On Fairness Concerns in the Blockchain Ecosystem

A dissertation submitted towards the degree
Doctor of Engineering
of the Faculty of Mathematics and Computer Science of
Saarland University

by
Johnnatan Messias Peixoto Afonso

Saarbrücken
2023

Date of Colloquium:	April 25th, 2024
Dean of Faculty:	Univ.-Prof. Dr. Roland Speicher
Chair of the Committee:	Prof. Dr. Anja Feldmann
Reporters	
First Reviewer:	Prof. Dr. Krishna P. Gummadi
Second Reviewer:	Prof Dr. Ingmar Weber
Third Reviewer:	Prof. Dr. Balakrishnan Chandrasekaran
Fourth Reviewer:	Prof. Dr. Patrick Loiseau
Academic Assistant:	Dr. Abhisek Dash

Abstract

Blockchains revolutionized centralized sectors like banking and finance by promoting decentralization and transparency. In a blockchain, information is transmitted through transactions issued by participants or applications. Miners crucially select, order, and validate pending transactions for block inclusion, prioritizing those with higher incentives or fees. The order in which transactions are included can impact the blockchain final state.

Moreover, applications running on top of a blockchain often rely on governance protocols to decentralize the decision-making power to make changes to their core functionality. These changes can affect how participants interact with these applications. Since one token equals one vote, participants holding multiple tokens have a higher voting power to support or reject the proposed changes. The extent to which this voting power is distributed is questionable and if highly concentrated among a few holders can lead to governance attacks.

In this thesis, we audit the Bitcoin and Ethereum blockchains to investigate the norms followed by miners in determining the transaction prioritization. We also audit decentralized governance protocols such as Compound to evaluate whether the voting power is fairly distributed among the participants. Our findings have significant implications for future developments of blockchains and decentralized applications.

Zusammenfassung

Blockchain-Technologien revolutionierten zentralisierte Bereiche wie Bankwesen und Finanzen, indem sie Dezentralisierung und Transparenz förderten. In einer Blockchain wird Informationen durch Transaktionen übertragen, die von Teilnehmern oder Anwendungen ausgestellt werden. Miner wählen Transaktionen aus, ordnen sie an und validieren sie für die Aufnahme in einen Block. Dabei priorisieren sie jene Transaktionen mit höheren Gebühren. Die Reihenfolge, in der Transaktionen aufgenommen werden, kann den endgültigen Zustand der Blockchain beeinflussen.

Anwendungen, die auf einer Blockchain laufen, oft auf Governance-Protokolle angewiesen, um die Entscheidungsbefugnis zur Änderung ihrer Kernfunktionalität zu dezentralisieren. Diese Änderungen können beeinflussen, wie Teilnehmer mit diesen Anwendungen interagieren. Da ein Token einem Stimmrecht entspricht, haben Teilnehmer mit mehreren Tokens eine höhere Abstimmungsbefugnis, um die vorgeschlagenen Änderungen zu unterstützen oder abzulehnen. Fraglich ist, inwieweit diese Abstimmungsbefugnis verteilt ist.

In dieser Arbeit prüfen wir die Bitcoin- und Ethereum-Blockchains, um die Normen zu untersuchen, denen Miner folgen, um die Priorisierung von Transaktionen festzulegen. Wir überprüfen dezentrale Governance-Protokolle wie Compound, um festzustellen, ob die Abstimmungsbefugnis fair unter den Teilnehmern verteilt ist. Unsere Ergebnisse haben wesentliche Auswirkungen auf zukünftige Entwicklungen von Blockchains und dezentralen Anwendungen.

Publications

Parts of this thesis have appeared in the following publications and technical reports.

- "Understanding Blockchain Governance: Analyzing Decentralized Voting to Amend DeFi Smart Contracts". **J. Messias**, V. Pahari, B. Chandrasekaran, K. P. Gummadi, and P. Loiseau. This work has been submitted and we are currently awaiting a decision.

- "Dissecting Bitcoin and Ethereum Transactions: On the Lack of Transaction Contention and Prioritization Transparency in Blockchains". **J. Messias**, V. Pahari, B. Chandrasekaran, K. P. Gummadi, and P. Loiseau. In *Proceedings of the 27th Financial Cryptography and Data Security (FC)*, Bol, Brač, Croatia, May 2023.

- "Selfish & Opaque Transaction Ordering in the Bitcoin Blockchain: The Case for Chain Neutrality". **J. Messias**, M. Alzayat, B. Chandrasekaran, K. P. Gummadi, P. Loiseau, and A. Mislove. In *Proceedings of the 21st ACM SIGCOMM Internet Measurement Conference (IMC)*, Virtual Event, November 2021.

- "On Blockchain Commit Times: An analysis of how miners choose Bitcoin transactions". **J. Messias**, M. Alzayat, B. Chandrasekaran, and K. P. Gummadi. In *2nd International KDD Workshop on Smart Data for Blockchain and Distributed Ledger (SDBD)*, Virtual Event, August 2020.

Additional publications and technical reports while at MPI-SWS.

- "Modeling Coordinated vs. P2P Mining: An Analysis of Inefficiency and Inequality in Proof-of-Work Blockchains". M. Alzayat, **J. Messias**, B. Chandrasekaran, K. P. Gummadi, and P. Loiseau. June 2021. (**Technical report**)

- "(Mis)Information Dissemination in WhatsApp: Gathering, Analyzing and Countermeasures". G. Resende, P. Melo, H. Sousa, **J. Messias**, M. Vasconcelos, J. Almeida, and F. Benevenuto. In *Proceedings of the 28th Web Conference (WWW)*, San Francisco, USA, May 2019.

- "WhatsApp Monitor: A Fact-Checking System for WhatsApp". P. Melo, **J. Messias**, G. Resende, K. Garimella, J. Almeida, and F. Benevenuto. In *Proceedings of the 13th International AAAI Conference on Web and Social Media (ICWSM)*, Munich, Germany, June 2019.

- "Search Bias Quantification: Investigating Political Bias in Social Media and Web Search". J. Kulshrestha, M. Eslami, **J. Messias**, M. B. Zafar, S. Ghosh, K. P. Gummadi, and K. Karahalios. In *Information Retrieval Journal*, Springer. Volume 22, Issue 1-2, April 2019.

- "On Microtargeting Socially Divisive Ads: A Case Study of Russia-Linked Ad Campaigns on Facebook". F. N. Ribeiro, K. Saha, M. Babaei, L. Henrique, **J. Messias**, F. Benevenuto, O. Goga, K. P. Gummadi, and E. M. Redmiles. In *Proceedings of the Conference on Fairness, Accountability, and Transparency (FAT*)*, Atlanta, Georgia. January 2019.

I dedicate my thesis to the following people who have played an important role in my life: my wife Aline, my father J. Missias (Jota Missias, 2021), my mother Varlene, and my brothers Joarlens and Jeanderson.

Acknowledgements

I would like to express my gratitude to my advisor, Krishna P. Gummadi, for his invaluable feedback and unwavering support throughout my PhD journey. I am also deeply thankful to Balakrishnan Chandrasekaran and Patrick Loiseau for their constructive insights and encouragement in my studies. I must also extend my appreciation to my previous advisor, Fabrício Benevenuto, whose consistent guidance and support carried me through this PhD journey.

Working alongside remarkable individuals from our Networked Systems group has been a privilege. I would like to thank Ayan Majumdar, Abhisek Dash, Camila Kolling, David Miller, Junaid Ali, Sepehr Mousavi, Till Speicher, Vabuk Pahari, Vedant Nanda, Nina Grgić-Hlača, and many others who have made an impact in my PhD journey.

I extend my thanks to the exceptional research assistants at MPI-SWS, who enthusiastically celebrated every milestone in my doctoral studies. A special mention goes to Mohamed Alzayat, whose feedback and collaboration were instrumental as I embarked on my research topics. I am also grateful to Ahana Ghosh, Andi Nika, Andrea Borgarelli, Angelica Goetzen, Chao Wen, Debasmita Lohar, George Tzannetos, Heiko Becker, Jan-Oliver Kaiser, Lennard Gäher, Matheus Stolet, Michael Sammler, Mihir Vahanwala, Nils Müller, Oshrat Ayalon, Pierfrancesco Ingo, Ralf Jung, Rati Devidze, Roberta De Viti, Victor Alexandru Padurean, Vaastav Anand, and many others for their steadfast support.

I am grateful for the vibrant interactions I had with interns who shared and discussed intriguing research ideas, with a special shoutout to Aleksa Sukovic, Ani Saxena, Ana-Andreea Stoica, Baltasar Dinis, Barbara Gomes, Daniel Kansaon, Diogo Antunues, Isadora Salles, Ignacio Tiraboschi, Lucas Costa De Lima, Maria Petrisor, Pedro Las-Casas, Ruchit Rawal, Sapana Chaudhary, and many others.

I am thankful to all MPI-SWS staff, in particular, Annika Meiser, Carina Schmitt, Christian Klein, Claudia Richter, Gretchen Gravelle, Krista Ames, Maria-Louise Albrecht, Rose Hoberman, and Sarah Naujoks for their sincere efforts and help during my PhD studies.

My appreciation extends to my wife Aline for her unwavering love and support throughout this journey. I also want to extend my thanks to my father, Jota Missias, my mother, Varlene, and my brothers, Joarlens and Jeanderson, for their unconditional support.

Table of contents

List of figures

List of tables

Chapter 1

Introduction

Blockchains have the potential to transform traditional and centralized sectors of great societal importance, such as banking and finance (Adams et al., 2021; Daian et al., 2020; Perez et al., 2021; Qin et al., 2021). They provide a secure means of ensuring compliance via contracts (i.e., established agreements) and tamper-proof mechanisms, especially in situations where participants cannot trust each other (Nakamoto, 2008; Sasson et al., 2014; Van Saberhagen, 2013; Wood et al., 2014). As a result, there are many blockchains available such as Bitcoin (Nakamoto, 2008), Ethereum (Wood et al., 2014), Polkadot (Wood, 2016), Zcash (Sasson et al., 2014), Monero (Van Saberhagen, 2013), among others. Bitcoin and Ethereum stand out as the most widely used blockchains, with market capitalization surpassing $536.72B and $228.57B as of May 2023 (CoinMarketCap, 2023), respectively. In addition, blockchains have not only been used to implement cryptocurrencies, but are increasingly being adopted across a wide range of domains including insurance (Martin Ruubel, 2018), education (Philipp Schmidt, 2015), healthcare (Ekblaw and Azaria, 2017), supply-chain management (Provenance, 2015; Robert Hackett, 2017), decentralized governance (Adams et al., 2021; Leshner and Hayes, 2019; MakerDAO, 2023), and decentralized finance (DeFi) applications through smart contracts such as exchanges (Daian et al., 2020; Uniswap, 2023), lending (Perez et al., 2021; Qin et al., 2021), and auctions (Ethereum Foundation, 2023c).

