and has since been included in Proof of Stake (PoS) Ethereum with mev-boost. The auction utilizes a private transaction pool and a sealed-bid block space auction mechanism. The auction is a first-price sealed-bid auction, which allows MEV extractors to privately communicate their bid and transaction order preference without paying for failed bids. The auction mechanism tries to increase validator payoffs, while non-winning participants remain anonymous due to the sealed auction format. The auction also provides guarantees such as pre-trade privacy and failed trade privacy. Bribes can either be paid through direct payment to the validator or higher gas fees. In this paper, we aim to understand the current market of MEV extraction in the described Flashbots auctions and how much bribe is required to win an auction. There exist several research studies analyzing the popularity of Flashbots auctions [34]. However, the bidding strategies of MEV bots using Flashbots private channels to submit transactions have not yet been thoroughly explored. This paper aims to fill that gap by providing a comprehensive analysis of the bidding practices of MEV bots on Ethereum that submit transactions through Flashbots private channels. We aim to evaluate the feasibility of using machine learning models in Flashbots auctions to predict winning bids while optimizing for profitability.

3 DATA COLLECTION

In this study, we aim to analyze current bribery practices and train machine learning models using a comprehensive dataset of successful transactions submitted through the Flashbots project. To achieve this, we employ the publicly accessible Flashbots Blocks API, which offers access to all successfully submitted Flashbots transactions, as well as the corresponding bribes paid to validators. This information allows us to better understand the market dynamics of MEV in Flashbots auctions. However, the API only provides transaction hashes and bribe amounts, which is insufficient for a thorough examination of MEV, as the specific intent of each transaction remains unknown. As a result, we developed a custom transaction classification tool in Python to identify and classify various MEV attack types. This tool leverages block traces from an Ethereum archive node with Flashbots

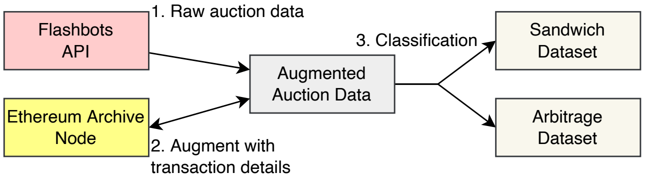

Figure 2: Schematic representation of the measurement setup, illustrating how data is collected from multiple sources.

data to categorize transactions. Through this classification process, we can detect specific MEV transactions, such as sandwich attacks and cyclic arbitrage, and associate potential revenue with bribes paid to validators. We illustrate the data collection process in Fig 2.

To enhance our understanding of bribe ratio fluctuations over time, we incrementally expanded the dataset. Ultimately, we gathered data on Ethereum transactions over a 200-day period, beginning with block number 14982026 (Jun-18-2022) and ending with block number 16326370 (Jan-03-2023). This range encompasses transactions before and after the Ethereum "Merge", enabling us to detect any additional influences during our analysis. After applying our custom classification tool to the collected data, we generated two distinct datasets—one for sandwich attacks and one for cyclic arbitrages. This separation was necessary to define and extract a unique set of features for each attack type, which is crucial for training machine learning models on the obtained data. Our classification tool adopts a heuristic approach, initially identifying all swaps and transfers executed within a block. We focused on the primary decentralized exchange (DEX) protocols available on Ethereum (UniSwap V2, UniSwap V3, Balancer, SushiSwap, and Curve) when detecting swaps. The information gathered from swaps and transfers is employed to apply distinct criteria to transactions, enabling the detection of cyclic arbitrage and sandwich attacks.

3.1 Cyclic Arbitrage

To recognize a transaction as a cyclic arbitrage trade, we established a set of criteria that must be satisfied [31]. Specifically, we consider transactions that involve multiple swaps executed within a single transaction. We define two distinct types of transactions that qualify as cyclic arbitrage.

First, for a transaction to be classified as an arbitrage trade, a sequence of swaps, denoted as $s_0, s_1, \ldots, s_n$, must form a loop. In this sequence, the input amount of swap s_i should match the output amount of the preceding swap s_{i-1}, and the input token of the first swap s_0 must be identical to the output token of the last swap s_n. The revenue opportunity is then calculated as the difference between the output amount s_n and the input amount s_0.

Second, an alternative cyclic arbitrage transaction comprises a sequence of swaps $s_0, s_1, \ldots, s_n$, where the input token and amount of swap s_0 equal the output amount and output token of s_n. Moreover, all but one pair of swaps should have matching input and output amounts. We regard the swap pair s_{j-1}, s_j, where the output/input condition does not hold, as the swap generating profit.

In our analysis, we focus on cyclic arbitrage trades with WETH as the profit token. This simplification eliminates the need for a price oracle to convert token prices to WETH at the time of the transaction. We also observed some Flashbots transactions exploiting multiple cyclic arbitrage opportunities simultaneously. As all these opportunities yielded WETH as profit, we consolidated them into one extensive cyclic arbitrage opportunity, representing it by aggregating the input and output amounts, summing the number of swaps, and concatenating the sets of involved protocols and tokens.

To evaluate whether a cyclic arbitrage trade is profitable, we compared the potential revenue opportunity against the costs of executing the trade and bribing the validator. Standard execution costs on the Ethereum network include the base gas fee consumed by the transaction, while bribing costs involve higher gas fees and direct ETH transfers to the validator. In our dataset, we exclusively consider profitable cyclic arbitrage opportunities, allowing us to focus on the most economically viable trades and better train our models. The dataset for the specified time window comprises 572,970 profitable cyclic arbitrage transactions.

3.2 Sandwich Attacks

A sandwich attack is a particular type of exploit where a malicious actor manipulates a transaction by placing it between two others. This differs from arbitrage transactions, which involve a single transaction containing multiple swaps. Sandwich attacks require analyzing multiple transactions to identify this attack type. In our study, we detected sandwich attacks based on the following criteria [32, 38]:

- Consider bundles that involve a "front-running" transaction T_F starting with WETH and a "back-running" transaction T_B ending with WETH.
- Both T_F and T_B should contain only one swap transaction each.
- Focus solely on sandwich attacks that involve exactly one "attacked" swap transaction T_A.

Additionally, any attacks classified as unprofitable were excluded from our dataset. To evaluate whether a sandwich attack is profitable, we compare the potential revenue opportunity against the cost of executing the trade and bribing the validator. This criteria enables us to concentrate on the most economically viable sandwich attacks and improve the training of our models.

We chose to exclude attacks targeting more than a single swap transaction to ensure a consistent feature set for model training. We discovered that this assumption led to the exclusion of a relatively small percentage of identified sandwich attacks. The final sandwich dataset contains 443,273 profitable sandwich attacks.

4 CYCLIC ARBITRAGE MEV AUCTIONS

Our first analysis focuses on the MEV from cyclic arbitrage. Our dataset, comprising 527,970 transactions, reveals a total revenue of 27,037.72 ETH. A comprehensive overview of the observed total revenue and per-transaction revenue is provided in Table 1. Notably, the median profit per transaction is approximately 0.00066 ETH, which equates to a mere 1.22 USD as of May 2023. This is a substantial decrease compared to the average profit reported by McLaughlin et al. [21] up until July 2022, where the median profit declined significantly from 0.002 ETH to 0.0006 ETH. This drastic drop underscores a more competitive arbitrage market, with cyclic arbitrage bots increasingly participating in opportunities that yield relatively modest gains. The average bribe ratio across the entire dataset is 79.18%, which is higher than the number reported by McLaughlin et al. [21]. This increase in bribe ratio might suggest a shift in the arbitrage strategies towards a more aggressive bidding behavior in an attempt to secure transaction priority, reflecting a heightened competitive landscape. Our findings thus highlight the evolving dynamics of the Ethereum arbitrage market, with implications for both individual actors and the broader ecosystem.

Figure 3 illustrates the progression of the average bribe ratio across a dataset of 1,000 Flashbots blocks throughout the studied period. Concurrently, Figure 4 demonstrates the quantity of active bots per 1,000 Flashbots blocks over the same duration. Generally, the average bribe ratio for arbitrage MEV auctions via Flashbots remains consistently above 60%, and there is a persistent presence of over 80 active bots seeking and exploiting arbitrage opportunities.

Moreover, these plots demonstrate the dynamic nature of the average bribe ratio, which exhibits fluctuations over time. Significantly, during the initial months following the introduction of Flashbots, the average bribe ratio swiftly rose from 70% to 80%. This trend persisted up to the point of the "Merge", with the average bribe ratio consistently surpassing 80%. This trend is indicative of the intense competition characteristic of these auctions, where numerous arbitrageurs contend for a limited quantity of lucrative opportunities. The continuous ascent in the average bribe ratio over time suggests a steady increase in the expected number of participants pinpointing arbitrage opportunities within these auctions.

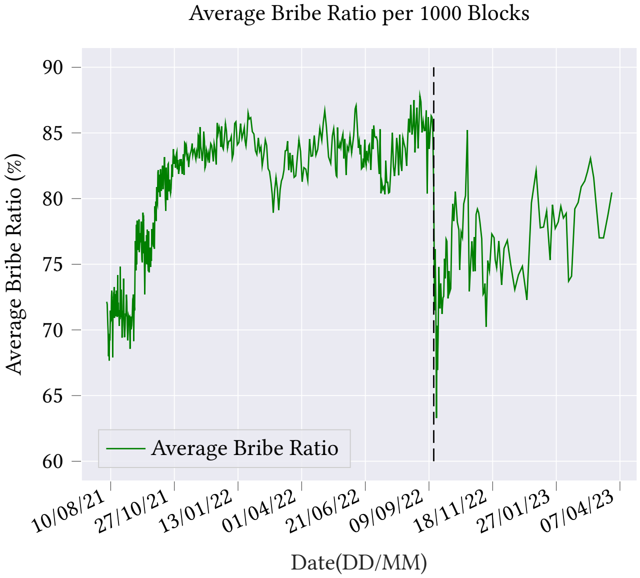

Figure 3: Cyclic Arbitrage: changes in the average bribe ratio per 1,000 Flashbots blocks across our dataset — the black dashed line shows the "Merge" (as is consistent in subsequent figures).

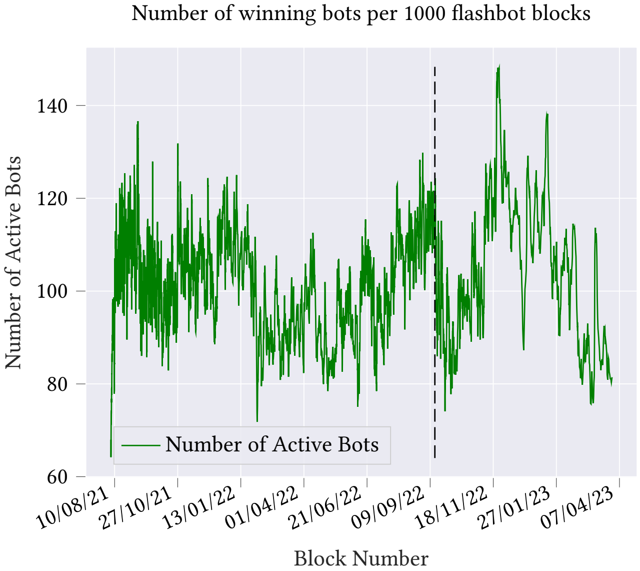

Figure 4: Cyclic Arbitrage: Active bots per 1,000 Flashbots blocks across our dataset.

The sharp decline in the average bribe ratio, marked by a dashed black line in Figure 3, corresponds to the time of the

Categories	Revenue	Base Fee	Bribe	Profit
Total	27037.72664	5987.86015	17614.24797	3435.61852
Mean	0.05121	0.01134	0.03336	0.00651
Median	0.01607	0.00834	0.00355	0.00066
Std Dev	1.22053	0.01125	1.15728	0.19188

Table 1: Statistic about Arbitrage Opportunities - information on arbitrage opportunities collected across all transactions, including both total values and values on a per-transaction basis, all measured in Ether.

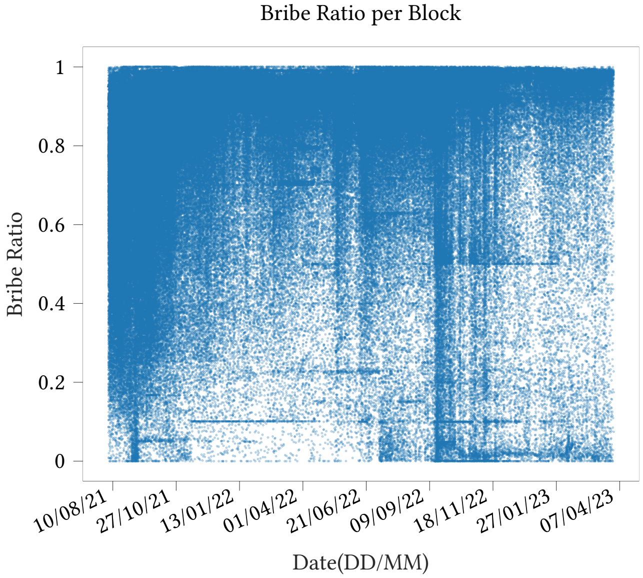

Figure 5: Cyclic Arbitrage: bribe ratio of transactions in our dataset — a blue dot represents the bribe ratio for one transaction.

“Merge”. At this juncture, the average bribe ratio witnessed a substantial decrease, plummeting from 85% to 75%, with increased fluctuations compared to the pre-Merge era. This period coincides with the introduction of Proposer/Builder Separation (PBS) through mev-boost[1], whereby the stakeholders responsible for selecting transactions to be included in blocks vary over time, alternating between builders and proposers. This inconsistency during the operation procedures may be identified as a potential contributing factor to the sharp decline and unstable bidding strategies in the average bribe ratio.

Additionally, we meticulously analyze the bribe ratio for each auction within our dataset over time, aiming to elucidate the evolution of bidding strategies at a granular level (cf. Figure 5). Our analysis reveals that the majority of auctions necessitate a bribe ratio approaching 100% in order to secure a win. However, after June 2022, more and more auctions result in a substantially lower winning bribe ratio, demonstrating the evolution of bribe strategies over time in arbitrage MEV auctions.

In the subsequent sections, we examine various factors that could potentially influence the bidding strategies employed in arbitrage auctions, with a particular focus on gauging the competitiveness of these auctions. To this end, our investigation encompasses both the correlation between profit and the bribe ratio of cyclic arbitrages, and the involvement of bots in the auctions.

4.1 Relationship of Bribe Ratio to Profit

Figure 6 demonstrates the relationship between the bribe ratio and the profit generated per transaction (measured in ETH), with distinct colors representing the bids made by different arbitrageurs. It's evident from the plot that bidding strategies in arbitrage MEV auctions are quite varied. The bribe ratio for auctions generating a profit of 10^{-3} ETH spans the entire bidding range. However, as the profit increases or decreases, the bribe ratio tends to converge towards 1. Interestingly, some of the highest profit opportunities have very low bribe ratios. This may indicate that not all competitors are aware of these opportunities and further in-depth analysis is needed.

Based on the graphical representation of the profit and bribe ratio, the strategies adopted by individual arbitrageurs can be categorized into several types. Further details about these categories will be discussed in the subsequent sections. For instance, bids plotted along a single horizontal line indicate that an arbitrageur consistently applies the same bribe ratio, regardless of the available profit opportunity. This suggests a risk-averse approach, prioritizing predictable costs over potential profits. Conversely, vertical lines reveal that certain arbitrageurs consistently aim for a fixed profit level, allocating the remaining amount to bribes. This strategy implies that as profit opportunities escalate, so too does the

[1] https://github.com/Flashbots/mev-boost

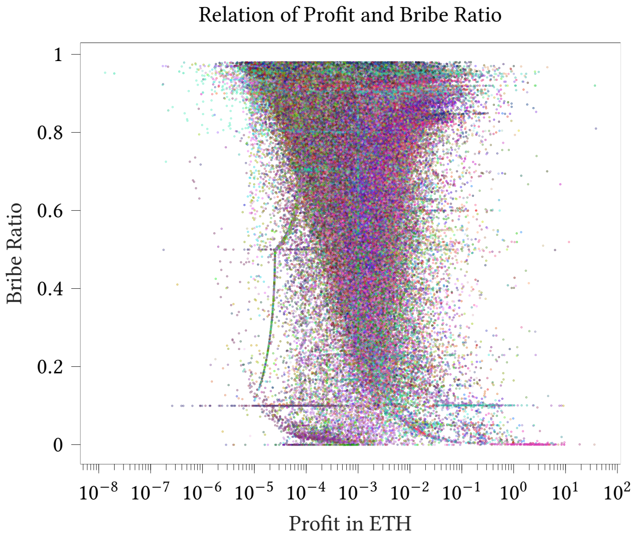

Figure 6: Relationship of bribe ratio to profit in ETH of all transactions. Each unique color corresponds to one specific arbitrageur.

bribe ratio, reflecting an opportunistic approach to capitalizing on lucrative arbitrage opportunities.

Our analysis also unearths more intricate strategies employed by particular arbitrageurs. We observe that some arbitrageurs modulate the bribe ratio following a specific function based on the anticipated profit. This approach represents a more nuanced strategy, indicating a level of sophistication that adjusts the bribe proportion in response to fluctuating profit scenarios, hinting at the dynamic and complex nature of the arbitrage market in the Ethereum ecosystem.