In the blockchain space multiple parties interact with each other. These include: (i) transaction issuers who are responsible for issuing transactions via interactions with the blockchain and its applications through smart contracts; (ii) miners or block validators who ensure the validity of forthcoming information or blocks for inclusion within all its transactions; and (iii) smart contract applications that are software programs running atop a blockchain, capable of executing predetermined actions, creating or transferring tokens, enabling voting for smart contract amendments, etc. In any real-world scenario involving a diverse group of individuals with varying roles, establishing a foundation

of trust and fairness becomes paramount to ensure that no one can take advantage of others.

Unlike the past, where interactions occurred mostly between individuals who knew and trusted each other, the rise of blockchain has enabled interactions within a decentralized system, devoid of inherent trust. However, in such a trustless environment, the potential for unfairness arises. One of the captivating aspects of blockchain systems is the interaction among participants who are strangers to each other, lacking pre-established trust. This raises questions about whether interactions are conducted fairly and what fairness concerns exist. For instance, in this thesis we focus on three primary unfairness issues: (i) Transaction ordering; (ii) Transaction transparency; and (iii) Fair distribution of voting power for smart contract amendments.

Fairness related to transaction ordering. The sequence in which transactions are processed is crucial, as everyone seeks timely processing. Ensuring fairness in this ordering presents challenges. For instance, how do we know that the ordering is fair?

In other words, noticeably absent from Bitcoin, Ethereum, and other decentralized blockchains is the requirement of any a priori trust between the users issuing transactions (i.e., registers persisted in the blockchain), the miners confirming transactions, and the peer-to-peer (P2P) nodes maintaining the blockchain. Despite their widespread use in ordering critical applications (Daian et al., 2020; Kharif, 2017; McCorry et al., 2017; Perez et al., 2021; Pilkington, 2016; Uniswap, 2023), blockchain protocols formally specify *neither* the manner by which miners should select transactions for inclusion in a new block from the set of all available transactions, *nor* the order in which they should be included in the block. While informal conventions or *norms* for prioritizing transactions exist, to our knowledge, before us no one has systematically verified if these norms were being followed by miners in practice (Messias et al., 2020, 2021, 2023a).

Studying this problem has significant implications for both blockchain users and miners. Specifically, when setting fees for their transactions, transaction issuers (i.e., through their wallet software) assume that the fees offered by all their competing transactions are fully transparent—our findings contradict this assumption. Similarly, when transactions offer different confirmation fees to different miners, it raises significant unfairness concerns with respect to the order in which these transactions are included. We also show that mining pools collude when prioritizing self-interested transactions for inclusion which can exacerbates the growing concerns about the concentration of hash rates amongst a few miners in proof-of-work blockchains (PoW) (Bahack, 2013; Gervais et al., 2016).

Fairness related to transaction transparency. Assumptions dictate that all participants can observe public transactions. This transparency impacts transaction prioritization

and fee-setting. However, reality often diverges, leading to concerns about fairness in transaction transparency.

For instance, the lack of transparency in blockchains arises from genuine concerns of transactions issuers, which cannot be overlooked. One significant concern is the risk of transactions being front-run by bots (Daian et al., 2020; Eskandari et al., 2020; Torres et al., 2021; Weintraub et al., 2022), which creates the need for transaction privacy. Mining pools that address this need also facilitate, unsurprisingly, off-chain payments via which transaction issuers can (privately) incentivize the miners (BTC.com, 2022; Messias et al., 2021; ViaBTC, 2022). We consider these developments as natural and logical steps in the evolution of blockchains and back our assertions with empirical observations. In contrast to prior research (Daian et al., 2020; Strehle and Ante, 2020), we argue that the fundamental threat to blockchain stability lies in the opacity of the overall fees issued by transaction issuers. Most wallet software and crypto-exchanges currently rely on reconstructing the current public Mempool state to suggest an appropriate fee to transaction issuers. As a result, transaction issuers cannot precisely determine the fee required to ensure the inclusion of their transaction in the next block. Consequently, miners can overcharge them, as the "real" fees are opaque to the rest of the network (Weintraub et al., 2022).

Fairness related to voting power to amend smart contracts. Smart contracts, serving as trust-enforcing mechanisms, entail participants' agreement with stipulated rules. However, as these smart contracts can be upgraded (or changed) it raises the question of who possesses the authority to modify them.

Put differently, blockchains face challenges related to decentralized decision-making processes for amending smart contracts. For example, blockchains have been explored by many prior works who studied different types of security vulnerabilities that arise from incorrect implementations or unintended (or undesired) executions of smart contracts over blockchains, particularly in the context of DeFi applications (Daian et al., 2020; Mike Dalton, 2022; Qin et al., 2021; Torres et al., 2021; Weintraub et al., 2022). However, few studies, if any, focused, however, on vulnerabilities that may originate in the design of the procedures to *amend*, i.e., change, smart contracts through *governance protocols*, and/or stem from the execution of these procedures in practice. These governance protocols intend to eliminate (or at least minimize) centralized decision-making in blockchains. Their effectiveness in achieving that goal can, however, be compromised depending on how the tokens (i.e., voting power—typically one token equals one vote) are distributed which can lead to voting concentration. Such concentration poses a threat to the overall governance of smart contract applications in blockchains leading, for example, to governance attacks (Mike Dalton, 2022).

Therefore, in this thesis, we also provide an in-depth analysis of the voting patterns, delegation practices, and outcomes of proposals in one of the widely used governance protocols: Compound (Compound Labs, Inc., 2022a; Leshner and Hayes, 2019). Since Compound records the votes cast transparently on a blockchain (i.e., it uses on-chain voting), we conduct measurements studies to analyze the extent to which this voting is decentralized, i.e., how small or large are the set of voters that determine the outcomes for the amendments.

It is important to acknowledge that our focus on these fairness concerns does not imply exclusivity. Additional concerns exist, such as fairness in compensating miners proportionately to their contributions. Nonetheless, in this thesis, we focus on the three aforementioned fairness concerns: (i) Transaction ordering; (ii) Transaction transparency; and (iii) Fair distribution of voting power for smart contract amendments.

1.1 Overview of thesis contributions

This thesis aims to address the fairness challenges mentioned above. We outline our research contributions in pursuit of this objective below.

1.1.1 Transaction prioritization norms (Messias et al., 2020, 2021)

The conventional wisdom today is that many miners follow the prioritization norms, implicitly, by using widely shared blockchain software like the Bitcoin Core (bitcoin.org, 2023; Coin Dance, 2021). In Bitcoin, the presumed "norm" is that miners prioritize a transaction for inclusion based on its offered *fee rate* or fee-per-byte, which is the transaction's fee divided by the transaction's size in bytes. We show evidence of this presumed norm in Figure 1.1. The norm is also justified as "incentive compatible" because miners wanting to maximize their rewards, i.e., fees collected from all transactions packed into a size-limited block, would be incentivized to include preferentially transactions with higher fee rates. Assuming that miners follow this norm, Bitcoin users are issued a crucial recommendation: To accelerate the confirmation of a transaction, particularly during periods of congestion, they should increase the transaction fees. We show that miners are, however, free to deviate from this norm and such norm violations cause irreparable economic harm to users.

We summarize our contributions as follows.

► To quantify the deviation from the norm, we propose two measures that we call ***signed position prediction error (SPPE)*** and ***position prediction error (PPE)***. These measures

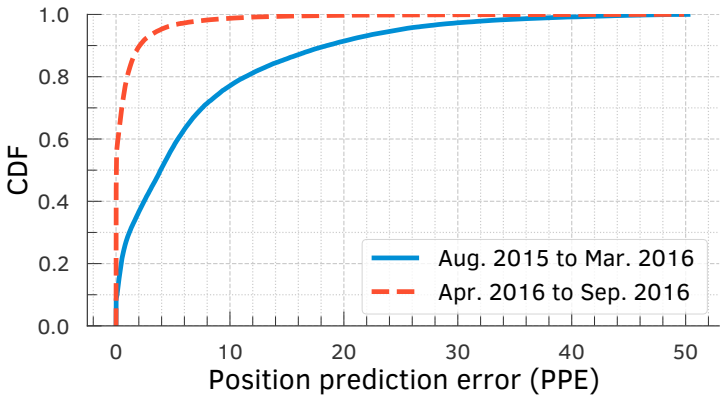

Figure 1.1: CDF of the error in predicting where a transaction would be positioned or ordered within a block according to the greedy fee-rate-based norm. Bitcoin Core code shifted completely to the fee-rate-based norm starting April 2016: Transaction ordering in Bitcoin closely tracks the fee-rate-based norm from April 2016, but differs significantly from it prior to April 2016 when a different norm was in place.

allow us to quantify the transaction deviation (or acceleration) for all transactions in block. It can also be applied to any blockchain that order transactions based on fees-incentive.

► We perform an extensive empirical audit of the miners' behavior to check whether they conform to the norms. At a high-level, we find that transactions are indeed primarily prioritized according to the assumed norms. We also, nevertheless, offer evidence of a non-trivial fraction of priority-norm violations amongst confirmed transactions. An in-depth investigation of these norm violations uncovered many highly troubling misbehavior by miners.

► Multiple large mining pools tend to *selfishly prioritize* transactions in which they have a vested interest; e.g., transactions in which payments are made from or to wallets owned by the mining pool operators. Some even *collude* with other large mining pools to prioritize their transactions.

► Many large mining pools accept additional *dark (opaque) fees* to accelerate transactions via non-public side-channels (e.g., their websites). Such dark-fee transactions violate an important, but unstated assumption in blockchains that confirmation fees offered by transactions are transparent and equal to all miners.

► We release the data sets and the scripts used in our analyses to facilitate others to reproduce our results (Messias, 2023b).

1.1.2 Transaction Prioritization and Contention Transparency (Messias et al., 2023a)

In the context of the *lack of contention transparency*, not all transactions are publicly broadcasted. Instead, users can submit transactions to a subset of miners or mining pools through *private channels* or *relays* that are not visible to the public. In this case, transactions remain private to the relay until they are committed into a block. Additionally, users may choose to submit their transactions exclusively to a particular mining pool that guarantees a fast commit time. This thesis aims to shed light on the growing prevalence of private mining practices, where transactions are submitted to only a subset of the miners. Furthermore, it analyzes the distinct characteristics of these private transactions.

Moreover, with the *lack of prioritization transparency*, the fees offered by a transaction can be significantly higher than what is publicly declared. For example, a transaction can privately offer additional fees to a miner in order to "accelerate" its inclusion in a block. Many such transaction-accelerator (or *front-running as a service (FRaaS)*) platforms exist for Bitcoin (BTC.com, 2022; ViaBTC, 2022) and Ethereum (Eskandari et al., 2020; Flashbots, 2022b; SparkPool, 2021; Strehle and Ante, 2020). Furthermore, the same transaction can offer different fees to different mining pools through their relays. The existence of these hidden or dark-fees can undermine the reliability of any fee prediction: Transaction issuers may end up paying considerably higher fees without receiving proportional or any reduction in commit delays. This thesis aims to characterize the prevalence of such dark-fee transactions and analyze the most popular private relay network available in Ethereum, Flashbots (Flashbots, 2022b). We also conduct active experiments in both Bitcoin and Ethereum to validate our assumptions regarding the prioritization transparency.

In addition to demonstrating the non-uniformity of transaction fees across miners, we argue that, given the lack of contention transparency, the lack of prioritization transparency may become even more widespread in the future.

We summarize our contributions as follows.

► We provide a comprehensive characterization of the lack of contention transparency in both Bitcoin and Ethereum. Our analysis reveals the widespread use of private channels or relay networks to submit transactions directly to a subset of miners. This practice has the potential to undermine prioritization transparency, as transaction issuers may not be able to estimate the appropriate fees once none of which is publicly visible.

► We investigate the prevalence of private transaction fees, with a particular focus on Flashbots bundles in Ethereum. Our findings indicate that Flashbots bundles represent a significant portion (52.11%) of all Ethereum blocks. This lack of prioritization

transparency may enable miners to overcharge users when they send their transactions privately.

► We investigate whether public transactions are bundled together with private transactions using the Flashbots private relay to exploit arbitrage opportunities through *Maximal Extractable Value (MEV)*. Interestingly, we find that the public transactions within these bundles are, for example, associated with oracle[1] updates, specifically involving the adjustment of prices for particular token pairs. This is made possible by the sequential execution of transactions within the bundle by miners. Consequently, the transaction that is executed right after the oracle update gain immediate access to the updated price information as soon as it is recorded on the blockchain.

► We demonstrate evidence of collusion among Bitcoin miners, collectively possessing more than 50% of the network's total hashing power, particularly concerning the inclusion of dark-fees transactions.

► To promote transparency and facilitate the scientific reproducibility of our results, we publicly release our data sets and scripts used in our analysis (Messias, 2023b).

1.1.3 Decision-making power distribution for blockchain governance (Messias et al., 2023b)

This thesis also provides an in-depth analysis of the voting patterns, delegation practices, and outcomes of proposals in one of the widely used governance protocols: Compound (Compound Labs, Inc., 2022a; Leshner and Hayes, 2019). Since Compound records the votes cast transparently on a blockchain (i.e., uses on-chain voting), we conducted measurements studies to analyze the extent to which this voting is decentralized, i.e., how small or large are the set of voters that determine the outcomes for the amendments. Our goal is to thoroughly examine this protocol to better understand how its governance mechanism operates and identify potential areas for improvement.

Compound regulates its voting process via the Compound (COMP) token, an ERC-20 asset, as follows. First, it allows token holders to participate in governance by proposing and voting on changes to the protocol through an on-chain voting mechanism where voting power of a user is proportional to the amount of delegated tokens held by them—one token equals one vote. Second, it permits its holders to delegate their tokens to other users, enabling users (who do not wish to exercise their voting rights) to delegate their voting power to others. The protocol essentially supports a form of *liquid democracy* that

[1]Decentralized Oracle Networks facilitate off-chain data access for blockchains, including exchanges prices, weather forecast, and more (Breidenbach et al., 2021). Typically, blockchain applications rely on these oracles to gather information they need.

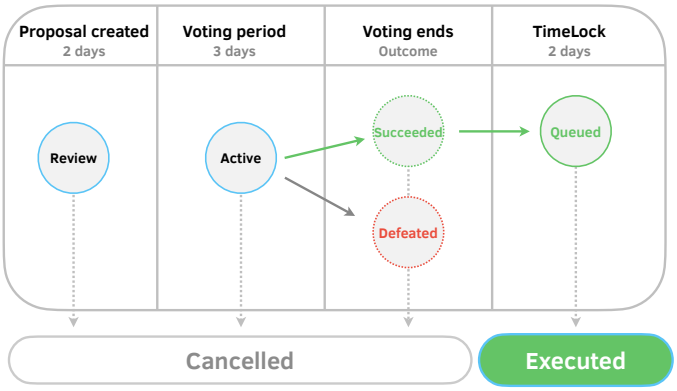

Figure 1.2: The lifecycle of a Compound proposal lasts 7 days. After a proposal is created, it waits for 2 days before the 3-day voting period begins. Once the outcome of the election is decided, it takes 2 more days for the proposal to be executed and become part of the Compound Governance protocol. Proposals can also be cancelled at any time before they are executed.

combines direct democracy and representative democracy, where voters can delegate their voting power to trusted representatives (Behrens, 2017; Blum and Zuber, 2016; Carroll, 1884). Some protocol changes that token holders can propose and vote on include adjustments to the borrowing and lending rates, changes to how new tokens are distributed, and changes to the parameters of the voting process. They can also change the duration of the proposal's life cycle (per Figure 1.2, currently taking on average 7 days) and other aspects of the protocol. To incentivize participation in token lending and borrowing through the protocol, Compound distributes 1234 COMP tokens *daily* to users and applications in various markets (e.g., ETH, DAI, and USDC), in proportion to the amount that they lend or borrow (Compound Labs, Inc., 2022c).

We summarize our contributions as follows.

▶ We characterize the Compound protocol's on-chain voting process, showing that it is active and regularly used, with a steady flow of proposals. The majority of the proposals receive significant support: on average, 89.39% of votes are in favor.

▶ We reveal a substantial variation in voting costs, from \$0.03 to \$294.02, with an average of \$7.88.[2] If we normalize the costs per vote by the count of tokens held by users, we obtain an average cost per vote unit of \$358.54. Voting costs on Compound can, hence, be *unfairly* expensive for small token holders, which has fairness implications for the decision-making process.

[2]All costs are in US dollars, taking into account the exchange rate at the time of casting the vote.

► We show that a small group of 10 voters holds a significant amount of voting power (57.86% of all tokens) and that proposals only required an average of 2.84 voters to obtain at least 50% of the votes. These observations strongly suggest that the voting outcomes in Compound may not reflect the preferences of the broader community.
► We also discover potential voting coalitions among the top voters, which could further exacerbate concerns of voting concentration.
► To foster reproducible research and inspire investigations into other aspects of governance protocols, we share our scripts and data sets in a GitHub repository (Messias, 2023a).

1.2 Thesis outline

In summary, this thesis addresses three important fairness concerns. First is the transaction ordering where we examine the prioritization of transactions by auditing the order in which they are included by miners. We investigate whether miners adhere to the existing norms or widely accepted practices in this regard. Second, is the transaction transparency where we explore whether miners ensure transparency in terms of transaction prioritization and transaction contention (i.e., the transaction and its content is accessible by all miners), ensuring equal access for all participants. Finally, is the fair distribution of voting power for smart contract amendments where we delve into the distribution of voting power in decentralized governance protocols, with a particular focus on Compound. Both of these concerns are crucial for establishing a fair blockchain ecosystem.

Specifically, this thesis is organized as follows:
► In Chapter 2, we begin by presenting the background of blockchains and smart contracts. Next, we delve into the background of transaction prioritization norms. Then, we discuss the background of transaction prioritization and contention transparency. Finally, we explore the background of decentralized governance protocols used for amending smart contracts.
► In Chapter 3, we perform an audit of miners' prioritization norms and evaluate the degree to which they comply with commonly accepted prioritization assumptions. Our findings reveal that miners tend to prioritize transactions based on self-interest, including their own transactions or those from friendly mining pools. Additionally, we highlight the presence of acceleration services provided by miners, enabling off-chain payments to expedite transactions. Unfortunately, the fees associated with these services are dark or opaque to other participants, making it challenging for them to estimate appropriate fees for timely transaction inclusion. Finally, we propose two metrics for

verifying transaction ordering misposition within a block. These metrics play a crucial role in determining whether miners adhere to the assumed norms and can be generalized to other blockchains.

► In Chapter 4, we conduct a data-driven analysis focused on the prioritization of dark-fees payments by miners. Our results reveal the existence and prevalence of private relay networks (e.g., Flashbots), which allow transaction issuers to privately send their transactions to miners, keeping them hidden from other participants. To assess the level of prioritization provided by miners, we performed two active experiments, where we accelerated the inclusion of transactions by offering dark-fee payments to miners. We discovered evidence of potential collusion among mining pools, where their combined hash rate accounted for over 50% of the network's hash rate.

► In Chapter 5, we conduct an in-depth audit of governance protocols, with a specific focus on Compound, to evaluate the extent to which they have succeeded in achieving their primary goal of decentralizing the decision-making process for amending smart contracts. Our analysis reveals that this goal may not have been fully achieved, as we observe a concerning concentration of voting power among a small number of participants, along with the existence of voting coalitions formed by powerful voters. This concentration stands in contrast to the fundamental objective of decentralized governance protocols, which aims to mitigate centralization in the decision-making process.

► In Chapter 6, we review the works in the literature that are relevant to this thesis.

► In Chapter 7, we present a detailed discussion of our work and its limitations. Additionally, we explore potential directions for future work.

Chapter 2

Background

In this chapter, we provide the necessary background on blockchains and smart contracts. We then present the context for transaction prioritization norms, followed by transaction prioritization and contention transparency. Next, we present the background information on decentralized governance. Lastly, we introduce the background details concerning Layer 2.0 solutions aimed at enhancing the scalability of blockchains.

2.1 Blockchains & smart contracts

A blockchain is a decentralized and distributed ledger of cryptographically linked records of transactions stored in blocks. A block is a set of zero[3] or more transactions. In the case of Bitcoin, a block also includes a Coinbase transaction, responsible for transferring rewards (or compensation for the miner' efforts) to the miner's wallet. These blocks are interconnected, tracing back to the original (or "Genesis") block. Figure 2.1 presents an illustration of a blockchain with a block inclusion rate of 10 minutes. The ledger is maintained and continually extended by the blockchain participants by carrying out various functions. *Transaction issuers*, for instance, issue transactions and share it with other participants through a peer-to-peer (P2P) network. Such transactions are deemed *unconfirmed* until they are added permanently to the blockchain. Some others, called *miners*[4] (in proof-of-work blockchains) or *block proposers* (in proof-of-stake blockchains), bundle the transactions into *blocks* for confirmation. They propose the new block over the P2P network where others can verify it and, if successful, add to their copy of the blockchain—thereby extending the ledger or chain of blocks. The miner typically collects a reward in the form of newly minted coins as an incentive for their contribution to

[3] As miners can also mine a block without including any transaction on it, we refer to those blocks as *empty blocks*.