4.2 Examining Individual Aribitrage Bots

In our analysis, we identified 1128 unique Ethereum addresses (bots) that emerged victorious in Flashbots auctions. Interestingly, the top 10 winning bots accounted for 43.63% of all auction victories. Figure 7 illustrates the number of won arbitrage MEV auctions and the cumulative profit for each bot. We observed that as the number of successful auctions increases, the average profit per arbitrageur tends to converge. Unlike certain arbitrageurs who exploit a very few MEV opportunities for exceptionally high profits, the arbitrageurs who regularly participate in arbitrage MEV auctions tend to focus on opportunities with average profits. This implies that the majority of arbitrage MEV auctions are conducted with similar profit expectations. This demonstrates the competitive nature of these auctions, where a consistent performance may be valued and sustainable over occasional high returns.

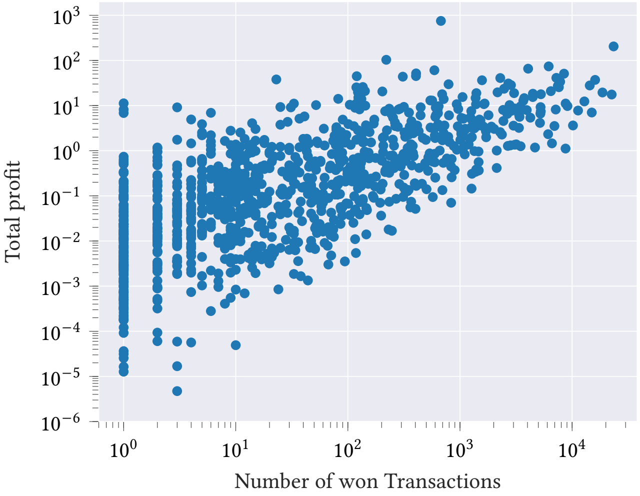

Figure 7: Relationship of the number of won transaction to the total profit per unique address.

We further illustrate the bribe ratio over time and the relationship between profit and bribe ratio for individual arbitrageurs in Appendix A.1. There are substantial differences in the strategies adopted by different arbitrageurs. For instance, Top 1 bot, who won the most number of MEV auctions, tends to employ a stable bidding strategy, which includes offering a fixed bribe ratio. This is represented by the horizontal line in the figure. Moreover, he manages to achieve a consistent profit despite fluctuations in revenue, as demonstrated by the vertical lines in the figure.

In contrast, Top 2 bot employed dynamic bidding strategies. Despite the varied nature of these strategies, they consistently yielded auction wins across the entire time frame. Notably, we observed that Top 2 bot started bidding with very low bribe ratios in the auctions recently, a significant deviation from his previous strategies which involved higher bids. This evidence demonstrates that individual arbitrageurs continuously adapt their bidding strategies over time to optimize profits in the auctions. This adaptability may prove to be a crucial factor in maintaining competitiveness in the MEV auction marketplace.

There are also arbitrageurs who do not appear to have a distinct bidding strategy. For instance, Top 3 and 4 bots have experimented with similar bidding practices during various time periods. Despite their disparate performances in winning MEV auctions, for auctions with profits ranging between 0.001 ETH and 0.01 ETH, their bribe ratios are relatively dispersed. Interestingly, their bids veer closer to 1 in MEV auctions with larger or smaller profits. This seemingly erratic behavior may indicate a lack of a specific strategy or could be a reflection of these arbitrageurs' efforts to experiment with a variety of approaches to identify the

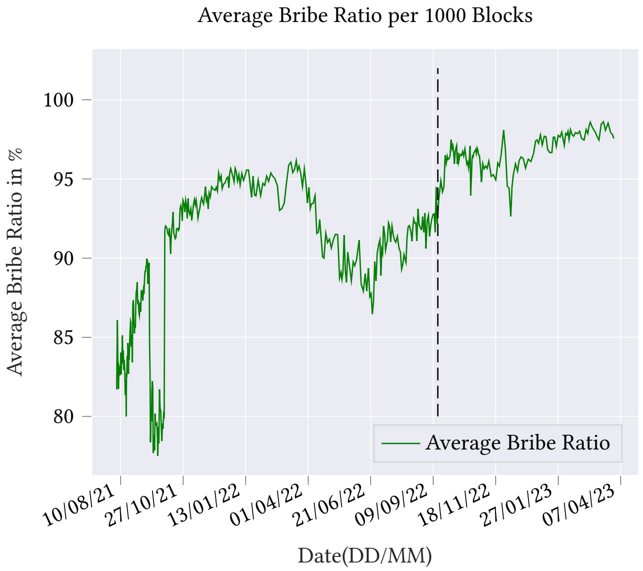

Figure 8: Sandwich attacks: average bribe ratio per 1,000 Flashbots blocks across our dataset.

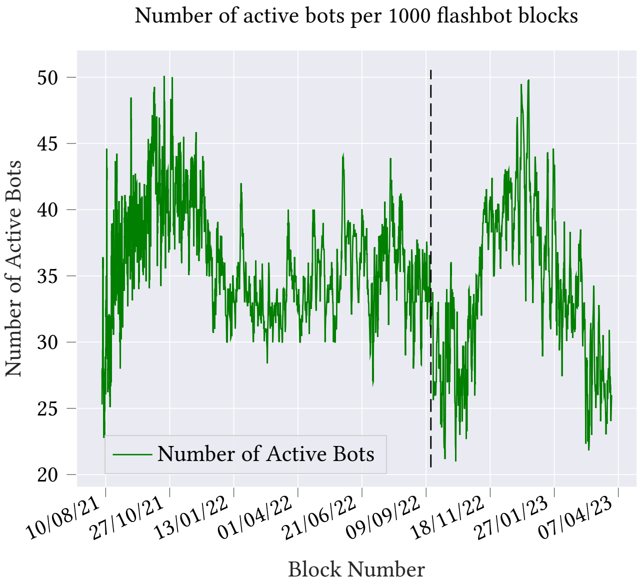

Figure 9: Sandwich attacks: active bots per 1,000 Flashbots blocks across our dataset.

most profitable one. It also underscores the wide range of strategies employed by participants in MEV auctions, further demonstrating the complexity and dynamism of this space.

5 SANDWICH ATTACKS MEV AUCTIONS

The second type of MEV strategy that we scrutinize in this research is the sandwich attack [32, 38]. In our dataset, we identified 443,273 successful sandwich attacks, executed by sandwich bots, resulting in a total profit of 4081.63 ETH. However, these attacks required substantial bribes totaling 34408.77 ETH. Detailed profit and bribe statistics can be seen in Table 2.

Interestingly, our dataset exhibits an average bribe ratio exceeding 91.41%, which is substantially higher than the bribe ratio observed in cyclic arbitrage. Despite the high bribe ratio, sandwich MEV strategies appear to generate significantly higher revenue for the attackers. Therefore, attackers still reap more profits from sandwich MEV auctions as compared to arbitrage MEV auctions. This suggests that sandwich attacks, despite their higher costs, can still be a lucrative strategy in the high-stakes, competitive landscape of MEV auctions.

Figure 8 traces the progression of the average bribe ratio for sandwich attacks over time, while Figure 9 illustrates the count of active bots per 1,000 Flashbots blocks during the same timeframe. Generally, the average bribe ratio for sandwich MEV auctions via Flashbots maintains a consistent level above 80%, with over 30 active bots persistently seeking and capitalizing on these opportunities. Although the bribe ratio consistently surpasses the winning bids in arbitrage MEV auctions, the number of active bots participating in these auctions is significantly lower.

Through an examination of the data, we can discern that the average bribe ratio required to win a sandwich auction can be categorized into three distinct phases. In the first phase, from the inception of Flashbots until April 2022, the bribe ratio steadily increases until it reaches 95%. In the second phase, it drops to 90% and persists at this level until the "Merge". Thereafter, in the final phase, the bribe ratio leaps back up to 95% and approaches 100%. This is different from what we observed in arbitrage MEV auctions where the average bribe ratio decreases after the “Merge”. The sandwich MEV auctions are generally involved more competitions among arbitrageurs as the average bribe ratio over time is strictly higher than arbitrage MEV auctions.

Figure 10 illustrates the bribe ratio on an auction-by-auction basis. Over time, we observe a declining trend in the number of opportunities. We hypothesize that this decrease is influenced by a range of factors. Firstly, major Ethereum DEXs and Aggregators such as 1Inch are improving measures to protect users against sandwich attacks. Secondly, changes in the macroeconomic environment have led to a substantial decline in trading volume on DEXs over the observed time window. On the other hand, the winning bids tend to converge to high bribe ratio close to one, which is also different from the trend with arbitrage MEV auctions where some winning bidders use a zero bribe ratio to miners. As we have

Categories	Revenue	Base Fee	Bribe	Profit
Total	42917.10427	4426.70603	34408.76744	4081.63080
Mean	0.09682	0.00999	0.07762	0.00921
Median	0.02901	0.00694	0.01684	0.00061
Std Dev	0.39174	0.01115	0.34394	0.08715

Table 2: Sandwich Opportunities in ETH, along with their revenue, bribe and profit.

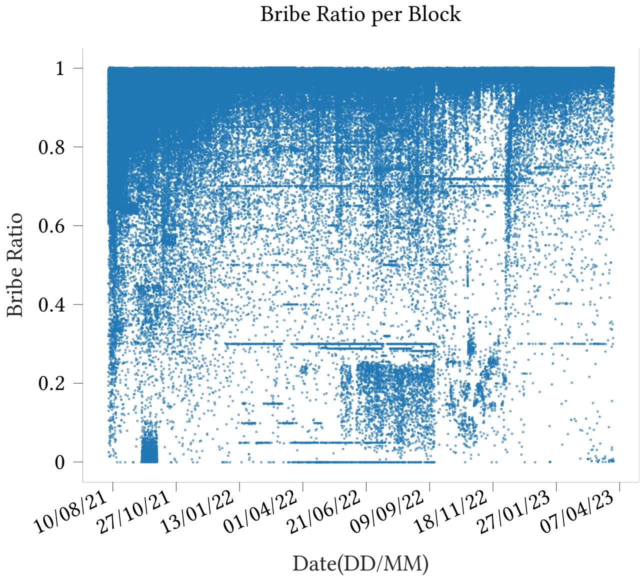

Figure 10: Sandwich attacks: bribe ratio of observed transactions across our dataset.

previously mentioned, we can only observe the winning auction bids, limiting our ability to determine the number of bots participating in an auction and the influence this may have on the bid outcome. Also, we lack of the information about the relationship between the winning attackers and the miners who generated the blocks containing the MEV transaction.

In the ensuing sections, we undertake an analysis of various factors that might shape the bidding strategies deployed in sandwich auctions, akin to the considerations for arbitrage transactions. Our investigation primarily zeroes in on both the correlation between profit and the bribe ratio of sandwich attacks, and the involvement of bots in the auctions.

5.1 Relationship of Bribe Ratio to Profit

We first explore the relationship between the bribe ratio and the profit derived from sandwich attacks, with the aim of identifying specific strategies employed by bots. Figure 12

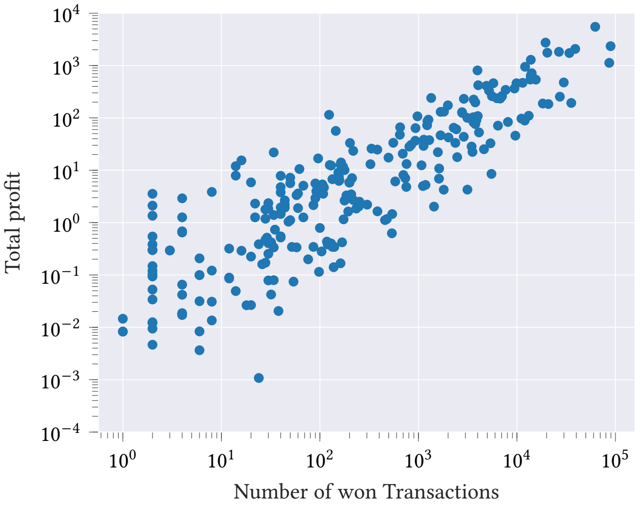

Figure 11: Relationship of the number of won transaction to the total profit per unique address.

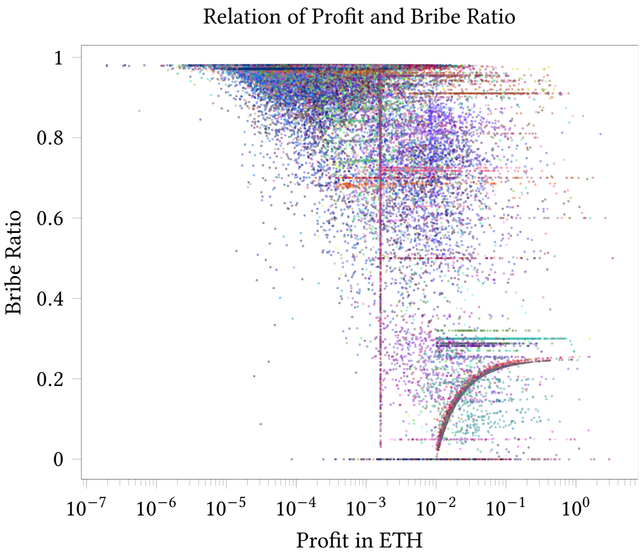

Figure 12: Sandwich attacks: the relationship between bribe ratio and profit in ETH.

provides a visual representation of these strategies, with each color denoting the bid associated with a distinct bot address.

Analogous to the cyclic arbitrage scenario, we endeavor to discern bot strategies through the analysis of bid patterns. Horizontal lines signify a bot consistently bidding a particular bribe ratio, while vertical lines indicate a bot targeting a specific profit regardless of the potential profit at stake. Compared to tactics observed in arbitrage MEV auctions, bidders' strategies in sandwich attacks exhibit several significant features. Firstly, the overall trend of the bribe ratio is inversely proportional to profit. More specifically, sandwich MEV auctions with lower returns typically result in a bribe ratio of 1, while those sandwich MEV auctions with higher returns, surpassing 0.001 ETH, tend to produce a significantly lower bribe ratio.

Secondly, winning strategies appear more dispersed in comparison to the arbitrage MEV auction. While there exists a multitude of different bribe ratios used to achieve the same profit in arbitrage MEV auctions, the strategies in sandwich MEV auctions appear more consolidated or aggregated. This implies a difference in the competitive landscape and tactics between the two types of MEV auctions, leading to different results and strategic patterns.

This discrepancy further emphasizes the complexity of the MEV landscape and the necessity for thorough and nuanced analysis to develop effective strategies for increasing profits while minimizing risks. As follows, we extend our investigation to consider individual bots, in order to better understand the nuances of their strategic approaches.

5.2 Examining Individual Sandwich Bots

Our dataset comprises 288 distinct sandwich bots, with the top 20 accounting for 68.35% of all successful auctions. Notably, most participating bots secured victories in only a few (1-10) auctions. The leading bot captured 10.07% of all auctions, the second-ranked bot secured 9.75%, and the third-ranked bot 7.03%. We illustrate the bribe ratio over time and the relationship between profit and bribe ratio for individual sandwich bots in Appendix A.2.

We delve into the relationship between the bribe ratio and profit, aiming to discern the strategies adopted by the top-performing bots. It is found that the top sandwich attackers implement more clear-cut and straightforward bidding strategies compared to the arbitrage exploiters. For instance, Top 1, 3, 6, 11, 12, and 14 bots – nearly half of the top bidders – adhered to constant bribe ratio strategies over time. This highlights that even simple bidding strategies can yield reasonably good performance in sandwich MEV auctions. The reason for this phenomenon could be the structure and dynamics of sandwich MEV auctions, which might lend themselves well to constant bribe ratio strategies. Nonetheless, this finding should be interpreted with caution as other factors such as market conditions, transaction timing, and the behavior of other bidders could also influence the effectiveness of different bidding strategies.

6 DYNAMIC BIDDING MODEL

In this section, we formulate machine learning models designed to predict the bribe ratio based on a specific set of features related to MEV transactions. Our objective is to design a model that can predict a winning bribe by effectively predicting the lowest winning bribe ratio. We focus on developing machine learning models for both cyclic arbitrage and sandwich attacks, requiring two separate models due to their distinct feature sets.

6.1 System Configuration

In constructing our prediction model, we consider only features that are available prior to the auction. This constraint is essential, as our ultimate goal is to develop a model capable of predicting bribes in real-time Flashbots auctions. Moreover, we examine the model's performance within a restricted context. Although our model may predict winning bribes, it cannot account for how other bots might alter their strategies based on our bids. Consequently, when evaluating our models, we assume a static environment in which bots maintain their existing strategies and do not adapt dynamically in response to changes.

6.2 Cyclic Arbitrage

6.2.1 Machine Learning Model. Before we optimize for profit, we need to model the bribe ratio. We have evaluated various machine learning models, ultimately finding that the best-performing model was an *LGBMRegressor.* Given a trained *LGBMRegressor*, we optimize for profit by adding a fixed offset to the output of the *LGBMRegressor.* We determine the optimal offset on the training dataset and leave it unchanged on the test dataset. Unfortunately, the underlying data distribution is skewed in the sense that a low number of auctions makes up a large part of the potential profit. For our training data, 0.1% of the auctions make up 30% of the total profit. For this reason, we deploy a separate strategy for all auctions with a profit further than three standard deviations away from the mean profit. We call this auctions *outliers.* For these auctions, a strategy with a constant bribe ratio turned out to work best. Consequently, we combine the two approaches to obtain the best outcome and call this approach the **hybrid model**. The prediction flow of our hybrid model proceeds as follows. First, we check is an auction is an outlier. If the auction is an outlier, bid using the constant bribe ratio. Otherwise, bid using the *LGBMRegressor* prediction plus the offset.