[4] Throughout this thesis, we use the terms *miner*, *mining pool*, *mining pool operator (MPO)*, and *block proposer* interchangeably.

the network along with the transactions fees provided by the issuers. We use the term "fee" to refer generally to the incentive offered by a user to miners for prioritizing the inclusion of their transaction in a block, albeit its exact form may vary, e.g., *fee rate* (or *fee-per-byte*) in Bitcoin and *gas price* in Ethereum[5]. To join a blockchain, miners use a software implementation (along with the hardware) which we refer as a *node*. A node allows a miner to receive broadcasts of transactions and blocks from their peers, validate the data, and mine a block. Nodes queue the unconfirmed transactions received via broadcasts in an in-memory buffer, called the *Mempool*, from where they are dequeued for inclusion in a size-limited block. One can also configure the node to skip mining and simply use it as an observer.

Some blockchains, such as Ethereum, allow for the execution of *smart contracts*—contractual agreements encoded in software programs that run atop the blockchain. Smart contracts are typically developed using Solidity, a domain-specific programming language (Solidity Team, 2023), and they can be executed in the Ethereum Virtual Machine (EVM). Their implementations abide by standards such as ERC-20 (Ethereum Foundation, 2023a) and ERC-721 (Ethereum Foundation, 2023b) to ensure compatibility and interoperability among them. These standards define, for instance, the key functions for creating and implementing smart contracts for fungible tokens (e.g., Compound's COMP token (Leshner and Hayes, 2019)) in the case of ERC-20 or non-fungible tokens (NFTs) in the case of ERC-721.

2.2 Transaction prioritization norms

A crucial detail absent in the design of a blockchain per (Nakamoto, 2008) is any notion of a formal specification of transaction prioritization. Said differently, Nakamoto's design does not formally specify how miners should select a set of candidate transactions for confirmation from all available unconfirmed transactions. Notwithstanding this shortcoming, "norms" have originated from miners' use of a shared software implementation: For example, in Bitcoin, miners predominantly use the Bitcoin Core (bitcoin.org, 2023) software for communicating with their peers (e.g., to advertise blocks and learn about new unconfirmed transactions) and reaching a consensus regarding the chain.

[5]Ethereum recently switched its consensus mechanism from proof-of-work to proof-of-stake with *The Merge* hard fork deployed on September 15, 2022, at block number 15537394 (Ethereum Foundation, 2022a,b).

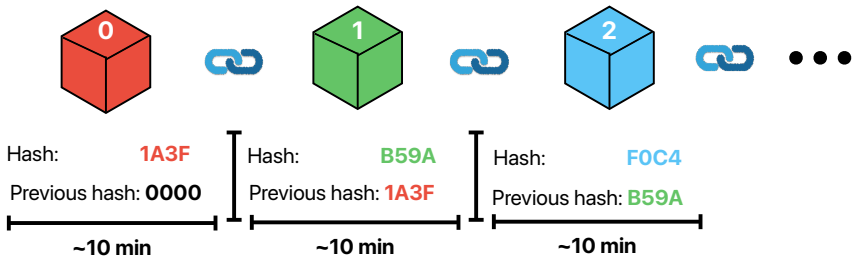

Figure 2.1: Illustration of a blockchain consisting of three blocks, with an extimation generation and inclusion time approximately 10 minutes. It is important to note that every block, except for the Genesis block which is the initial block in the blockchain, also uses the hash of its preceding block to compute its own hash.

Of particular note in the popular Bitcoin core's implementation is the `GetBlockTemplate (GBT)` mining protocol, implemented by the Bitcoin community around February 2012.[6] *GetBlockTemplate* rank orders transactions based on the fee-per-byte (i.e., transaction fees normalized by the transaction's size) metric (Bitcoin Wiki, 2023a).

In Bitcoin, the term *size*, refers to *virtual* size, each unit of which corresponds to four *weight units* as defined in the Bitcoin improvement proposal BIP-141 (Lombrozo et al., 2015). The size (or maximum capacity) of a block is also limited to 1 MB virtual size. The predominant use of GBT (through the use of Bitcoin core) by miners coupled with the fact that GBT is maintained by the Bitcoin community *implicitly* establishes two norms. A third norm stems from a configuration parameter of the Bitcoin core implementation. We now elucidate these three norms.

I. *When mining a new block, miners select transactions for inclusion, from the Mempool, based solely on their fee rates.*

II. *When constructing a block, miners order (place) higher fee rate transactions before lower fee rate transactions.*

III. *Transactions with fee rate below a minimum threshold are ignored and never committed to the blockchain.*

The GBT protocol implementation in Bitcoin core is the source of the first two norms. GBT's rank ordering determines both which set of transactions are selected for inclusion (from the Mempool) and in what order they are placed within a block. GBT dictates that a transaction with higher fee-per-byte *will* be selected before all other transactions with a lower fee-per-byte. It also stipulates that within a block a transaction with the highest fee-per-byte appears first, followed by next highest fee-per-byte, and so on.

[6]Even within mining pools, the widely used Stratum protocol internally uses the *GetBlockTemplate* mechanism (Braiins, 2021a).

The third norm stems from the fee-per-byte threshold configuration parameter. Bitcoin core, by design, will not accept any transaction with fee rate below this threshold, essentially filtering out low-fee-rate transactions from even being accepted into the Mempool. The default (and recommended) value for this configurable threshold is set to $1\,\mathrm{sat/B}$.[7]

2.3 Transaction prioritization and contention transparency

The rate at which users issue transactions in permissionless blockchains, e.g., Bitcoin (Nakamoto, 2008) and Ethereum (Wood et al., 2014), is often much higher than the rate at which miners can include them in a block (Easley et al., 2019; Huberman et al., 2021; Lavi et al., 2019; Messias et al., 2020, 2021). Figure 2.2 shows that 50% of all Bitcoin transactions were added to the blockchain in just 3 years. Similarly, in Ethereum, 80% of the transactions were added in just 1.5 years. Users typically issue transactions using a wallet software, whose primary functionality is determining an "appropriate" fee for a given transaction. This (prioritization) fee varies, unsurprisingly, as a function of the level of congestion in the blockchain (Messias et al., 2021) as well as the distribution of fees across available transactions. Inferring either of these is, however, deceptively complicated.

At first glance, these tasks appear straightforward, since every transaction is broadcast to all miners in the blockchain. A user could simply gather all transactions broadcast over time and reconstruct the set of uncommitted transactions available to a miner (i.e., contents of the miner's Mempool) at any point of time (Messias et al., 2020). We refer to this assumption of a public and uniform view (across miners) of all available transactions as *contention transparency*. If contention transparency exists, a user could rank order available transactions by their fee (based on which miners should select transactions for inclusion) and estimate the commit delay of any transaction (Messias et al., 2021). Consequently, they could determine the fee that they must pay to guarantee inclusion of their transaction in a given block. We label this assumption that the (prioritization) fee offered by a transaction is only that publicly declared by that transaction as *prioritization transparency*. Neither the contention transparency nor the prioritization transparency, however, holds today in permissionless blockchains.

[7] One Bitcoin (BTC) is equal to 10^8 satoshi (sat).

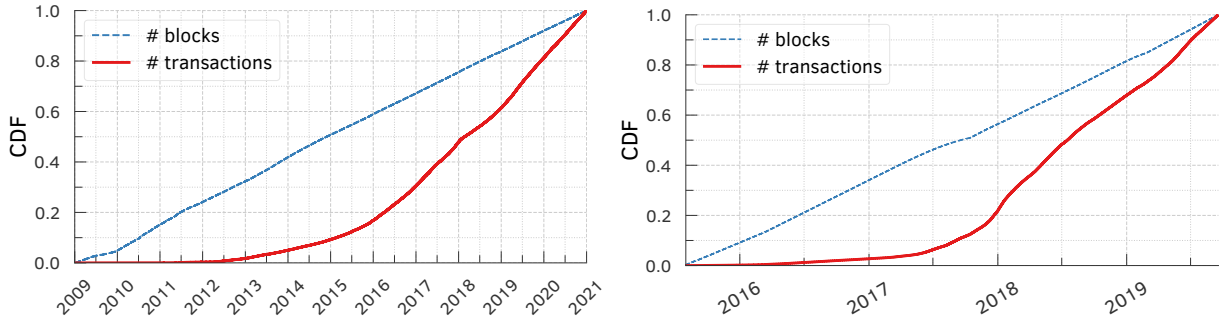

(a) Bitcoin blockchain **(b)** Ethereum blockchain

Figure 2.2: Volume of transactions issued and blocks mined as a function of time, showing that transactions have been issued at high rates for both (a) Bitcoin and (b) Ethereum blockchains.

2.4 Decentralized governance

Smart contracts underpin many DeFi applications today (Adams et al., 2021; Daian et al., 2020; Ethereum Foundation, 2023c; Perez et al., 2021; Qin et al., 2021), and it is only natural to have a mechanism for updating these (software) contracts to fix bugs or evolve them over time to cater for new use cases (David Siegel, 2013; Liu et al., 2018; Zhou et al., 2023). If decisions concerning such updates are made in a centralized manner, e.g., by a regulatory body, or a cabal of developers or miners, it undermines users' trust in the applications that these contracts support. The updates could, for instance, be tailored to benefit the centralized regulatory body at the expense of others. Governance protocols address this issue by distributing the decision-making power among all the users of the application or smart contract being updated.

A governance protocol establishes rules and (transparent) mechanisms for changing smart contracts. It defines the required procedures for creating, voting on, and executing proposals to amend smart contracts. It facilitates users of a protocol (or, more aptly, *token holders* who hold one or more tokens of the protocol) to propose changes. The changes are then vetted by and voted by other users, and implemented only if the proposals receive the majority of favorable votes. The protocols also grant voting power to a user based on the number of tokens held by them—typically one token equals one vote, essentially capturing the user's stake and/or participation in the protocol. Some protocols such as Compound (Leshner and Hayes, 2019) and Uniswap (Adams et al., 2021) allow token holders who do not wish to exercise their voting power to delegate their voting power (i.e., tokens) to others. This delegation is a form of *liquid democracy*, where voters can participate in decision-making either directly by voting or indirectly by delegating their

voting rights to trusted representatives (Behrens, 2017; Blum and Zuber, 2016; Carroll, 1884). Governance protocols give every participant the right to propose, support, or oppose any proposal. They are, hence, crucial for ensuring absolute decentralization of applications running atop blockchains.