6.2.2 Features. The *LGBMRegressor* was trained on the following features: Block number, Potential profit, Revenue with base fee, Number of swaps involved, Start amount (arbitrage), End amount (arbitrage), Base gas cost, Gas required, Protocols involved (one-hot encoded), Token categories (one-hot encoded). These features were selected based on their availability and ease of computation when an arbitrage transaction is identified. We assume that the gas units required can be estimated with reasonable accuracy.

6.3 Arbitrage Evaluation

The model was trained on 80% of the collected data and evaluated on the remaining 20%. To develop a machine learning model that performs well on future auctions, we train the model on older data and use the most recent data from the dataset for prediction. This approach ensures that the model performs well over an unseen time horizon. We consider a static model that is not retrained while making predictions on the test data.

6.3.1 Metrics. We evaluated our models on the following two metrics: (i) Number of auctions won and (ii) Total profit from auctions. Optimizing solely for the number of auctions won, a model can easily win nearly all auctions by using a bribe ratio of 1. However, the profit in this case would be 0. Optimizing only for total profit might result in winning just one auction that yields a very high profit with a low bribe ratio. However, it is highly unlikely that this exact opportunity will recur. Thus, we aim to find a model that performs well in both metrics.

6.3.2 Baselines. To better understand the performance of our machine learning model, we compare it to three different baselines based on simple heuristics.

Baseline 1 We define a static heuristic that uses a fixed bribe ratio for each transaction. We compute the optimal bribe ratio on the training data and apply it to the test data.

Baseline 2 In the second baseline, we again define a static heuristic. However, instead of optimizing the bribe ratio for profit on the training data, we choose to bid the average bribe ratio of the training dataset, 79.94% of the available profit, in all arbitrage opportunities.

Baseline 3 As a third baseline, we consider a dynamic bribing strategy. We compute the average of the last 1,000 bribe ratios and use this to bid in a new auction. Essentially, this strategy employs the moving average to determine the next bid.

6.3.3 Hybrid Model. It used a positive offset of 0.02 on the output of the *LGBMRegressor* and a constant bribe ratio of 12.5% for the outlier auctions.

Type	Total Profit ETH	Auctions Won In %
Maximum	968.424	100%
Baseline 1	94.970	56.59%
Baseline 2	198.950	33.86%
Baseline 3	244.789	32.06%
Hybrid Model	500.733	50.21%

Table 3: Arbitrage model performance comparison.

6.3.4 Performance Overview. Table 3 provides an overview of the performance of various models, demonstrating the efficacy of our hybrid model in comparison to the baseline models. The hybrid model stands out in terms of total profit, securing over 50% of the available profit. Out of the total profit 389.990 ETH are contributed by *outlier* auctions, with 110.743 ETH coming from the large part of the remaining auctions. This considerable profit margin underlines the model's ability to predict and win potentially lucrative MEV auctions effectively. By forecasting the winning bids and adjusting to auction dynamics, the hybrid model enables resource allocation and a profit increase. On the other hand, it's interesting to observe the distinct performances of the baseline models. Baseline 3, which employs a dynamic bidding strategy, wins a similar amount of auctions as Baseline 2, following a static bidding strategy. Despite this similarity, Baseline 3 significantly surpasses Baseline 2 in terms of achieved profit. This suggests that dynamically adjusting bidding strategies in response to changing auction conditions can yield substantial dividends. In contrast, Baseline 1 clinches the highest percentage of auction wins. However, this victory ratio comes at a steep price of considerably diminished profit. It could be inferred that Baseline 1 possibly opts for a high frequency, low-profit strategy, trying to win as many auctions as possible, even if the profit margins for individual transactions are relatively low. While this strategy does result in more frequent wins, it appears to sacrifice the overall profitability, reinforcing the importance of strategic bidding in MEV auctions.

6.3.5 Feature Importance. Our machine learning model, specifically the *LGBMRegressor*, allows us to assess the importance of various features when predicting the bribe ratio in arbitrage MEV auctions, as depicted in Figure 13.

Among the most significant findings is the pivotal role played by potential profit. It appears that the projected profit from a given arbitrage MEV auction directly influences the bribe ratio: a higher potential profit is likely associated with a higher bribe ratio. This finding aligns with the intuitive understanding that extractors are willing to spend more to secure potentially lucrative transactions.

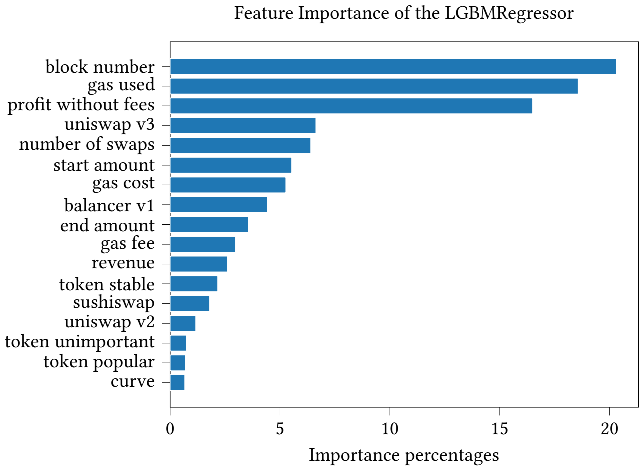

Figure 13: Cyclic arbitrage feature importance.

In contrast, the type of token involved in the transaction seems to have a minimal impact on the bribe ratio, according to our model. This suggests that the specific assets being arbitraged may not greatly influence the bidding strategy in the Flashbots auction, though this doesn't necessarily negate the importance of token type in other aspects of MEV extraction.

Interestingly, the block number emerges as the most influential feature in our model. This leads us to infer that the dynamics of the Flashbots auctions have evolved significantly over time. Factors such as changing market conditions, technological developments, and shifts in the competitive landscape may all contribute to this temporal variability.

Gas used, representing the complexity of the arbitrage implementation, also emerges as an influential feature. This could reflect the fact that more complex arbitrage operations require more computational resources (gas) to execute, thereby affecting the cost-benefit analysis and the ensuing bribe ratio.

Lastly, while other features may not play prominent roles in our machine learning model, it's crucial to note that their importance shouldn't be entirely dismissed. Though these features might not universally influence the bribe ratio, they could still have significant impacts under specific circumstances or scenarios. This emphasizes the multifaceted and context-dependent nature of MEV auctions and the strategies employed by participating bots.

6.4 Sandwich Attacks

For sandwich attacks, we employ a similar setup as for cyclic arbitrage. Although we train on a different dataset (i.e., using a different feature set), we predict the bribe ratio in the same way. Additionally, we found that the hybrid model used for cyclic arbitrage also performs reasonably well for sandwich attacks.

6.4.1 Machine Learning Model. As mentioned, the model we use to predict the bribe ratio for sandwich attacks is identical to the one employed for cyclic arbitrage. Consequently, the prediction flow and other aspects remain the same in both cases. We utilize the following features for sandwich attacks: Block number, Potential profit, Base gas fee, Frontrun gas used, Backrun gas used, Frontrun swap in amount, Frontrun swap out amount, Sandwiched swap in amount, Sandwiched swap out amount, Sandwiched transaction gas used, Backrun swap in amount, Backrun swap out amount, Protocols involved (one-hot encoded), Token Type (one-hot encoded).

6.5 Sandwich Evaluation

We assess the performance of our model in comparison to fundamental baselines. Since we are employing a similar setup, we utilize the same baselines as in the cyclic arbitrage case, but we evaluate them on the sandwich dataset. In this instance, we again allocate 80% of the available data for training and 20% for testing. Similarly, we detect outliers further than 3 standard deviations away from the mean profit of the training data.

6.5.1 Baselines (analogous to cyclic arbitrage). The heuristics behind establishing baselines for sandwich attacks follow a similar approach as those used for cyclic arbitrage. In the case of Baseline 2, we employ a constant average bribe ratio derived from the training data. Remarkably, this ratio stands at an elevated 90.05%.

6.5.2 Hybrid Model. It used a positive offset of 0.02 on the output of the *LGBMRegressor* and a constant bribe ratio of 72.5% for the outlier auctions.

6.5.3 Performance Overview. Table 4 presents a comparative evaluation of the different baselines. Interestingly, Baseline 1 secures the highest total profit, despite winning a meager 2.66% of auctions. This low winning percentage hints at a possible high volatility in the profit profile of Baseline 1. It also underscores the pronounced impact that a small number of high-profit opportunities can exert on the overall auction landscape. Contrastingly, our hybrid model, while realizing a slightly lower profit compared to Baseline 1, demonstrates a far superior capability in securing more auction wins. The dip in profit could potentially be attributed to some high-profit outliers in the test dataset that our model criteria may not have accurately identified.

There exists a subtle yet significant trade-off between securing more auction wins and achieving higher overall profits in the context of MEV auctions. At first glance, it might seem that winning more auctions would directly translate

Type	Total Profit ETH	Auctions Won In %
Maximum	209.585	100%
Baseline 1	62.624	2.66%
Baseline 2	28.608	5.59%
Baseline 3	17.272	18.46%
Hybrid Model	55.440	31.01%

Table 4: Sandwich model performance comparison.

to accruing more profits. However, on closer inspection, this isn't always the case. MEV opportunities, particularly those with high profit potential, are not uniformly distributed nor easily discoverable. These opportunities are sporadic and often require keen monitoring and swift action to capitalize upon. In the case of Baseline 1, despite winning only about 3% of auctions, it yielded the highest total profit. However, the sporadic nature of high-profit opportunities presents a substantial risk. Relying solely on infrequent, high-profit opportunities can lead to instability in profit generation and potential losses during periods of scarcity. Furthermore, sandwich opportunities, which often offer high profits, present an additional challenge. Given their lucrative nature, miners may choose to exploit these opportunities themselves instead of auctioning them off. This adds another layer of complexity to the already competitive landscape and could potentially decrease the chances of MEV bots securing these opportunities. Therefore, a well-rounded strategy would be to balance between pursuing high-profit opportunities and ensuring a consistent win rate in auctions. While high-profit opportunities offer substantial gains, their sporadic occurrence necessitates a strategy that also targets more frequent, albeit lower profit, opportunities. This way, bots can maintain a steady profit stream while still capitalizing on high-profit opportunities when they arise. A carefully calibrated blend of these strategies can lead to sustained profitability in the ever-evolving and dynamic landscape of MEV auctions.

6.5.4 *Key Challenges.*

The nature of sandwich auctions presents distinct challenges when constructing a robust machine learning model for MEV auctions. These challenges, primarily, stem from the heightened competitiveness and the dynamic transaction-profit relationship observed in sandwich auctions.

Competitive Landscape: Sandwich auctions have a higher average bribe ratio compared to arbitrage auctions, indicating a more competitive landscape. Bots are willing to pay a higher percentage of potential profit as a bribe to secure the transaction. This competition increases the unpredictability of the bribe ratio and makes it more difficult for a machine learning model to accurately forecast. Moreover, the intense competition might also lead to bots employing more complex, adaptive strategies, adding an extra layer of complexity to the prediction task.

Transaction-Profit Correlation: As Figure 11 illustrates, there's a stronger correlation between the number of won transactions and total profit in sandwich auctions compared to arbitrage auctions. This implies that securing more transactions generally leads to higher profits in sandwich auctions. While this could simplify profit prediction, it also suggests that there might be less room to exploit other bots' strategies for improved performance. If the majority of bots are playing optimally or near-optimally, a machine learning model may struggle to find exploitable inefficiencies.

Identifying High-Value Outliers: Identifying transactions with unusually high profit potential – the outliers – is a significant challenge. These opportunities don't occur frequently, and when they do, they're often quickly seized by bots using a constant high bribe ratio strategy. This could lead to these high-value opportunities being underrepresented in the data used to train the machine learning model. The model, therefore, may struggle to recognize and accurately predict the bribe ratio for these lucrative opportunities when they do arise.

7 RELATED WORKS

MEV MEV is the value miners can extract from their privileged position in the blockchain ecosystem, mainly through practices like front-running. Notable research on MEV and front-running includes Eskandari et al. [11], who offer a systematic overview of front-running attacks in blockchain systems, and Daian et al. [8], who introduces the concept of "MEV". Works that concentrate on quantifying MEV extraction include Qin et al. [25], who devise a framework for measuring MEV extraction, and Torres et al. [30], who suggest a methodology for assessing the volume and distribution of front-running activity. Several studies also focus on DeFi lending and borrowing protocols [5, 6, 10, 16–19, 22, 24, 33] and decentralized exchanges [8, 15, 25, 26, 31, 32, 37, 38], examining aspects such as economics, security, and formal modeling. Note that, MEV extraction can negatively impact user experience and, more significantly, the blockchain's underlying incentive structure, harming blockchain security [25, 36].

Proposer Builder Seperation (PBS), Private Transactions and Relayers Most blockchain transactions are propagated using the public P2P network. However, Ethereum experienced a significant shift with "the merge" in September 2022, transitioning from a PoW to a PoS leader election algorithm and introducing the proposer builder separation (PBS). PBS separates the tasks of creating new blocks and appending blocks to the blockchain [25]. In PBS, the role of a “validator” (previously "miner" in PoW) is split between distinct entities: “block builders” and “block proposers” (the validators themselves). Relays are also introduced to mediate and establish trust between block builders and proposers. Research on MEV extraction in private pools like Flashbots features studies by Weintraub et al. [35], Piet et al. [23], and Capponi et al. [7]. Furthermore, Zhou et al. investigated the use of relayers in DeFi attacks [39].

Machine Learning based Bidding While the application of machine learning techniques to MEV extraction is new, machine learning has been successfully employed in other domains to optimize bidding strategies [13, 27].

8 CONCLUSION

This paper analyses sandwich attack and cyclic arbitrage MEV strategies. Using a custom collection tool, we gather a dataset to gain insights into the current state of MEV bribing practices. Our analysis indicates that the average bribe ratio changes over time and that the available profit does not have a significant impact on the bribe ratio. We find that the bribe ratio for sandwich attacks is significantly higher than that of cyclic arbitrage opportunities. Building on these insights, we develop machine learning models to participate in Flashbots auctions. Our models were able to win more than 50% of the available auctions on the test data and outperformed our baselines, yielding higher profits. These results highlight the potential of machine learning in the MEV space and its ability to extract value.

REFERENCES

[1] [n.d.]. Alpha Homora V2. Available at: https://alphafinancelab.gitbook.io/alpha-homora-v2/.

[2] [n.d.]. Flashbots. https://github.com/flashbots.

[3] Hayden Adams, Noah Zinsmeister, Dan Robinson, Moody Salem, and River Keefer. 2020. Uniswap v3 Core. https://uniswap.org/whitepaper.pdf/. [Online; accessed 8-March-2023].

[4] Hayden Adams, Noah Zinsmeister, Moody Salem, River Keefer, and Dan Robinson. 2021. Uniswap v3 Core Whitepaper.

[5] Massimo Bartoletti, James Chiang, Tommi Junttila, Alberto Lluch Lafuente, Massimiliano Mirelli, and Andrea Vandin. 2022. Formal analysis of lending pools in decentralized finance. In *Leveraging Applications of Formal Methods, Verification and Validation. Adaptation and Learning: 11th International Symposium, ISoLA 2022, Rhodes, Greece, October 22–30, 2022, Proceedings, Part III*. Springer, 335–355.

[6] Massimo Bartoletti, James Hsin-yu Chiang, and Alberto Lluch Lafuente. 2021. SoK: lending pools in decentralized finance. In *International Conference on Financial Cryptography and Data Security*. Springer, 553–578.

[7] Agostino Capponi, Ruizhe Jia, and Ye Wang. 2022. The evolution of blockchain: from lit to dark. *arXiv preprint arXiv:2202.05779* (2022).

[8] Philip Daian, Steven Goldfeder, Tyler Kell, Yunqi Li, Xueyuan Zhao, Iddo Bentov, Lorenz Breidenbach, and Ari Juels. 2019. Flash boys 2.0: Frontrunning, transaction reordering, and consensus instability in decentralized exchanges. *arXiv preprint arXiv:1904.05234* (2019).

[9] Philip Daian, Steven Goldfeder, Tyler Kell, Yunqi Li, Xueyuan Zhao, Iddo Bentov, Lorenz Breidenbach, and Ari Juels. 2020. Flash boys 2.0: Frontrunning in decentralized exchanges, miner extractable value, and consensus instability. In *2020 IEEE Symposium on Security and Privacy (SP)*. IEEE, 910–927.

[10] Michael Darlin, Nikolaos Papadis, and Leandros Tassiulas. 2020. Optimal Bidding Strategy for Maker Auctions. *arXiv preprint arXiv:2009.07086* (2020).

[11] Shayan Eskandari, Seyedehmahsa Moosavi, and Jeremy Clark. 2020. Sok: Transparent dishonesty: front-running attacks on blockchain. In *Financial Cryptography and Data Security: FC 2019 International Workshops, VOTING and WTSC, St. Kitts, St. Kitts and Nevis, February 18–22, 2019, Revised Selected Papers 23*. Springer, 170–189.

[12] Ethereum Foundation. 2023. Ethereum Whitepaper. https://ethereum.org/en/whitepaper/. [Online; accessed 8-March-2023].

[13] Etan A Green and E Barry Plunkett. 2022. The science of the deal: Optimal bargaining on ebay using deep reinforcement learning. In *Proceedings of the 23rd ACM Conference on Economics and Computation*. 1–27.

[14] Dan Robinson Hayden Adams, Noah Zinsmeister. 2020. Uniswap v2 Core. https://uniswap.org/whitepaper.pdf/. [Online; accessed 8-March-2023].

[15] Lioba Heimbach, Ye Wang, and Roger Wattenhofer. 2021. Behavior of liquidity providers in decentralized exchanges. *arXiv preprint arXiv:2105.13822* (2021).

[16] Hsien-Tang Kao, Tarun Chitra, Rei Chiang, and John Morrow. 2020. An analysis of the market risk to participants in the compound protocol. In *Third International Symposium on Foundations and Applications of Blockchains*.

[17] Ariah Klages-Mundt, Dominik Harz, Lewis Gudgeon, Jun-You Liu, and Andreea Minca. 2020. Stablecoins 2.0: Economic foundations and risk-based models. In *Proceedings of the 2nd ACM Conference on Advances in Financial Technologies*. 59–79.

[18] Ariah Klages-Mundt and Andreea Minca. 2019. (In) Stability for the Blockchain: Deleveraging Spirals and Stablecoin Attacks. *arXiv preprint arXiv:1906.02152* (2019).

[19] Ariah Klages-Mundt and Andreea Minca. 2022. While stability lasts: A stochastic model of noncustodial stablecoins. *Mathematical Finance* (2022).

[20] Robert Leshner and Geoffrey Hayes. 2019. Compound Finance Whitepaper.

[21] Robert McLaughlin, Christopher Kruegel, and Giovanni Vigna. 2023. A Large Scale Study of the Ethereum Arbitrage Ecosystem. In *32th USENIX Security Symposium (USENIX Security 23)*.

[22] Daniel Perez, Sam M Werner, Jiahua Xu, and Benjamin Livshits. 2021. Liquidations: DeFi on a Knife-edge. In *International Conference on Financial Cryptography and Data Security*. Springer, 457–476.

[23] Julien Piet, Jaiden Fairoze, and Nicholas Weaver. 2022. Extracting godl [sic] from the salt mines: Ethereum miners extracting value. *arXiv preprint arXiv:2203.15930* (2022).

[24] Kaihua Qin, Liyi Zhou, Pablo Gamito, Philipp Jovanovic, and Arthur Gervais. 2021. An empirical study of defi liquidations: Incentives, risks, and instabilities. In *Proceedings of the 21st ACM Internet Measurement Conference*. 336–350.

[25] 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.

[26] Kaihua Qin, Liyi Zhou, Benjamin Livshits, and Arthur Gervais. 2021. Attacking the defi ecosystem with flash loans for fun and profit. In *International Conference on Financial Cryptography and Data Security*. Springer, 3–32.

[27] Kan Ren, Weinan Zhang, Ke Chang, Yifei Rong, Yong Yu, and Jun Wang. 2017. Bidding machine: Learning to bid for directly optimizing profits in display advertising. *IEEE Transactions on Knowledge and*

Data Engineering 30, 4 (2017), 645–659.

[28] Flashbot Team. [n.d.]. Flashbot Auction Overview. https://docs.flashbots.net/flashbots-auction/overview/. [Online; accessed 8-March-2023].

[29] Flashbot Team. [n.d.]. Welcome to Flashbots. https://docs.flashbots.net/. [Online; accessed 9-March-2023].

[30] Christof Ferreira Torres, Ramiro Camino, et al. 2021. Frontrunner jones and the raiders of the dark forest: An empirical study of frontrunning on the ethereum blockchain. In *30th USENIX Security Symposium (USENIX Security 21)*. 1343–1359.

[31] Ye Wang, Yan Chen, Haotian Wu, Liyi Zhou, Shuiguang Deng, and Roger Wattenhofer. 2022. Cyclic arbitrage in decentralized exchanges. In *Companion Proceedings of the Web Conference 2022*. 12–19.

[32] Ye Wang, Patrick Zuest, Yaxing Yao, Zhicong Lu, and Roger Wattenhofer. 2022. Impact and User Perception of Sandwich Attacks in the DeFi Ecosystem. In *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems*. 1–15.

[33] Zhipeng Wang, Kaihua Qin, Duc Vu Minh, and Arthur Gervais. 2022. Speculative multipliers on defi: Quantifying on-chain leverage risks. In *Financial Cryptography and Data Security: 26th International Conference, FC 2022, Grenada, May 2–6, 2022, Revised Selected Papers*. Springer, 38–56.

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

[35] Ben Weintraub, Christof Ferreira Torres, Cristina Nita-Rotaru, and Radu State. 2022. A flash (bot) in the pan: measuring maximal extractable value in private pools. In *Proceedings of the 22nd ACM Internet Measurement Conference*. 458–471.

[36] Liyi Zhou, Kaihua Qin, Antoine Cully, Benjamin Livshits, and Arthur Gervais. 2021. On the just-in-time discovery of profit-generating transactions in defi protocols. In *2021 IEEE Symposium on Security and Privacy (SP)*. IEEE, 919–936.

[37] Liyi Zhou, Kaihua Qin, and Arthur Gervais. 2021. A2mm: Mitigating frontrunning, transaction reordering and consensus instability in decentralized exchanges. *arXiv preprint arXiv:2106.07371* (2021).

[38] Liyi Zhou, Kaihua Qin, Christof Ferreira Torres, Duc V Le, and Arthur Gervais. 2021. High-frequency trading on decentralized on-chain exchanges. In *2021 IEEE Symposium on Security and Privacy (SP)*. IEEE, 428–445.

[39] Liyi Zhou, Xihan Xiong, Jens Ernstberger, Stefanos Chaliasos, Zhipeng Wang, Ye Wang, Kaihua Qin, Roger Wattenhofer, Dawn Song, and Arthur Gervais. 2022. SoK: Decentralized Finance (DeFi) Attacks. *arXiv preprint arXiv:2208.13035* (2022).

A TOP PERFORMING BOTS

A.1 Top Arbitrage Bots

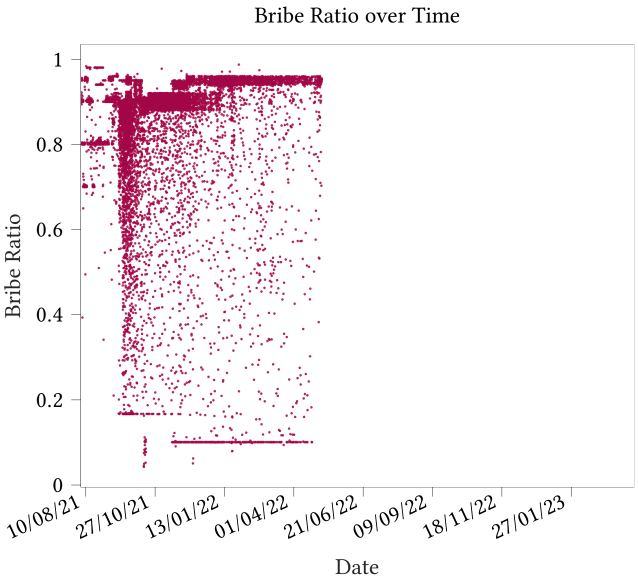

Figure 14: Bribe ratio over time for the Top 1 bot for cyclic arbitrage.

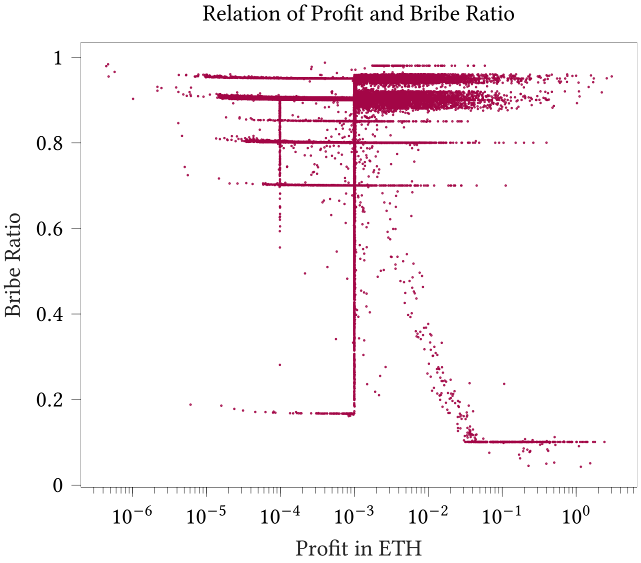

Figure 15: Bribe ratio compared to profit for the Top 1 bot for cyclic arbitrage.

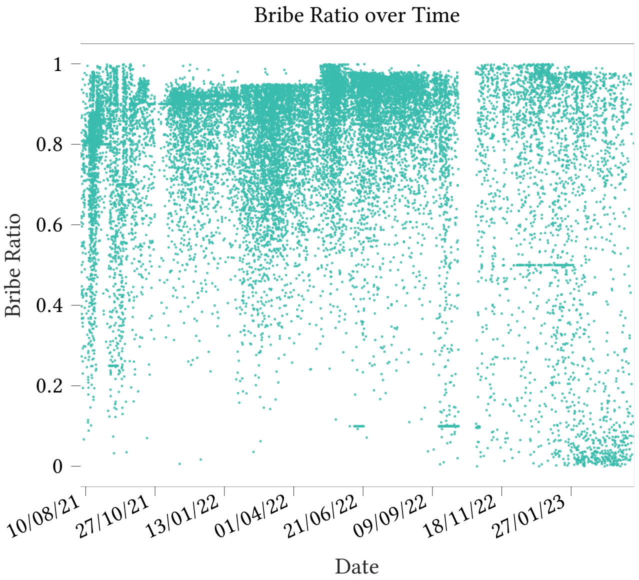

Figure 16: Bribe ratio over time for the Top 2 bot for cyclic arbitrage.

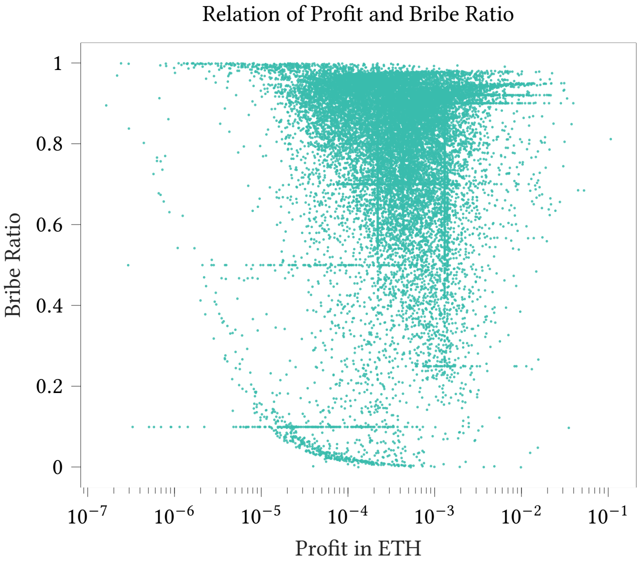

Figure 17: Bribe ratio compared to profit for the Top 2 bot for cyclic arbitrage.

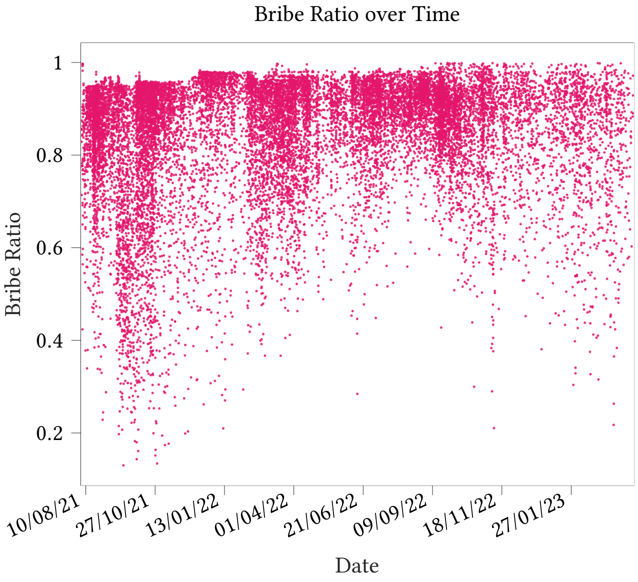

Figure 18: Bribe ratio over time for the Top 3 bot for cyclic arbitrage.

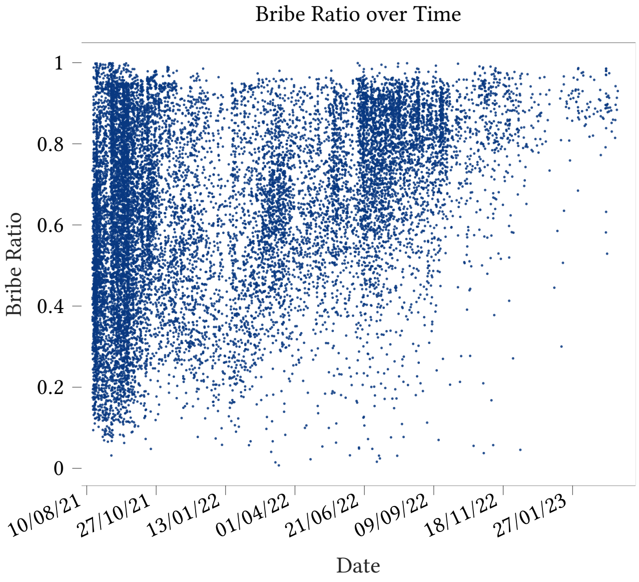

Figure 20: Bribe ratio over time for the Top 4 bot for cyclic arbitrage.

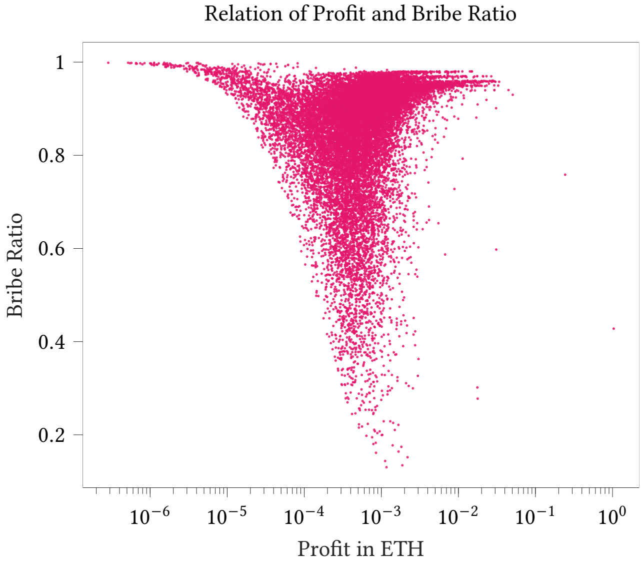

Figure 19: Bribe ratio compared to profit for the Top 3 bot for cyclic arbitrage.

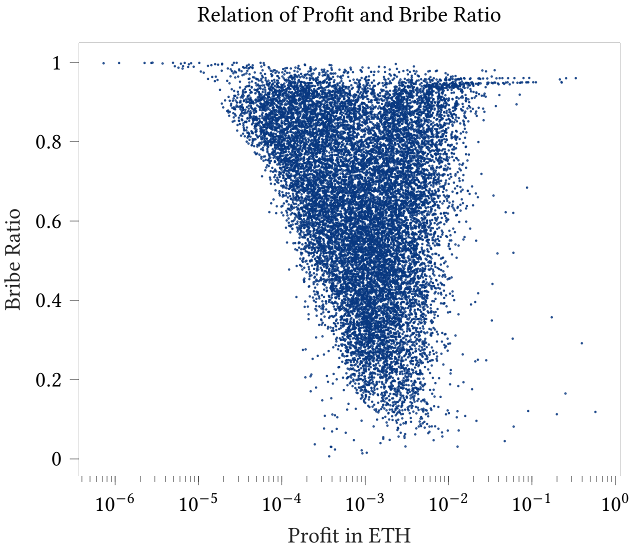

Figure 21: Bribe ratio compared to profit for the Top 4 bot for cyclic arbitrage.

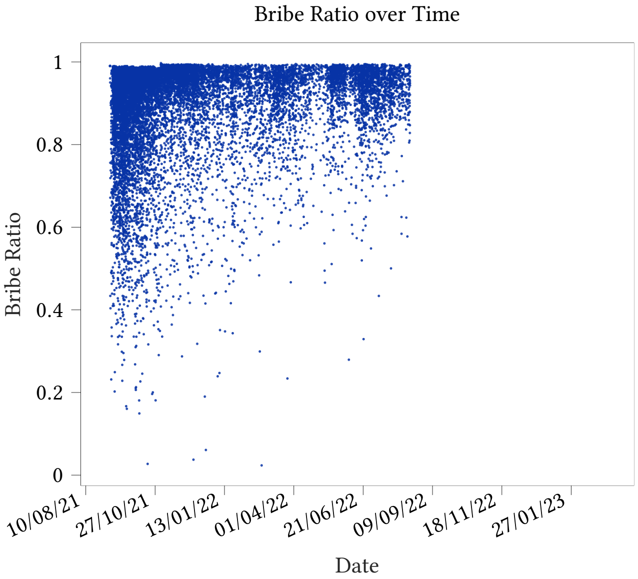

Figure 22: Bribe ratio over time for the Top 5 bot for cyclic arbitrage.