2.4.1 Voting modalities

Proposals to change a governance protocol takes birth in the protocol's community forum. Community members suggest and discuss potential changes to the protocol in the forum and may even conduct an informal poll to gauge the community's support for a proposal. The proposer then either amends the proposal to incorporate the community's feedback and submit it for a formal vote, or simply abandon it. The formal voting on the proposal has two modes: *on-chain* and *off-chain* voting.

On-chain voting In this voting system, participants make all governance decisions via smart contracts on a blockchain. Under this system, participants cast a vote by issuing a transaction (and paying a fee for committing it) to the blockchain. The system allows only participants with at least a threshold amount of (governance) tokens to create a proposal, albeit any token holding participant can vote on that proposal. It executes the proposal on the blockchain only if it receives a significant number of votes in favor and reaches a *quorum*.[8] This voting system thus facilitates making transparent and tamper-proof changes to the protocol. Decentralized governance protocols such as AAVE (AAVE, 2023), Compound (Compound Labs, Inc., 2022a), and Uniswap (Uniswap Labs, 2023) use on-chain voting.

Off-Chain Voting This system conducts voting on an off-chain third-party platform and, as a consequence, also establishes the rules for voting, aggregating votes, and determining the results off-chain. Protocols such as Balancer (Balancer.fi, 2023) and Convex Finance (Convex, 2023a), for instance, use Snapshot (Labs, 2023b) for off-chain voting. Snapshot stores the voting data on a P2P network called InterPlanetary File System (IPFS) (Labs, 2023a). The voting process does not require voters to pay any fees and (unlike on-chain voting) promotes participation across all participants, regardless of their level of participation or investment in the protocol. After the voting, this system uses a *multi-signature* contract to enact the off-chain voting outcome on the blockchain. Typically, an *n-of-m multisig* contract requires the transaction to be signed by at least n out of the m "admins" to be executed on-chain. The system trusts the multisig "admins", who

[8] In Compound, in order for a proposal to be executed, it needs to meet two requirements. First, it must receive a minimum of $400{,}000$ votes in favor of the proposal. This number corresponds to 4% of the total supply and is known as the *quorum*. Second, the majority of the votes cast must be in favor of the proposal.

are well-known in the community, to implement the voting outcome on the blockchain truthfully. The admins can also, however, refuse the proposal. In Convex Finance, for instance, the admins can choose not to execute a proposal if they deem it harmful, even if it had received the majority of votes and reached a quorum (Convex, 2023b). On-chain voting systems, in contrast, prevent such manipulation of voting outcomes by one or more individuals (after the voting process), since all governance decisions (e.g., voting and execution) happen on the chain.

Token delegations Some governance protocols (e.g., Compound, Uniswap, and AAVE) require a user to own a certain amount of governance tokens for casting a vote. Users must delegate their tokens either to themselves or, if they do not wish to vote, delegate them to others. The ability to delegate voting power to others facilitates a form of liquid democracy; the token holder who delegates or sells their tokens to another loses their voting power. Delegations allow anyone to buy (or sell) tokens and gain (or lose) voting power instantly. Justin Sun, the founder of (stablecoin) TrueUSD (TrueUSD, 2023), for instance, allegedly borrowed COMP tokens to create and vote for Compound proposal #84 (team, 2022), resulting in a *governance attack* (Thurman, 2022); this proposal was, however, defeated.

Token "locking" Protocols such as Balancer and Curve (Curve, 2023) mandate that a user "lock" their tokens into a smart contract for a specified period of time to gain their right to vote. The user cannot withdraw the locked tokens until the lock-up period expires. The voting power of a user in this system is proportional to the amount of tokens locked as well as the lock-up period. In Balancer, for instance, a user receives 1 unit of voting power if they lock 1 token into the contract for 1 year, and only half that voting power if they lock it instead for 6 months.

Continuous voting A few protocols (e.g., MakerDAO (MakerDAO, 2023)) allow voters to change their votes at any time during the voting period. Users propose a protocol change by developing a new implementation via a smart contract. The new implementation is accepted if it receives more votes than the current one, i.e., the winning implementation must always receive the majority of the votes (or tokens). MakerDAO requires a user to deposit their (MKR) tokens into the (Maker) Governance contract for casting a vote. The more tokens they deposit, the more voting power they obtain, and they vote for their desired implementation by specifying it as a protocol parameter in the smart contract. Since the voting process is continuous, if a user withdraws their MKR tokens from the governance contract, their vote will no longer count towards the implementations for which they previously voted.

We refer the reader to §C.4 (particularly Tab. C.1) for a characterization of the voting methods, delegation approaches, and proposal executions in various other governance protocols. Understanding how voting is conducted (i.e., whether it is on-chain or off-chain) and proposals are executed is fundamental for analyzing how voters, proposers, and others interact with these governance protocols.

2.5 Blockchain Scalability with Layer 2.0 Solutions

As observed in Figure 2.2, 50% of all Bitcoin transactions were added to the blockchain in just 3 years. Meanwhile, in Ethereum, 80% of the transactions were added in just 1.5 years. This highlights a crucial scalability challenge inherent in blockchains: blocks will not be able to accommodate all transaction available. Unfortunately, this issue is projected to aggravate each year. Consequently, the need for Layer 2.0 solutions has arisen to mitigate the escalating blockchain scalability dilemma.

Layer 2.0 consists of an off-chain network, system, or technology that is built on top of a blockchain (Chainlink Foundation, 2023). In other words, this approach execute transactions in batches, together, and off-chain. This process helps accelerate the execution of the transactions leaving the main blockchain for persisting the validity (or proofs) that the transactions were correctly executed. As a result, this method accelerates both transaction execution and confirmation processes.

There are two most popular types of Layer 2.0 solutions: Zero Knowledge (ZK) (Goldreich and Oren, 1994; Goldwasser et al., 2019) and Optimistic (Ethereum Foundation, 2023d) rollups.

Within the Zero Knowledge (ZK) rollup, transactions are grouped into batches for execution, leading to a reduction in execution costs. The outcome, which is a proof of transaction validity, gets stored back on the main blockchain (or Layer 1.0). Therefore, rather than persisting the complete transaction data for each individual transaction on the blockchain, this approach conserves space by storing only a proof that verifies the validity of an entire batch of transactions executed on Layer 2.0. Example of ZK rollup is ZKSync (Matter Labs, 2023).

On the other hand, within the Optimistic rollup, it operates under the assumption that all transactions are valid by default, unless certain participants provide evidence to the contrary. Therefore, it is an optimistic approach. Examples of Layer 2.0 using Optimistic rollups are Optimism (Optimism Foundation, 2023) and Arbitrum (Offchain Labs, 2023).

In this thesis, we delve into a Layer 2.0 solution designed for Bitcoin, known as Omni Layer (Omni Layer, 2023). This aims to shed light on off-chain transaction acceleration

facilitated by miners. Unlike the batch approach used in both ZK and Optimistic rollups, Omni Layer does not batch transactions. Instead, it creates a Bitcoin transaction for every transaction issued on its layer. This enables us to analysis the extent to which miners accelerate the inclusion of these individual transactions.

CHAPTER 3

Transaction Prioritization Norms

In this chapter, we delve into the research questions, methodology, and the implications of analyzing and comprehending the norms employed by miners when prioritizing transactions for inclusion. Understanding these prioritization norms is key to enable transaction issuers to determine the appropriate fees for their transactions to be included within an expected number of blocks.

The order in which miners select transactions for inclusion have significant implications for the ultimate outcome of the transaction execution. As a result, we aim to audit the Bitcoin blockchain to address the following research questions.

► **RQ 1**: *Which transactions are allowed or transmitted over the public P2P network?* This research question pertains to the *norm III*, as described in §2.2. It assumes that transactions with fee rate below a minimum threshold are discarded and never committed to the blockchain. If miners are considering transactions with incentives below the established threshold, their view of available transactions for inclusion may differ.

► **RQ 2**: *How do miners select transactions for inclusion in a block once they enter the miners' Mempool?* This research question pertains to *norm I*, which states that miners select transactions for inclusion, from the Mempool, based *solely* on their fee rates. However, if some miners accept additional incentives to accelerate the transaction confirmation such as dark or opaque fees, the miners' view of the offered fees may vary. As a result, transaction issuers may struggle to estimate the appropriate fees needed for their transaction to be prioritized, as the fees' distribution may differ based on individual miner's view of the system.

► **RQ 3**: *In what order do miners include transactions within a block?* This question aligns with *norm II*, where higher fee rate transactions are prioritized and placed before lower fee rate transactions during the block construction. If miners misplace transactions within a block, it can result in different outcomes for those transactions when the miner's block is selected as the next block. Transaction issuers may be susceptible to well-known

Table 3.1: *Bitcoin data sets ($\mathcal{A}$ and $\mathcal{B}$) used for testing miners' adherence to transaction-prioritization norms and ($\mathcal{C}$) for investigating the behaviour of miners with respect to transaction acceleration. Child-pays-for-parent (CPFP) transactions are transactions that depend on other transactions to be included into a block.*

Attributes	*Data set $\mathcal{A}$*	*Data set $\mathcal{B}$*	*Data set $\mathcal{C}$*
Time span	Feb. 20th – Mar. 13th, 2019	Jun. 1st – 30th, 2019	Jan. 1st – Dec. 31st, 2020
Block height	563,833 – 566,951	578,717 – 583,236	610,691 – 663,904
Number of blocks	3119	4520	53,214
# of txs. issued	6,816,375	10,484,201	112,489,054
% of CPFP-txs.	26.45%	23.17%	19.11%
# of empty-blocks	38	18	240

attacks such as front-running (Daian et al., 2020; Torres et al., 2021) or even sandwich attacks (Qin et al., 2021; Zhou et al., 2023).

Addressing these research questions is crucial for understanding how miners actually prioritize transactions for inclusion. Hence, accurate knowledge of the order in which miners include transactions can assist in estimating appropriate fees that issuers need to offer to prioritize their transactions. Next, we discuss our methodology for gathering the necessary data sets and detecting accelerated or highly prioritized transactions.

Relevant publication

The results presented in this chapter have been published in (Messias et al., 2020, 2021).

3.1 Methodology

To understand the importance of transaction ordering to issuers and to investigate when and how miners violate the transaction prioritization "norms," we resort to an empirical, data-driven approach. Below, we briefly describe three different data sets that we curated from Bitcoin and highlight how we use the data sets in different analyses in the rest of this thesis. Furthermore, we present our methodology for detecting transaction acceleration, which can be applied to any fee-based incentive blockchains.