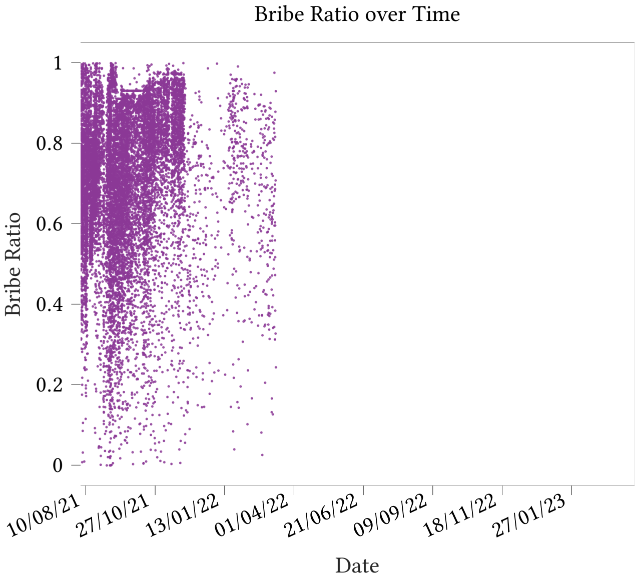

Figure 24: Bribe ratio over time for the Top 6 bot for cyclic arbitrage.

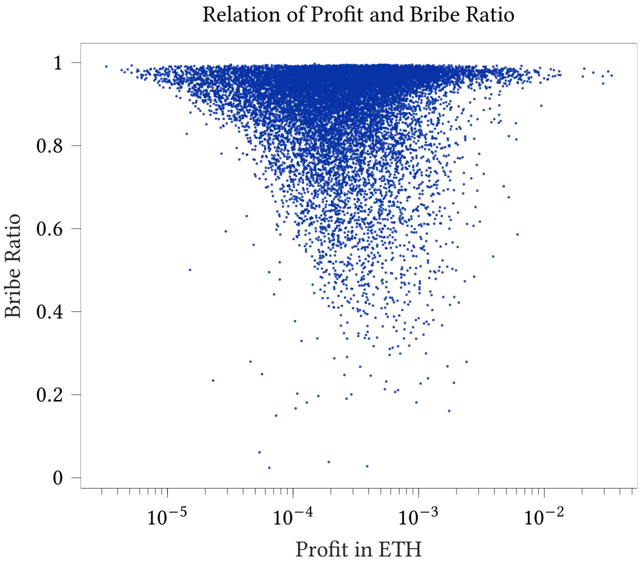

Figure 23: Bribe ratio compared to profit for the Top 5 bot for cyclic arbitrage.

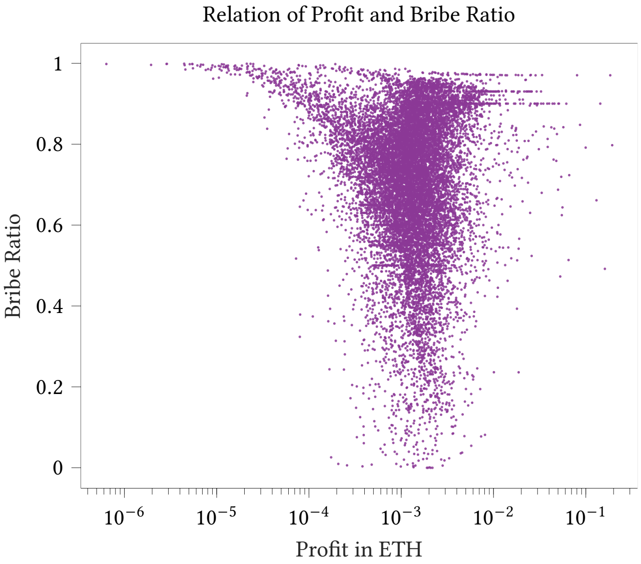

Figure 25: Bribe ratio compared to profit for the Top 6 bot for cyclic arbitrage.

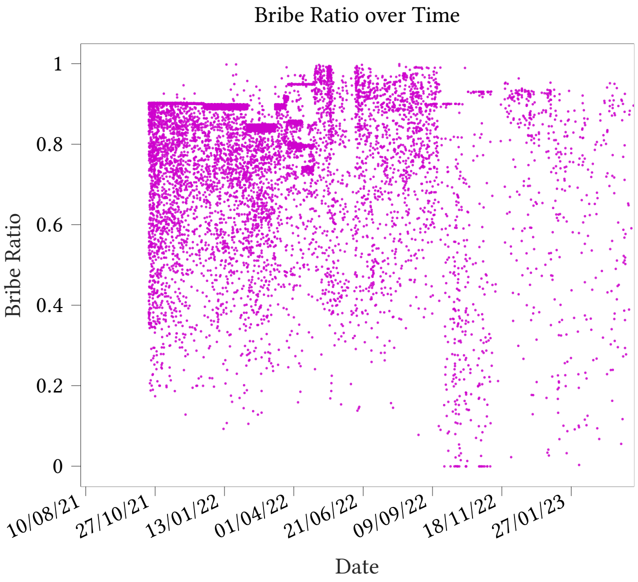

Figure 26: Bribe ratio over time for the Top 7 bot for cyclic arbitrage.

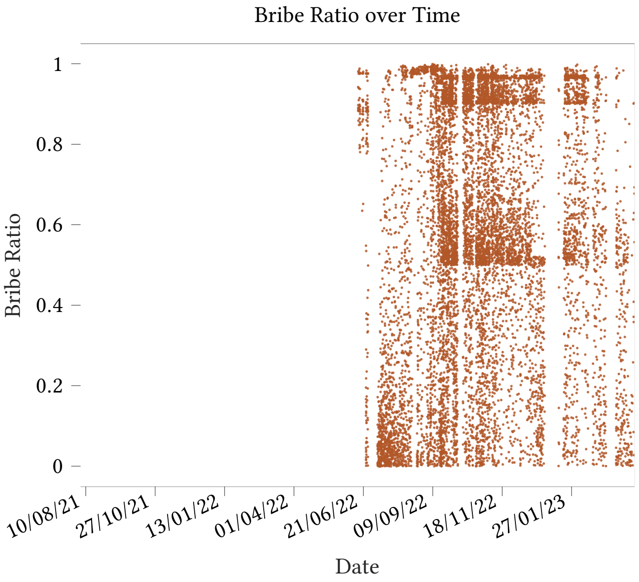

Figure 28: Bribe ratio over time for the Top 8 bot for cyclic arbitrage.

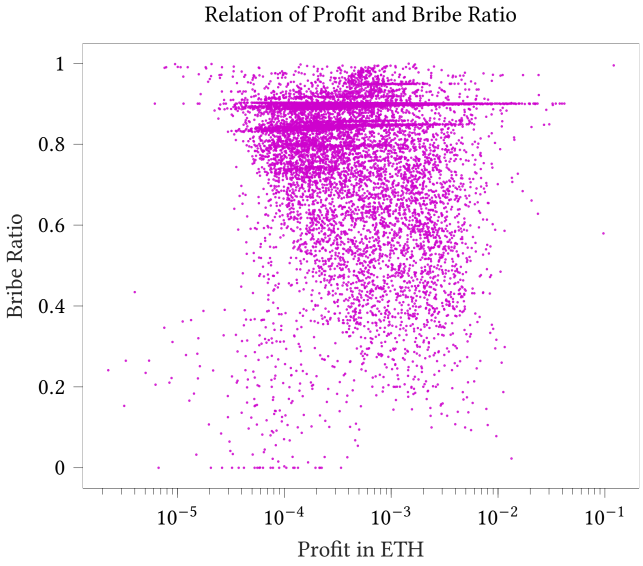

Figure 27: Bribe ratio compared to profit for the Top 7 bot for cyclic arbitrage.

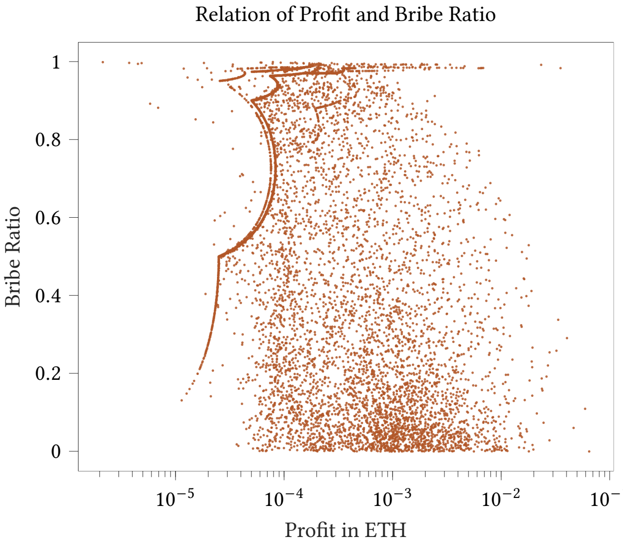

Figure 29: Bribe ratio compared to profit for the Top 8 bot for cyclic arbitrage.

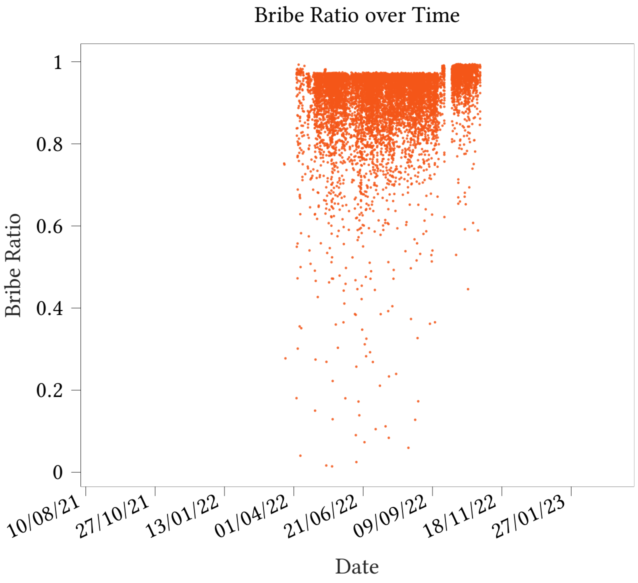

Figure 30: Bribe ratio over time for the Top 9 bot for cyclic arbitrage.

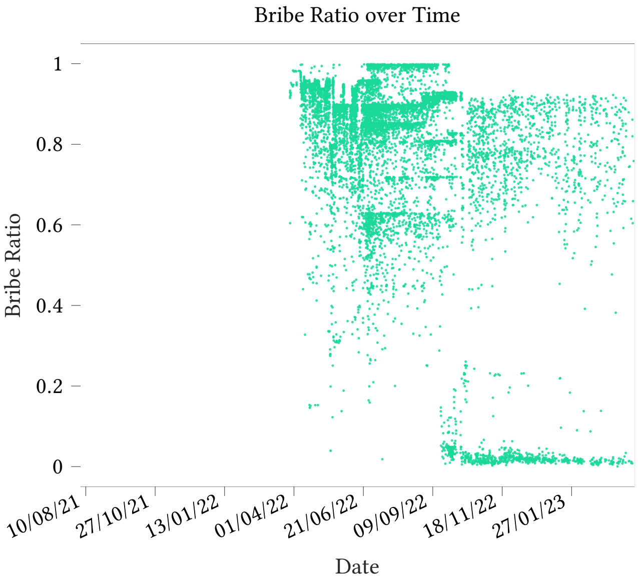

Figure 32: Bribe ratio over time for the Top 10 bot for cyclic arbitrage.

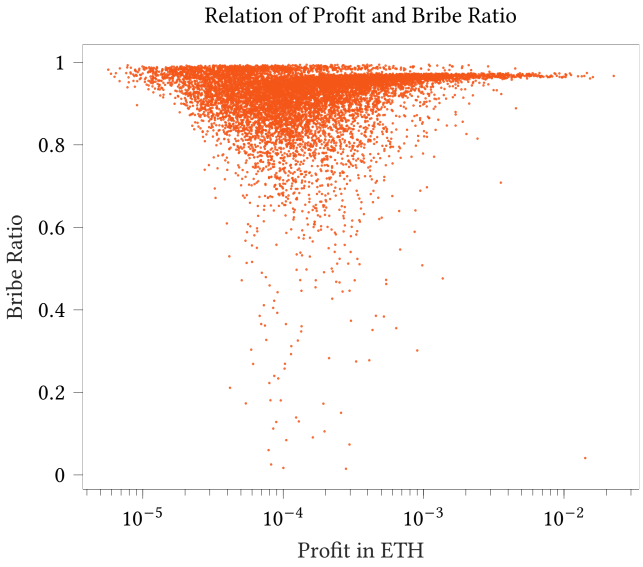

Figure 31: Bribe ratio compared to profit for the Top 9 bot for cyclic arbitrage.

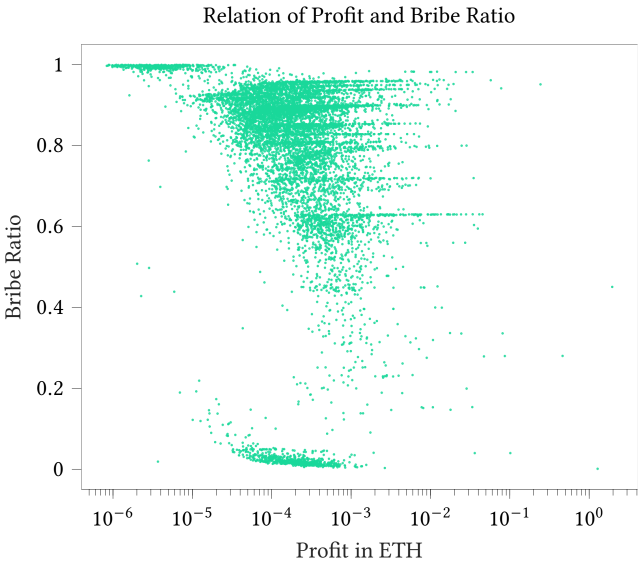

Figure 33: Bribe ratio compared to profit for the Top 10 bot for cyclic arbitrage.

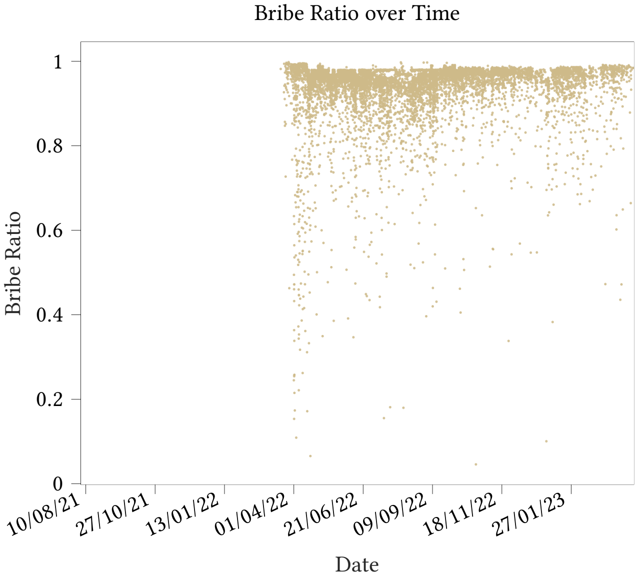

Figure 34: Bribe ratio over time for the Top 11 bot for cyclic arbitrage.

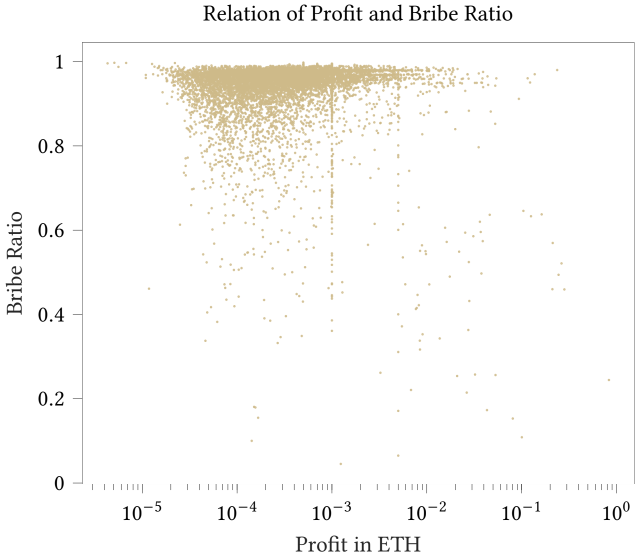

Figure 35: Bribe ratio compared to profit for the Top 11 bot for cyclic arbitrage.

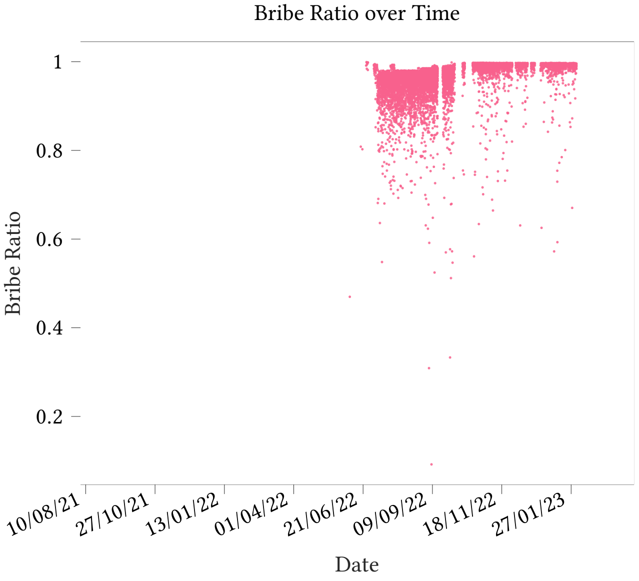

Figure 36: Bribe ratio over time for the Top 12 bot for cyclic arbitrage.

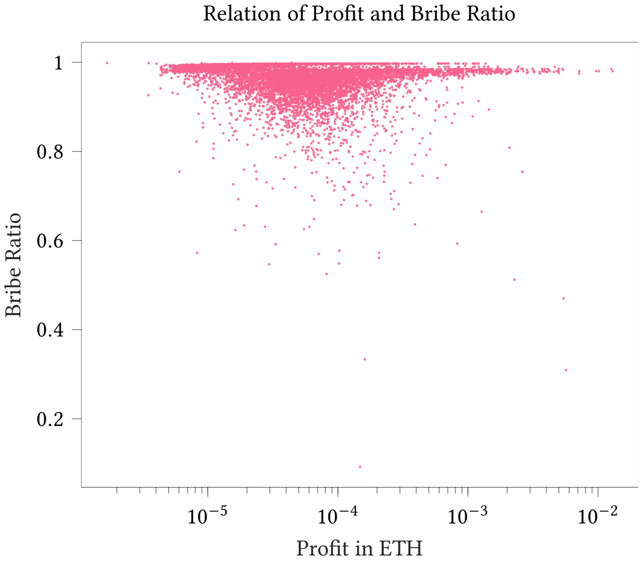

Figure 37: Bribe ratio compared to profit for the Top 12 bot for cyclic arbitrage.

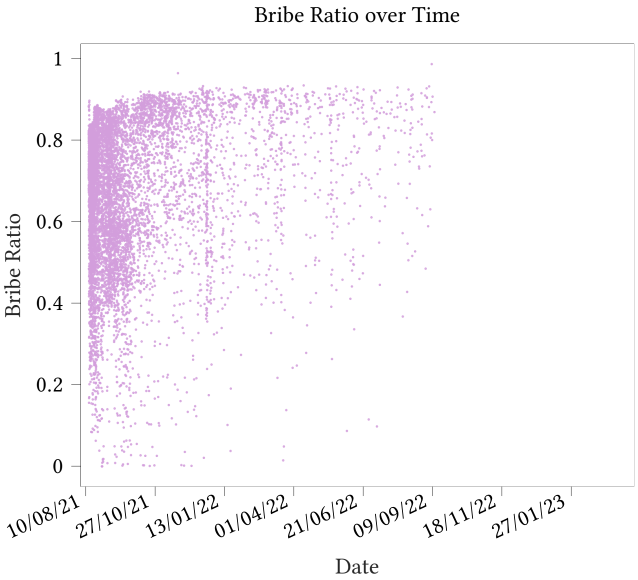

Figure 38: Bribe ratio over time for the Top 13 bot for cyclic arbitrage.

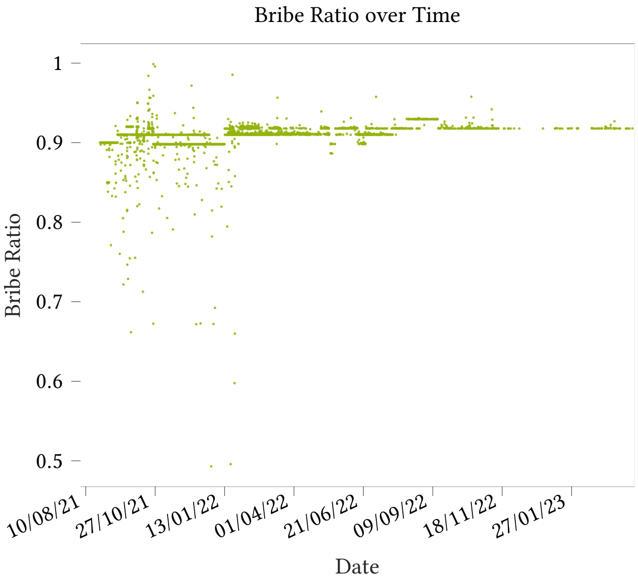

Figure 40: Bribe ratio over time for the Top 14 bot for cyclic arbitrage.

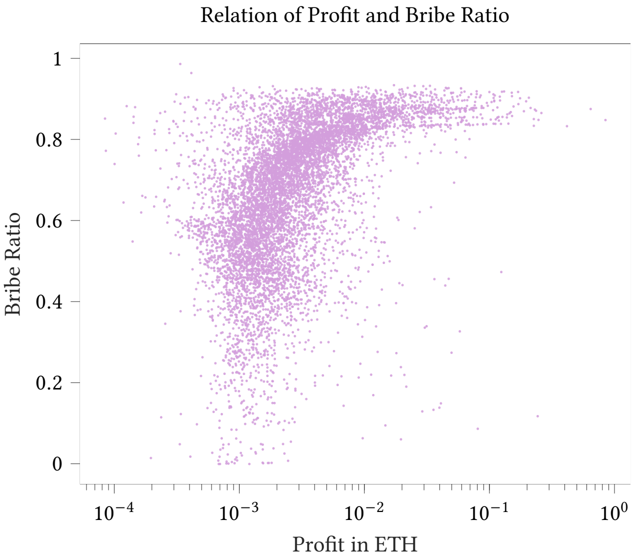

Figure 39: Bribe ratio compared to profit for the Top 13 bot for cyclic arbitrage.

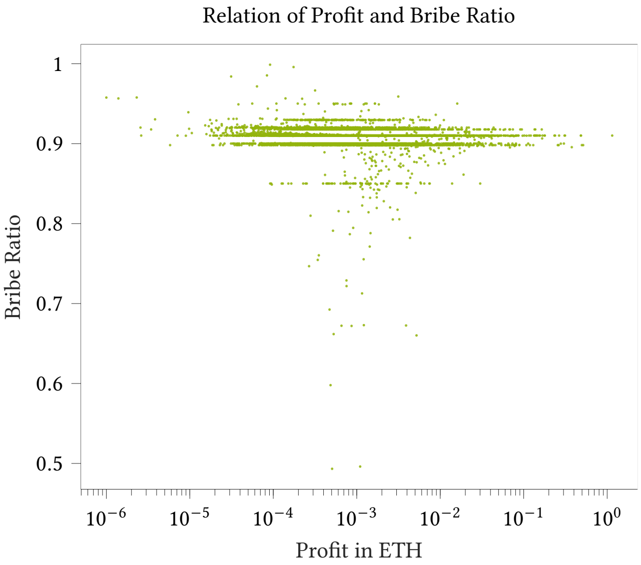

Figure 41: Bribe ratio compared to profit for the Top 14 bot for cyclic arbitrage.

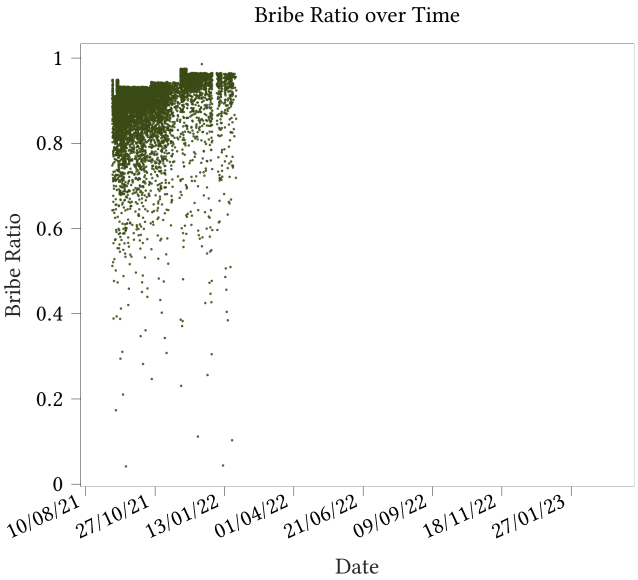

Figure 42: Bribe ratio over time for the Top 15 bot for cyclic arbitrage.

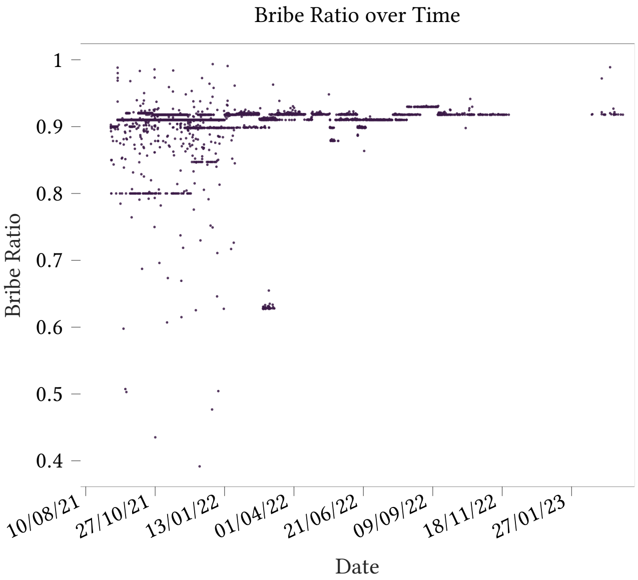

Figure 44: Bribe ratio over time for the Top 16 bot for cyclic arbitrage.

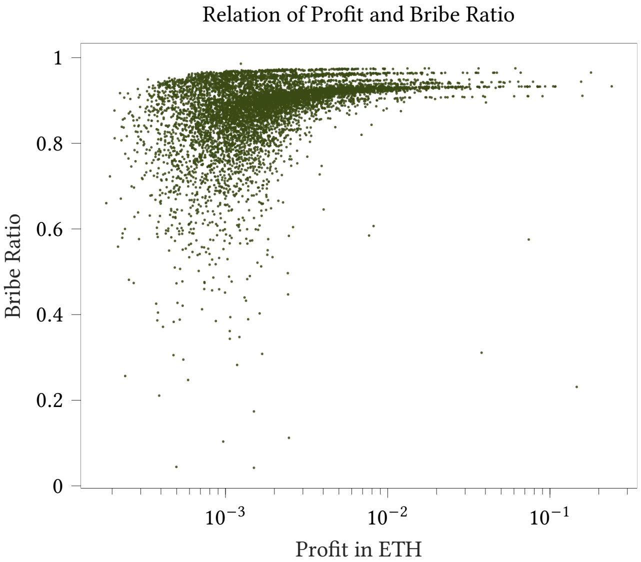

Figure 43: Bribe ratio compared to profit for the Top 15 bot for cyclic arbitrage.

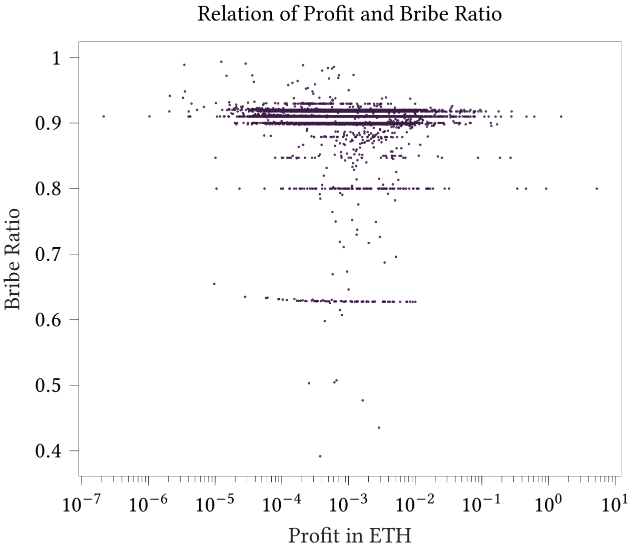

Figure 45: Bribe ratio compared to profit for the Top 16 bot for cyclic arbitrage.

A.2 Top Sandwich Bots

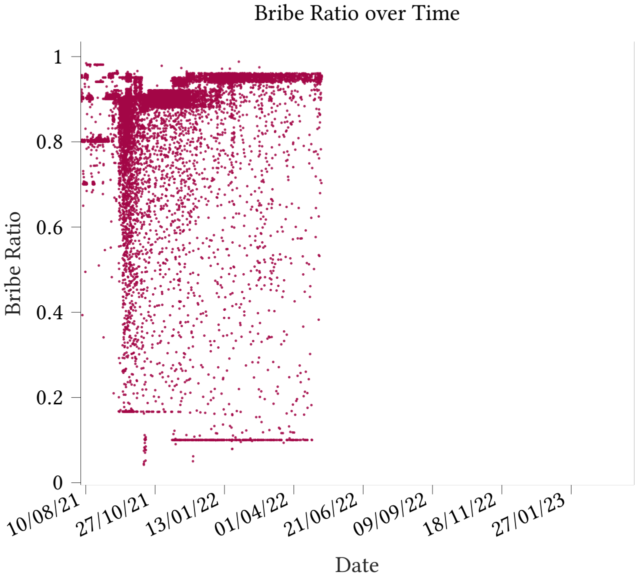

Figure 46: Bribe ratio over time for the Top 1 bot for sandwich attacks.

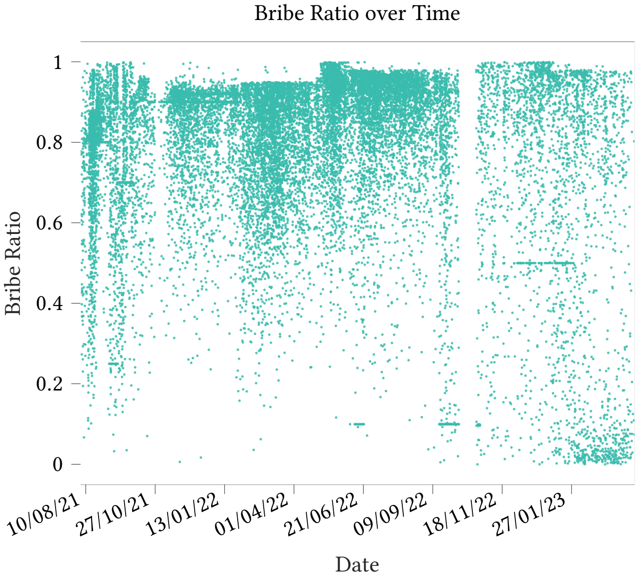

Figure 48: Bribe ratio over time for the Top 2 bot for sandwich attacks.

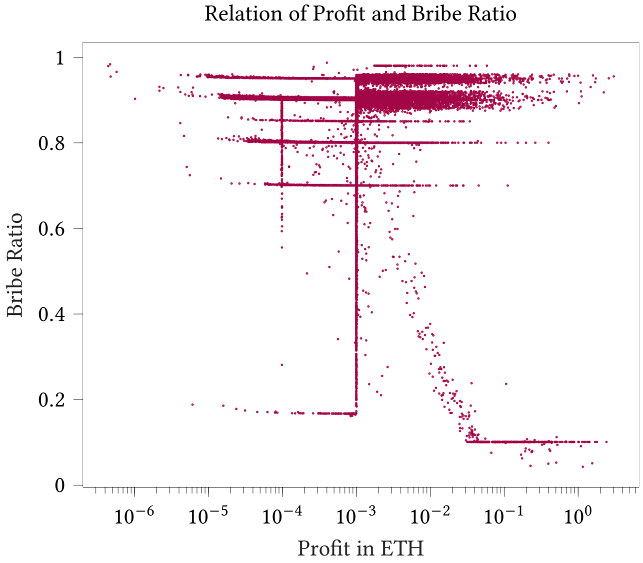

Figure 47: Bribe ratio compared to profit for the Top 1 bot for sandwich attacks.

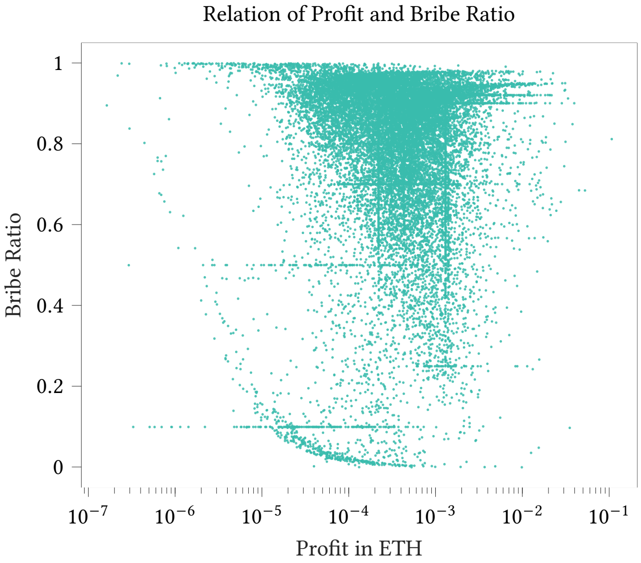

Figure 49: Bribe ratio compared to profit for the Top 2 bot for sandwich attacks.