3.1.1 Data set collection

Data set $\mathcal{A}$. To check miners' compliance to prioritization norms in Bitcoin, we analyzed all transactions and blocks issued in Bitcoin over a three-week time frame from

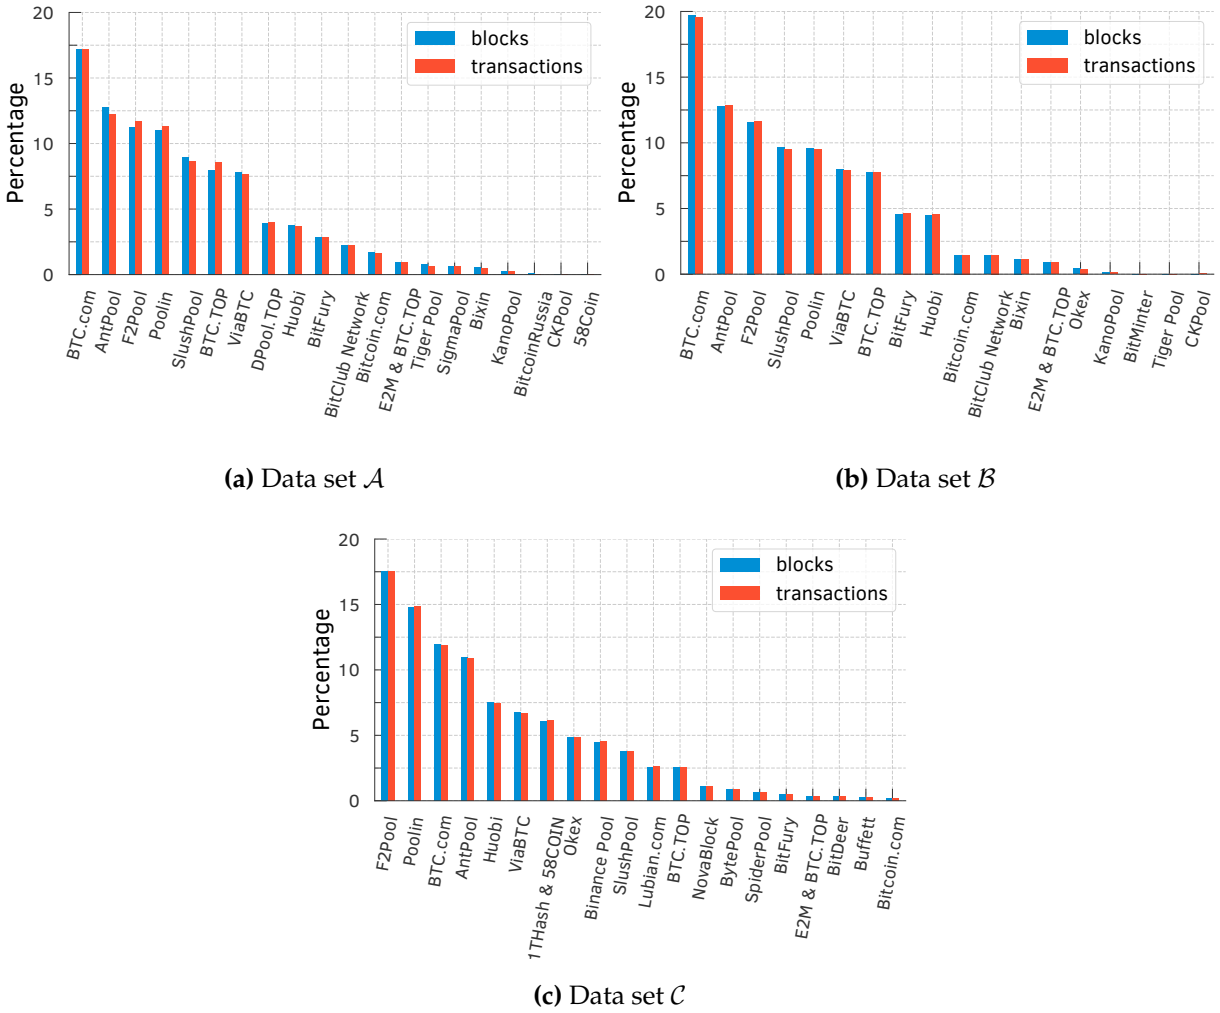

(a) Data set $\mathcal{A}$

(b) Data set $\mathcal{B}$

(c) Data set $\mathcal{C}$

Figure 3.1: Distribution of blocks mined and transactions confirmed by the top-20 MPOs in data sets $\mathcal{A}$, $\mathcal{B}$, and $\mathcal{C}$. Their combined normalized hash-rates account for 94.97%, 93.52%, and 98.08% of all blocks mined in data set $\mathcal{A}$, $\mathcal{B}$, and $\mathcal{C}$, respectively.

February 20 through March 13, 2019 (see Table 3.1). We obtained the data by running a *full* node, a Bitcoin software that performs nearly all operations of a miner (e.g., receiving broadcasts of transactions and blocks, validating the data, and re-broadcasting them to peers) except for mining. The data set also contains a set of periodic *snapshots*, recorded once per 15 seconds for the entire three-week period, where each snapshot captures the state of the full node's Mempool. We plot the distribution of the count of blocks and transactions mined by the top-20 MPOs for data set $\mathcal{A}$ in Figure 3.1a. If we rank the MPOs in data set $\mathcal{A}$ by the number of blocks (B) mined (or, essentially, the approximate hashing capacity h), the top five MPOs turn out to be BTC.com (B: 536; h: 17.18%), AntPool (B: 399; h: 12.79%), F2Pool (B: 352; h: 11.29%), Poolin (B: 344; h: 11.03%), and

SlushPool (B: 279; h: 8.94%). We use this data for checking whether miners adhere to prioritization norms when selecting transactions for confirmation or inclusion in a block.

Data set $\mathcal{B}$. Differences in configuration of the Bitcoin software may subtly affect the inferences drawn from data set $\mathcal{A}$. A full node connects to 8 peers, for instance, in the default configuration, and increasing this number may reduce the likelihood of missing a transaction due to a "slow" peer. The default configuration also imposes a minimum fee rate threshold of $1\,\mathrm{sat/B}$ (or 1 satoshi-per-byte) for accepting a transaction. We instantiated, hence, another full node to expand the scope of our data collection. We configured this second node, for instance, to connect to as many as 125 peers. We also removed the fee rate threshold to accept even zero-fee transactions. $\mathcal{B}$ contains Mempool snapshots of this full node, also recorded once per 15 s, for the entire month of June 2019 (refer Table 3.1). We notice that 99.7% of the transactions received by our Mempool were included by miners. Figure 3.1b shows the distribution of the count of blocks and transactions mined by the top-20 MPOs for data set $\mathcal{B}$. The top five MPOs are BTC.com (B: 889; h: 19.67%), AntPool (B: 577; h: 12.77%), F2Pool (B: 523; h: 11.57%), SlushPool (B: 438; h: 9.69%), and Poolin (B: 433; h: 9.58%).

Data set $\mathcal{C}$. The insights derived from the above data motivated us to shed light on the aberrant behavior of mining pool operators (MPOs). To this end, we gathered all (53,214) Bitcoin blocks mined and their 112,542,268 transactions from January 1st to December 31st 2020. These blocks also contain one Coinbase transaction per block, which the MPO creates to receive the block and the fee rewards. This data set, labeled $\mathcal{C}$, contains 112,489,054 issued transactions (see Table 3.1). MPOs typically include a *signature* or *marker* in the Coinbase transaction, probably to claim their ownership of the block. Following prior work (e.g., (Judmayer et al., 2017; Romiti et al., 2019)), we use such markers for identifying the MPO (owner) of each block. We failed to identify the owners of 703 blocks (or approximately 1.32% of the total), albeit we inferred 30 MPOs in our data set. In this thesis, we consider only the top-20 MPOs whose combined normalized hash-rates account for 98.08% of all blocks mined. Figure 3.1c shows the count of blocks mined by the top-20 MPOs according to $\mathcal{C}$. The top five MPOs in terms of the number of blocks (B) mined are F2Pool (B: 9326; h: 17.53%), Poolin (B: 7876; h: 14.80%), BTC.com (B: 6381; h: 11.99%), AntPool (B: 5832; h: 10.96%), and Huobi (B: 3990; h: 7.5%).

3.1.2 Detecting accelerated transactions.

Given the high fees demanded by acceleration services, we anticipate that *accelerated transactions would be included in the blockchain with the highest priority*, i.e., in the first few blocks mined by the accelerating miner and amongst the first few positions within the

block. We would also anticipate that *without the acceleration fee, the transaction would not stand a chance of being included in the block based on its publicly offered transaction fee*. The above two observations suggest a potential method for detecting accelerated transactions in the Bitcoin blockchain: An accelerated transaction would have a very high ***signed position prediction error (SPPE)*** (refer to §3.3.1), as its predicted position based on its public fee would be towards the bottom of the block it is included in, while its actual position would be towards the very top of the block.

To test the effectiveness of our method, we analyzed all 6381 blocks and $13{,}395{,}079$ transactions mined by BTC.com mining pool in data set $\mathcal{C}$. We then extracted all transactions with SPPE greater or equal than 100%, 99%, 90%, 50%, 1% and checked what fraction of such transactions were accelerated, according to the BTC.com transaction accelerator API (BTC.com, 2022).

3.2 Analyzing norm adherence

In this section, we analyze whether Bitcoin miners adhere to prioritization norms, when selecting transactions for confirmation. To this end, we first investigate whether transaction ordering matters to Bitcoin users in practice, i.e., *are there times when transactions suffer extreme delays and do users offer high transaction fees in such times to confirm their transactions faster?* We then conduct a progressively deeper investigation of the norm violations, including potential underlying causes, which we investigate in greater detail in the subsequent sections.

3.2.1 Does transaction ordering matter?

A congestion in the Mempool leads to contention among transactions for inclusion in a block. Transactions that fail to contend with others (i.e., win a spot for inclusion) experience inevitable delays in commit times. Transaction ordering, hence, has crucial implications for users when the Mempool experiences congestion. For instance, the Bitcoin Core code and most of the wallet software rely on the distribution of transactions' fee rates included in previous blocks to suggest to users the fees that they should include in their transactions (bitcoin.org, 2023; Coinbase, 2021; Lavi et al., 2019). Such transaction-fee predictions from any predictor, which assume that miners follow the norm, will be misleading.[9] Below, we examine whether Mempool in a real-world blockchain

[9]Coinbase, one of the top cryptocurrency exchanges, does not allow users to set transaction fees manually. Instead, it charges a fee based on how much they expect to pay for the concerned transaction, which in turn relies on miners following the norm (Coinbase, 2021).