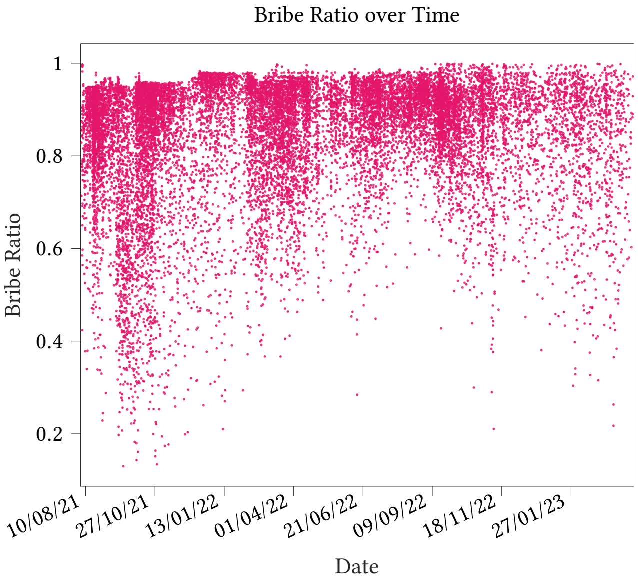

Figure 50: Bribe ratio over time for the Top 3 bot for sandwich attacks.

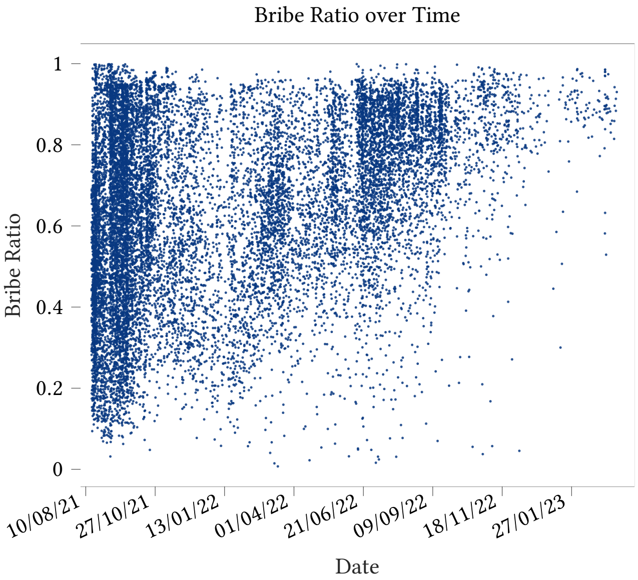

Figure 52: Bribe ratio over time for the Top 4 bot for sandwich attacks.

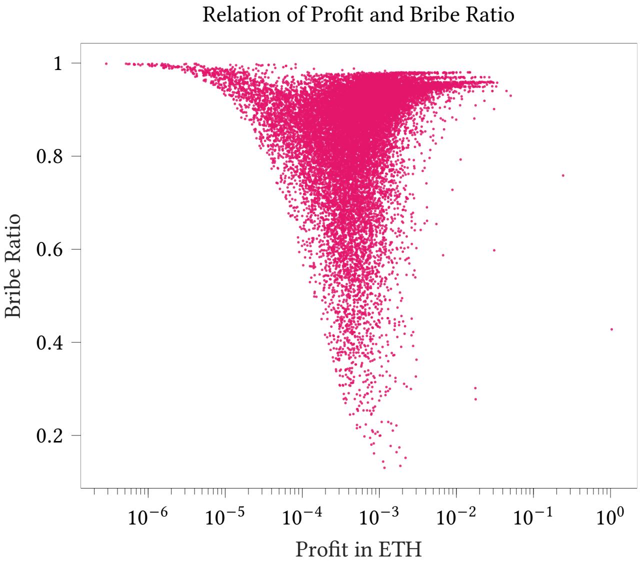

Figure 51: Bribe ratio compared to profit for the Top 3 bot for sandwich attacks.

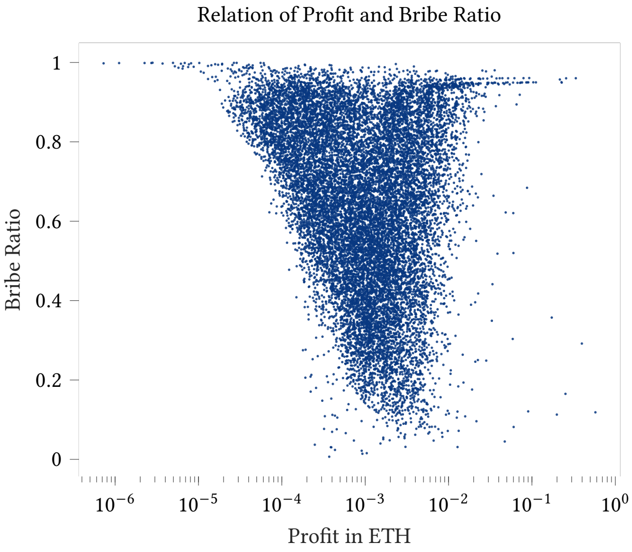

Figure 53: Bribe ratio compared to profit for the Top 4 bot for sandwich attacks.

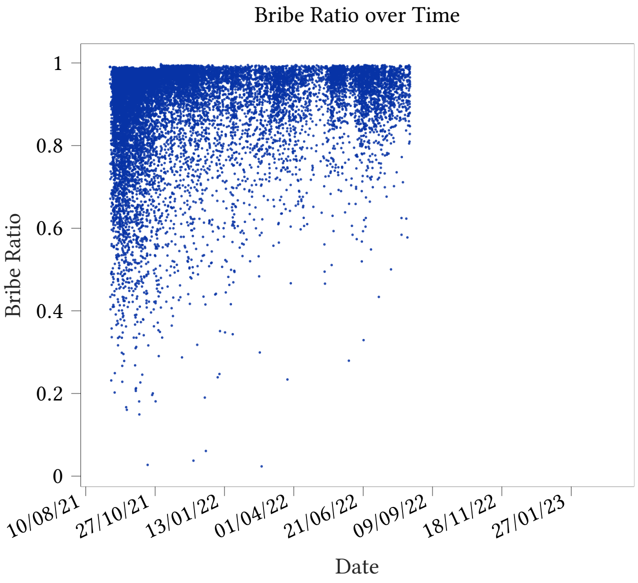

Figure 54: Bribe ratio over time for the Top 5 bot for sandwich attacks.

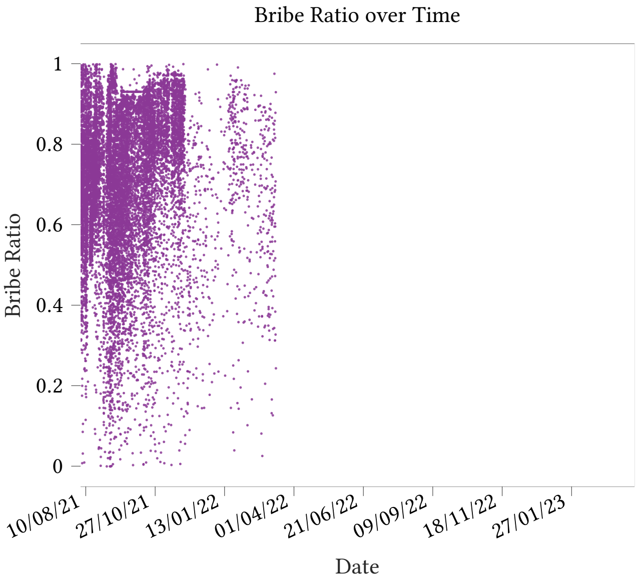

Figure 56: Bribe ratio over time for the Top 6 bot for sandwich attacks.

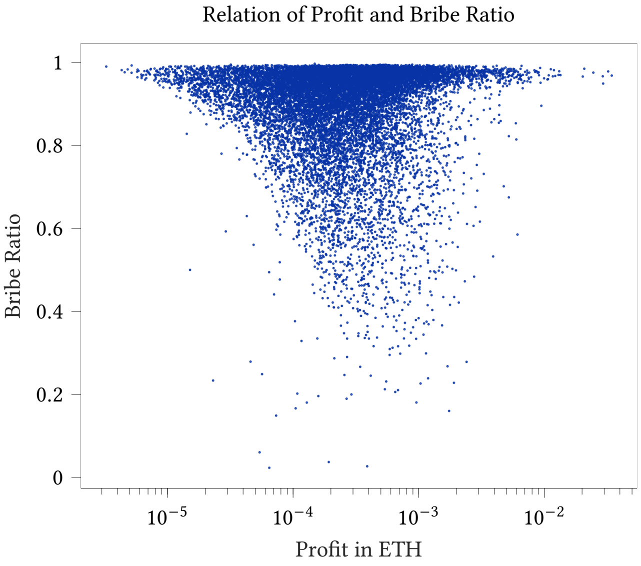

Figure 55: Bribe ratio compared to profit for the Top 5 bot for sandwich attacks.

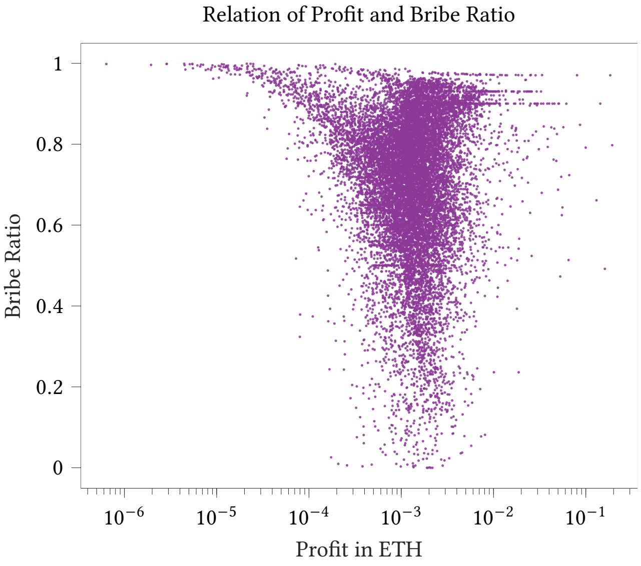

Figure 57: Bribe ratio compared to profit for the Top 6 bot for sandwich attacks.

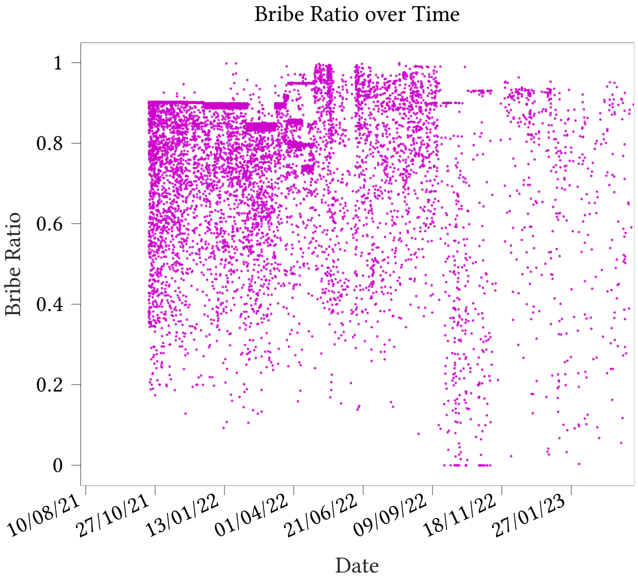

Figure 58: Bribe ratio over time for the Top 7 bot for sandwich attacks.

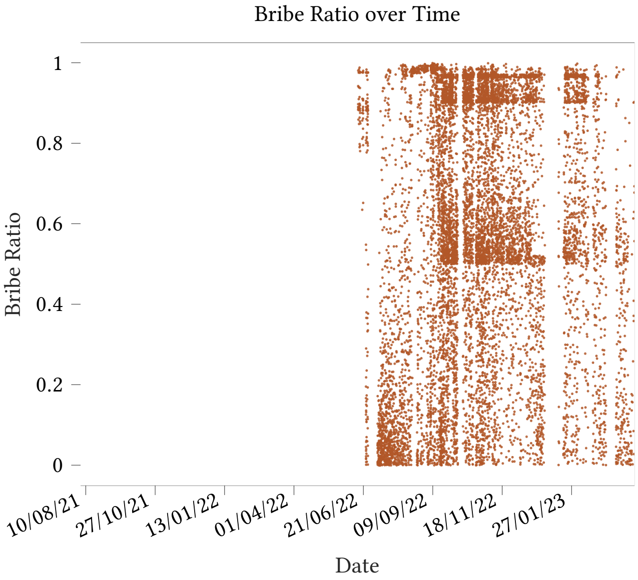

Figure 60: Bribe ratio over time for the Top 8 bot for sandwich attacks.

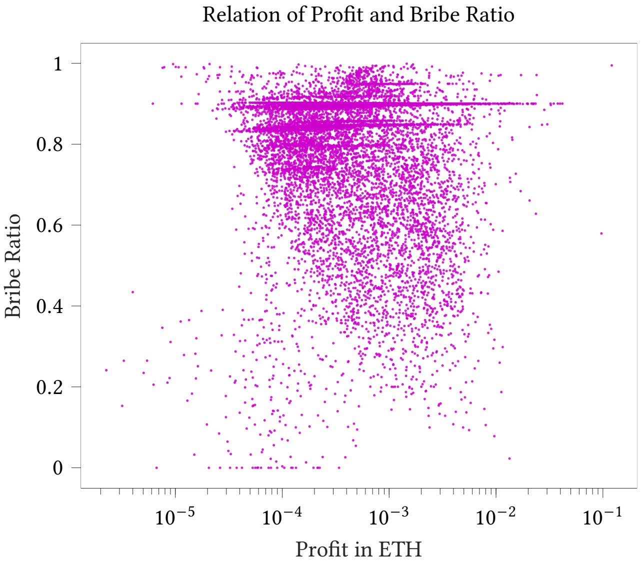

Figure 59: Bribe ratio compared to profit for the Top 7 bot for sandwich attacks.

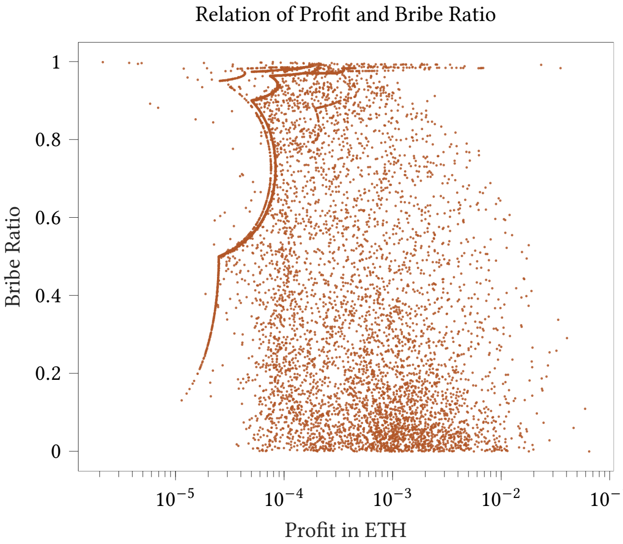

Figure 61: Bribe ratio compared to profit for the Top 8 bot for sandwich attacks.

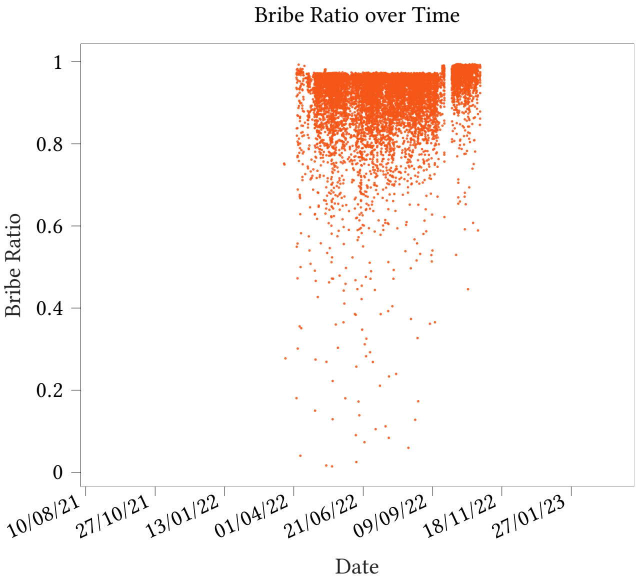

Figure 62: Bribe ratio over time for the Top 9 bot for sandwich attacks.

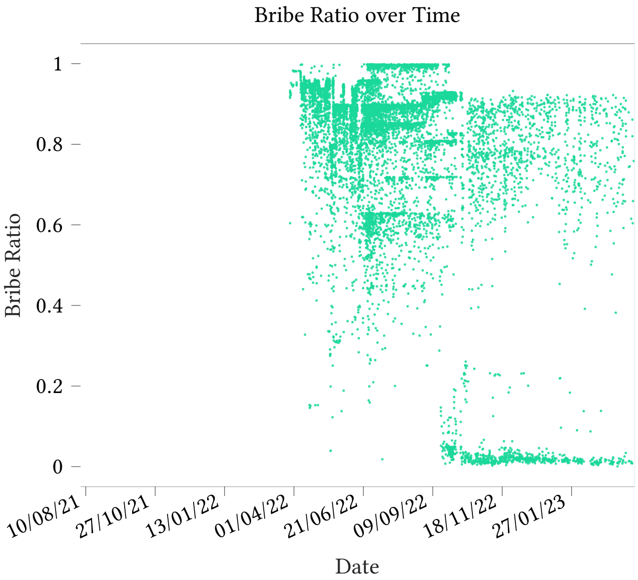

Figure 64: Bribe ratio over time for the Top 10 bot for sandwich attacks.

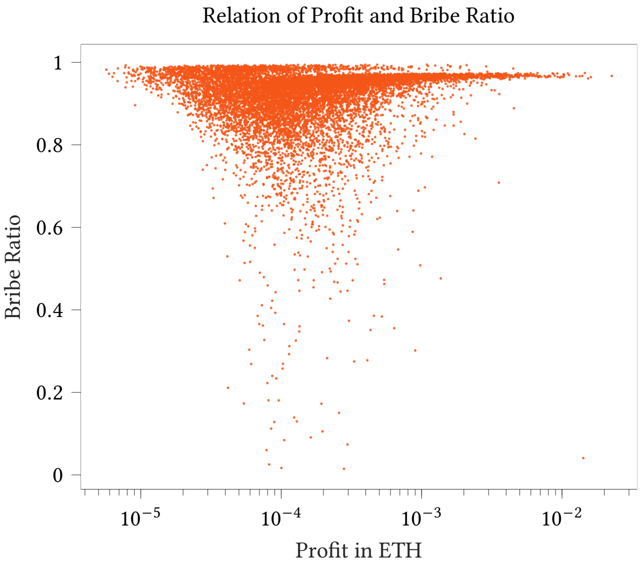

Figure 63: Bribe ratio compared to profit for the Top 9 bot for sandwich attacks.

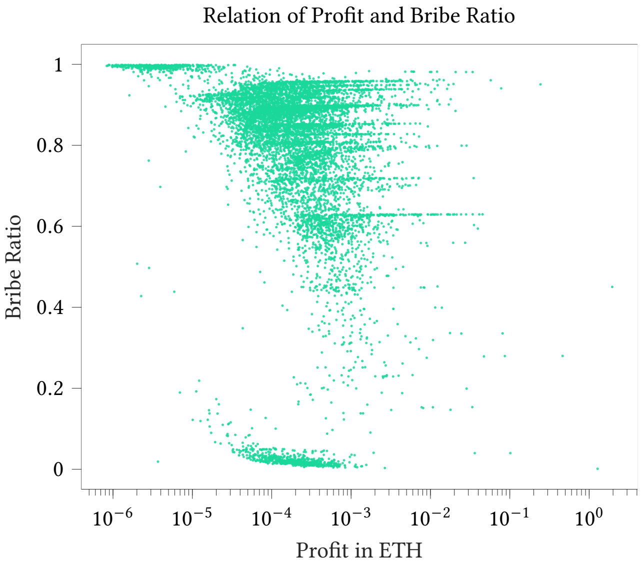

Figure 65: Bribe ratio compared to profit for the Top 10 bot for sandwich attacks.

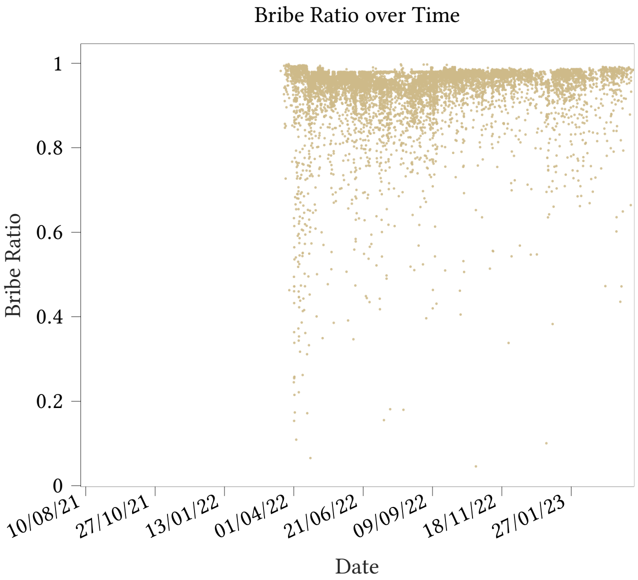

Figure 66: Bribe ratio over time for the Top 11 bot for sandwich attacks.

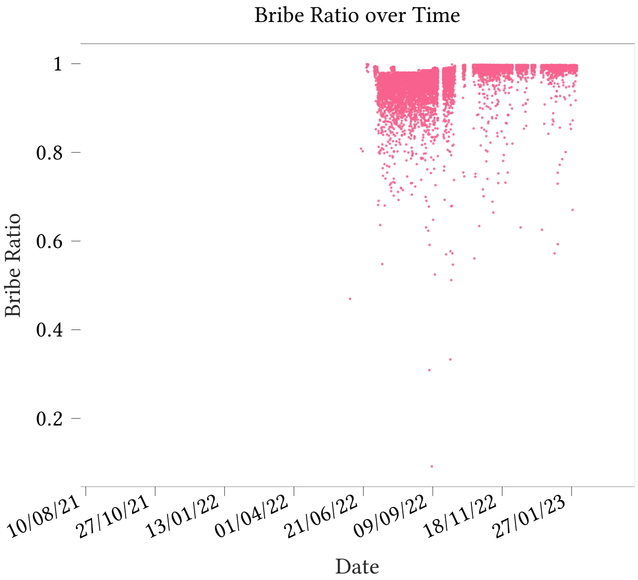

Figure 68: Bribe ratio over time for the Top 12 bot for sandwich attacks.

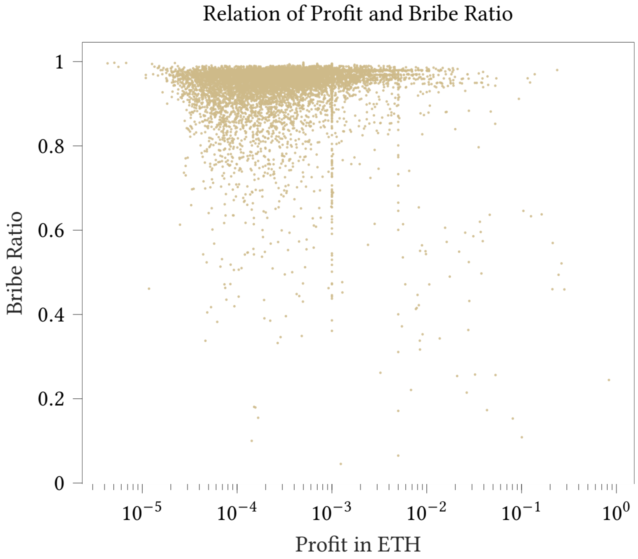

Figure 67: Bribe ratio compared to profit for the Top 11 bot for sandwich attacks.

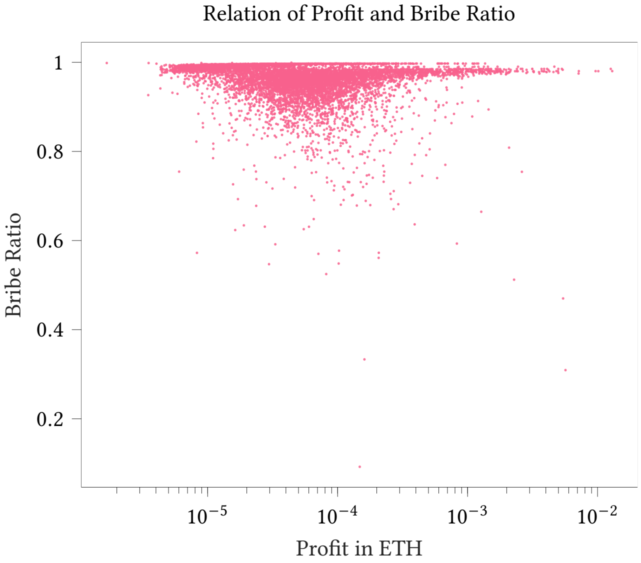

Figure 69: Bribe ratio compared to profit for the Top 12 bot for sandwich attacks.

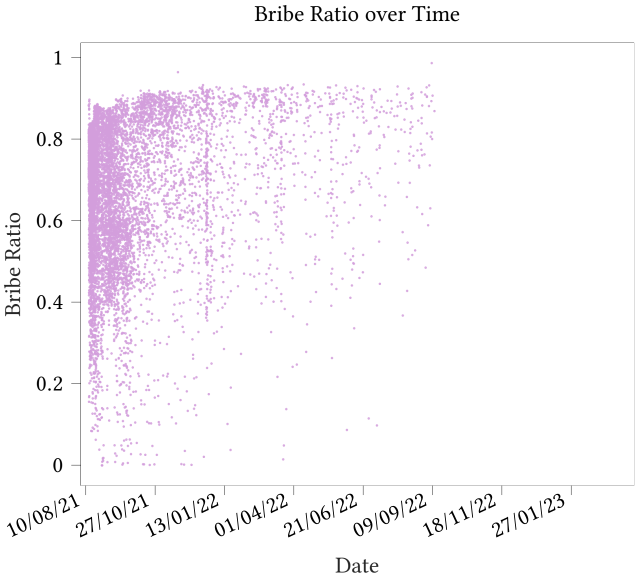

Figure 70: Bribe ratio over time for the Top 13 bot for sandwich attacks.

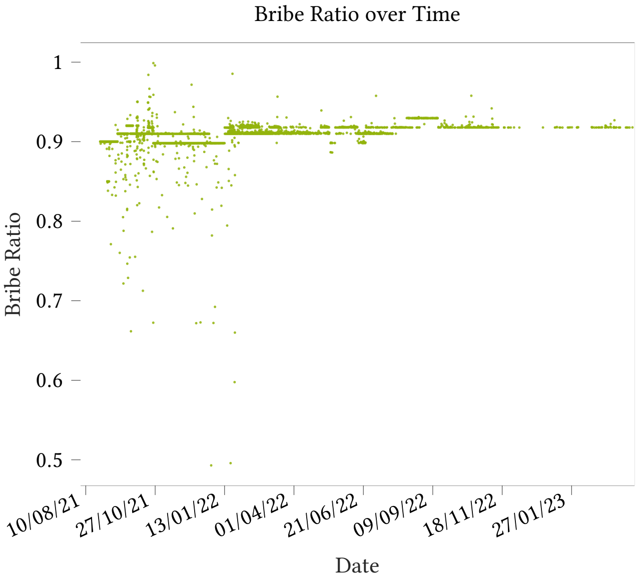

Figure 72: Bribe ratio over time for the Top 14 bot for sandwich attacks.

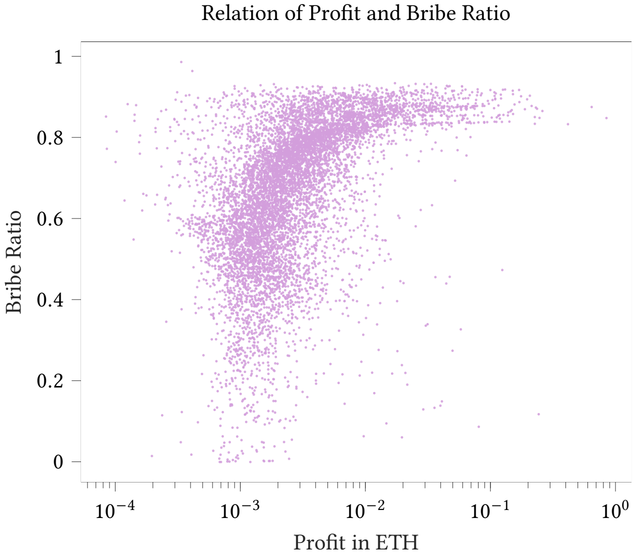

Figure 71: Bribe ratio compared to profit for the Top 13 bot for sandwich attacks.

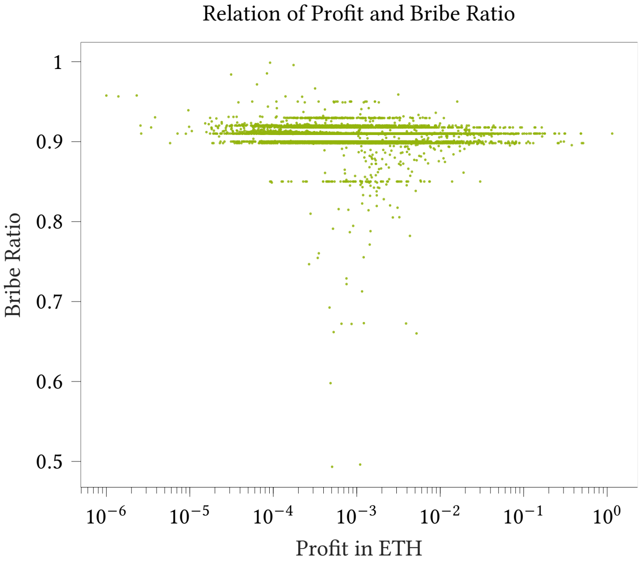

Figure 73: Bribe ratio compared to profit for the Top 14 bot for sandwich attacks.

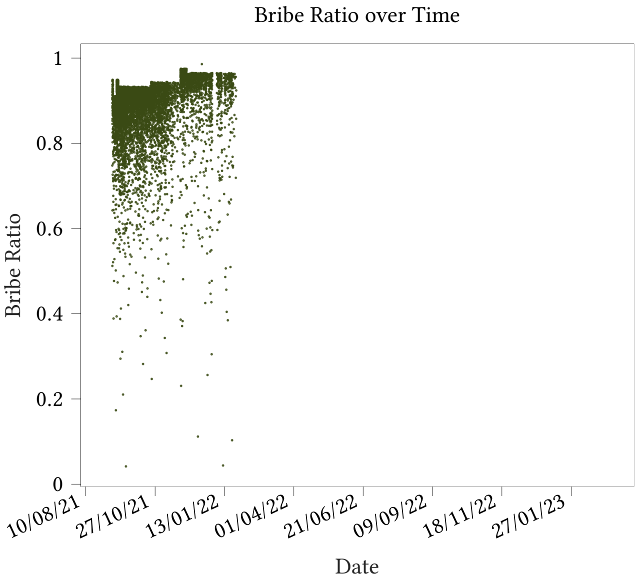

Figure 74: Bribe ratio over time for the Top 15 bot for sandwich attacks.

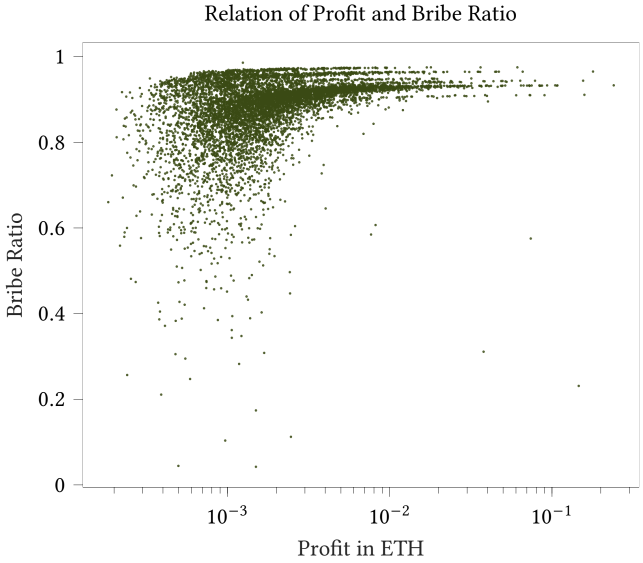

Figure 75: Bribe ratio compared to profit for the Top 15 bot for sandwich attacks.

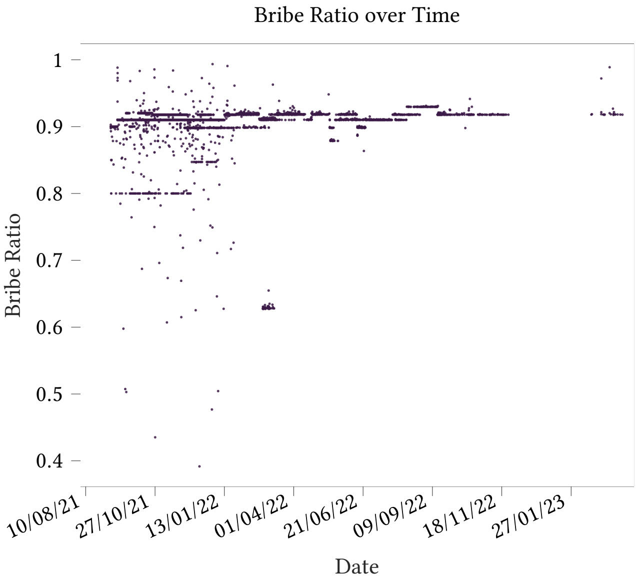

Figure 76: Bribe ratio over time for the Top 16 bot for sandwich attacks.

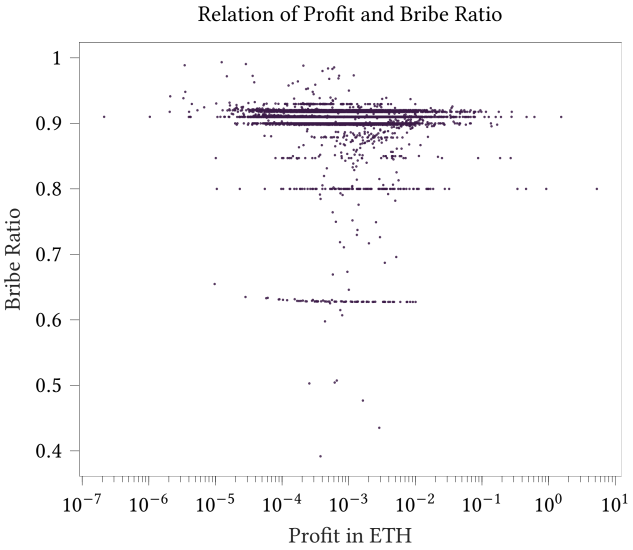

Figure 77: Bribe ratio compared to profit for the Top 16 bot for sandwich attacks.