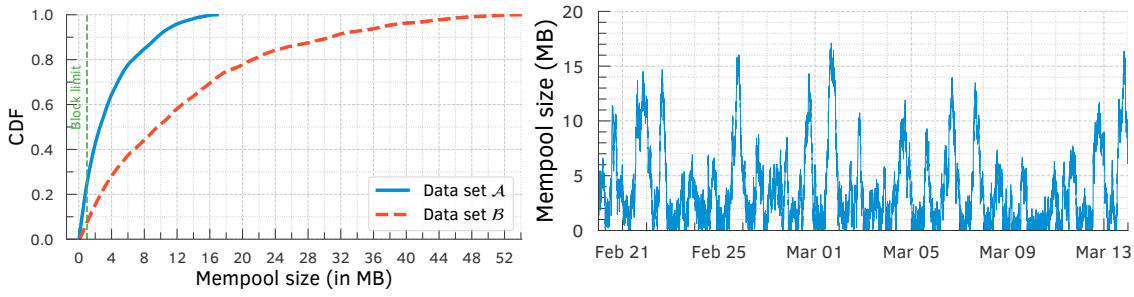

(a) **(b)**

Figure 3.2: (a) Distributions of Mempool size in both data sets $\mathcal{A}$ and $\mathcal{B}$; and (b) the size Mempool in $\mathcal{A}$ as a function of time, both indicating that congestion is typical in Bitcoin.

deployment experiences congestion and its impact on transaction-commit delays. We then analyze whether, and how, users adjust transaction fees to cope with congestion, and the effect of these fee adjustments on commit delays.

Congestion and delays

Bitcoin's design—specifically, the adjustment of hashing difficulty to enforce a constant mining rate—ensures that there is a steady flow of currency generation in the network. The aggregate number of size-limited blocks mined in Bitcoin, consequently, increases linearly over time (Figure 2.2a). Transactions, however, are *not* subject to such constraints and have been issued at much higher rates, particularly, according to Figure 2.2a, since mid-2017: 60% of all transactions ever introduced were added in only in the last 3.5 years of the nearly decade-long life of the cryptocurrency. Should this growth in transaction issued continue to hold, transactions will increasingly have to contend with one another for inclusion within the limited space (of $1\,\mathrm{MB}$) in a block. Below, we empirically show that this contention among transactions is already common in the Bitcoin network.

Using the data sets $\mathcal{A}$ and $\mathcal{B}$ (refer §3.1.1), we measured the number of unconfirmed transactions in the Mempool, at the granularity of $15\,\mathrm{s}$. Per Figure 3.2a, congestion in Mempool is *typical* in Bitcoin: During the three-week period of $\mathcal{A}$, the aggregate size of all unconfirmed transactions was above the maximum block size (of $1\,\mathrm{MB}$) for nearly 75% of the time; per data set $\mathcal{B}$ the Mempool was congested for nearly 92% of the time period. Figure 3.2b provides a complementary view of the Mempool congestion in $\mathcal{A}$, by plotting the Mempool size as a function of time. The measurements reveal a huge variance in Mempool congestion, with size of unconfirmed transactions at times exceeding 15-times the maximum size of a block. Transactions queued up during such periods of high

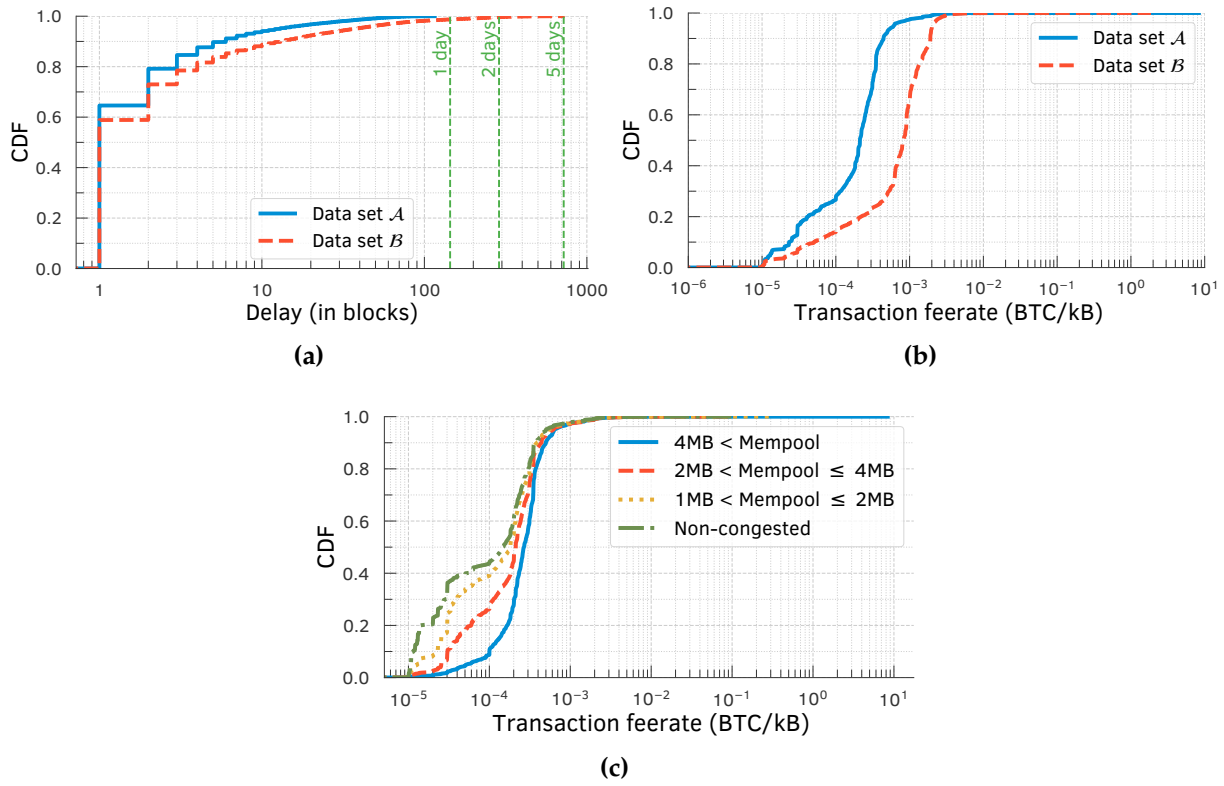

Figure 3.3: (a) Distributions of delays until transaction inclusion show that a significant fraction of Bitcoin transactions experience at least 3 blocks (or approximately 30 minutes) of delay; Distributions of fee rates for (b) all transactions and (c) transactions (in $\mathcal{A}$) issued at different congestion levels clearly indicate that users incentivize miners through transaction fees.

congestion will have to contend with one another until the Mempool size drains below 1 MB. These observations also hold in data set $\mathcal{B}$, the details of which are in §A.1.

The Mempool congestion, which in turns leads to the contention among transactions for inclusion in a block, has one serious implication for users: delays in transaction-commit times. While 65% (60%) of all transactions in data set $\mathcal{A}$ ($\mathcal{B}$) get committed in the next block (i.e., in the block immediately following their arrival in the Mempool), Figure 3.3a shows that nearly 15% (20%) of them wait for at least 3 blocks (i.e., 30 minutes on average). Moreover, 5% (10%) of the transactions wait for 10 or more blocks, or 100 minutes on average, in data set $\mathcal{A}$ ($\mathcal{B}$). While no transaction waited for more than a day in data set $\mathcal{A}$, a small percentage of transactions waited for up to five days (because of the high levels of congestion in June 2019) in data set $\mathcal{B}$.

Takeaways. Mempool is typically congested in Bitcoin. Transactions, hence, typically contend with one another for inclusion in a block. The Mempool congestion has non-trivial implications for transaction-commit times.

Transaction fee rates and delays

To combat the delays and ensure that a transaction is committed "on time" (i.e., selected for inclusion in the earliest block), users may include a transaction fee for incentivizing the miner. While the block reward from May 11, 2020 is $6.25\,\mathrm{BTC}$, the aggregate fees accrued per block is becoming considerable (i.e., 6.29% of the total miner revenue in 2020 per Table A.1 in §A.2). Prior work also show that revenue from transaction fees is clearly increasing (Easley et al., 2017). With the volume of transactions growing aggressively (Figure 2.2a) over time and the block rewards, in Bitcoin, halving every four years, it is inevitable that transaction fees will be an important, if not the only, criterion for including a transaction, leading possibly to undercutting attacks (Carlsten et al., 2016). Below, we analyze whether Bitcoin users incentivize miners via transaction fees and if such incentives are effective today.

Per Figure 3.3b the transaction fee rate of committed transactions in both data sets $\mathcal{A}$ and $\mathcal{B}$ exhibits a wide range, from 10^{-6} to beyond $1\,\mathrm{BTC/kB}$. The fee rate distributions of committed transactions also do not vary much between different mining pool operators (refer Figure A.2 in §A.3). A few transactions (0.001% in $\mathcal{A}$ and 0.07% in $\mathcal{B}$) were committed, despite offering fee rates less than the recommended minimum of 10^{-5} BTC/kB. A non-trivial percentage of transactions offered fee rates that are two orders of magnitude higher than the recommended value; particularly, in data set $\mathcal{B}$, perhaps due to the comparatively high levels of congestion (cf. Figure 3.2b and Figure A.1), 34.7% of transactions offered fee rates higher than 10^{-3} BTC/kB. Approximately 70% (51.3%) of the transactions in data set $\mathcal{A}$ ($\mathcal{B}$) offer fee rates between 10^{-4} and 10^{-3} BTC/kB, i.e., between one and two orders of magnitude more than the recommended minimum. Such high fee rates clearly capture the users' intents to incentivize the miners.

Our premise is that the (high) fee rates correlate with the level of Mempool congestion. Said differently, we hypothesize that users increase the fee rates to curb the delays induced by congestion. To test this hypothesis, we separate the Mempool snapshots (cf. §3.2.1) into 4 different bins. Each bin corresponds to a specific level of congestion identified by the Mempool size as follows: lower than $1\,\mathrm{MB}$ (*no congestion*), in $(1, 2]$ MB (*lowest congestion*), in $(2, 4]$ MB, and higher than $4\,\mathrm{MB}$ (*highest congestion*). The fee rates of transactions observed in the different bins or congestion levels, in Figure 3.3c, then validates our hypothesis: Fee rates are strictly higher (in distribution, and hence also on average) for higher congestion levels.

Figure 3.4 shows that users' strategy of increasing fee rates to combat congestion seems to work well in practice. Here, we compare the CDF of commit delays of transactions with low (i.e., less than 10^{-4} BTC/kB), high (i.e., between 10^{-4} and 10^{-3} BTC/kB),

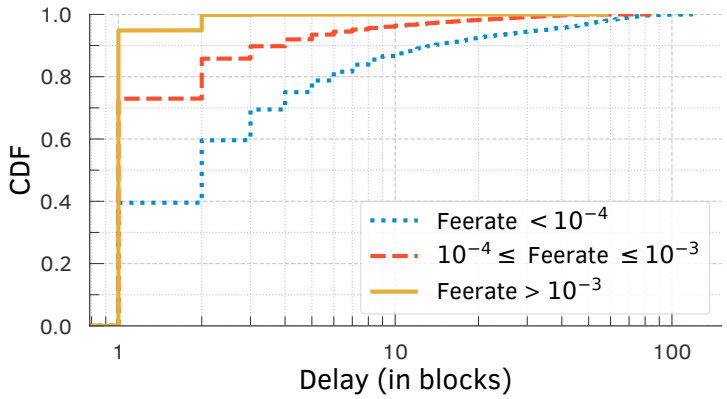

Figure 3.4: Distributions of transaction-commit delays for different fee rates for transactions in $\mathcal{A}$; incentivizing miners via fee rates works well in practice.

and exorbitant (i.e., more than 10^{-3}) fee rates, in data set $\mathcal{A}$. Similar analysis with data set $\mathcal{B}$ is provided in §A.4. We observe that an increase in the transaction fee rates is consistently rewarded (by miners) with a decrease in the commit delays. This observation suggests that, at least to some extent, miners prioritize transactions for inclusion based on fee rates or the fee-per-byte metric.

Takeaways. A significant fraction of transactions offers fee rates that are well above the recommended minimum (i.e., 10^{-5} BTC/kB or simply $1\,\mathrm{sat/B}$). Fee rates are typically higher at higher congestion levels, and reduce the commit delays. These observations suggest that users are indeed willing to spend money to decrease commit delays for their transactions during periods of congestion.

3.2.2 Do miners follow the norms?

Whether miners follow the transaction prioritization norms (as widely assumed) has implications for both Bitcoin and its users: The software used by users, for instance, assumes an adherence to these norms when suggesting a transaction fee to the user (bitcoin.org, 2023; Coinbase, 2021; Lavi et al., 2019). Deviations from these norms, hence, have far-reaching implications for both the blockchain and crucially for Bitcoin users.

Fee rate based selection when mining new blocks

Our finding above show that transactions offering higher fee rates experience lower confirmation delays suggests that miners tend to account for transaction fee rates when choosing transactions for new blocks. We now want to check, however, if transaction fee rate is the primary or the sole determining factor in transaction selection. To this end, we check our data sets for transaction pairs, where one transaction was issued earlier

and has a higher fee rate than the other, but was committed later than the other. The existence of such transaction pairs would unequivocally show that fee rate alone does not explain the order in which they are selected.

We sampled 30 Mempool snapshots, uniformly at random, from the set of all available snapshots in data set $\mathcal{A}$. Suppose that, in each snapshot, we denote, for any transaction i, the time at which it was received in the Mempool by t_i, its fee rate by f_i, and the block in which it was committed by b_i. We then selected, from each snapshot, all pairs of transactions (i, j) such that $t_i < t_j$ and $f_i > f_j$, but $b_i > b_j$. Such pairs clearly constitute a violation of the fee-rate-based transaction-selection norm.

Figure 3.5a shows a cumulative distribution of the fraction of all transaction pairs (line labeled "$\star$") violating the norm over all sampled snapshots. Across all snapshots, a small but non-trivial fraction of all transaction pairs violate the norm. One potential explanation for violations might be that the transactions are received by the mining pools in different order than the one in which our Mempool receives. To account for such differences, we tighten the time constraint as $t_i + \epsilon < t_j$ and use an ϵ of either 10 seconds or 10 minutes. Even with the tightened time constraints, Figure 3.5a shows that a non-trivial fraction of all transaction pairs violate the norm.

Another potential source of violations is Bitcoin's dependent (or, parent and child) transactions, where the child pays a high fee to incentivize miners to also confirm the parent from which it draws its inputs. This mechanism enables users to "accelerate" a transaction that has been "stuck" because of low fee (CoinStaker, 2018). As the existence of such *child-pays-for-parent (CPFP)* transactions (formally defined in §A.5) would introduce false positives in our analysis we decided to discard them. Figure 3.5b shows that the violations exist even after discarding all such dependent transaction pairs.

Fee rate based ordering within blocks

We now turn our attention to transaction ordering within individual (mined) blocks in Bitcoin. If a miner followed GBT, transactions would be ordered based on their fee rate. In this case, given the set of non-CPFP transactions $T = \{T_1, T_2, ...T_n\}$ included in a block B, we should be able to predict their position in the block by simply ordering the transactions based on their fee rate (as specified in the GBT implementation in Bitcoin Core). To quantify the deviation from the norm, we compute a measure that we call ***position prediction error (PPE)***: PPE of a block B is the average absolute difference between the predicted and the observed (actual) positions for all transactions in block B, normalized by the size of the block (n) and expressed as a percentage. More precisely,

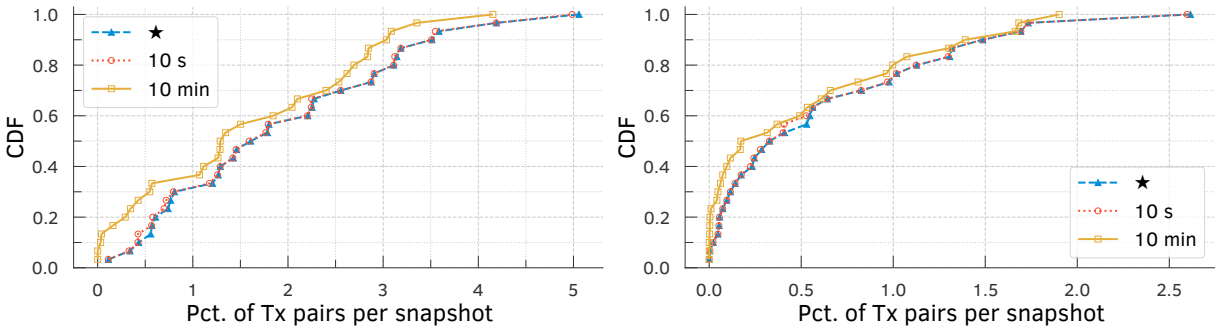

(a) All transactions **(b)** Only Non-CPFP transactions

Figure 3.5: There exists a non-trivial fraction of transaction pairs violating the norm across all snapshots, clearly indicating that miners do not adhere to the norm.

$$PPE(B) = \sum_{i=1}^{n} \frac{(|T_i^p - T_i^o)|) \cdot 100}{n}$$

where T_i^p and T_i^o are the predicted and observed positions of a transaction, respectively.

Figure 3.6a shows the cumulative distribution of PPE values for each block in our data set $\mathcal{C}$, containing $53{,}214$ blocks. 80% of the blocks have PPE values less than 4.03%. The mean PPE across all blocks is 2.65%, with a standard deviation of 2.89. Per this plot the position of a transaction within a block can be predicted with very high accuracy (within a few percentile position error), suggesting that transactions are by and large ordered within a block based on their fee rate. Figure 3.6b shows PPE values separately for each of the 6 largest mining pools in data set $\mathcal{C}$. The plots show that all mining pools by and large follow the norm, though some like ViaBTC seems to deviate slightly more from the norm compared to the other mining pools.

Fee rate threshold for excluding transactions

In their default configuration, many nodes in the Bitcoin P2P network drop (i.e., ignore) transactions that offer less than a threshold fee rate (typically, 10^{-5} BTC/kB). As miners select transactions for inclusion from their local Bitcoin P2P node, this (default) norm would result in such low-fee transactions never being included in the blockchain, even during periods of non-congestion (when blocks have spare capacity to accommodate additional transactions).

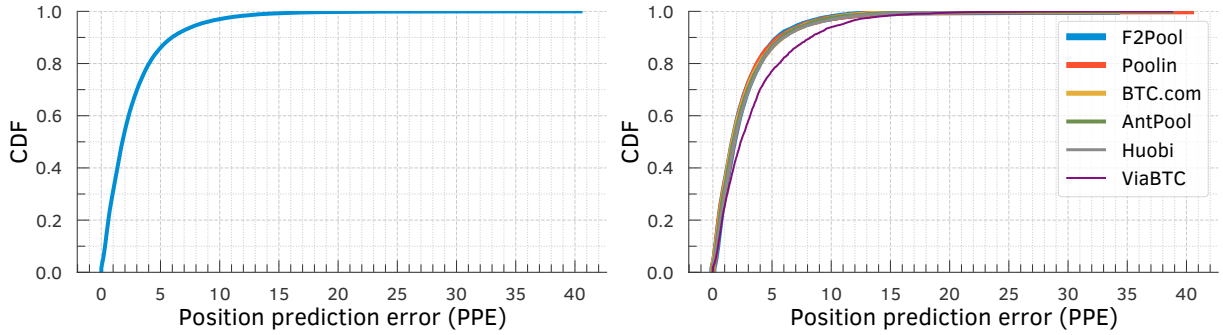

(a) Overall position prediction error **(b)** Position prediction error of the top-6 MPOs

Figure 3.6: Position prediction error (PPE). (a) There are 52,974 (99.55%) blocks with at least one non-CPFP txs. The mean PPE is 2.65%, with an std of 2.89. 80% of all blocks has PPE less than 4.03%. (b) The PPEs of blocks mined by the top-6 MPOs according to their normalized hash rate.

We collected data set $\mathcal{A}$ using a default Bitcoin node, and our node, hence, did not accept or record low-fee transactions. When gathering data set $\mathcal{B}$, however, we configured our Bitcoin node to accept all transactions, irrespective of their fee rates. In data set $\mathcal{B}$, our node, consequently, received 1084 transactions that offered less than the recommended fee rate and 489 (45.11%) of them were zero-fee transactions. From these low fee rate transactions, only 53 (4.89%) were confirmed in the Bitcoin blockchain; 9 (16.98%) were confirmed months after they were observed in our data set. In contrast, the vast majority (99.7%) of the transactions that offered greater than or equal to the recommended fee rate were all (eventually) confirmed. Interestingly, the low-fee transactions were confirmed by just three mining pools: F2Pool, ViaBTC, and BTC.com included 38, 14, and 1 low-fee transactions, respectively. Our findings suggest that while the norm of ignoring transactions offering less than the recommended fee rate is being by and large followed by all miners, a few occasionally deviate from the norm.

3.3 Investigating norm violations

Our analysis so far showed that while Bitcoin miners by and large follow transaction-prioritization norms, there are many clear instances of norm violations. Our next goal is to develop a deeper understanding of the underlying reasons or motivations for miners to deviate from the fee rate based norms, at least for some subset of all transactions. To this end, we focus our investigation on the following three types of transactions, where we hypothesize miners might have an incentive to deviate from the current norms, which are well-aligned towards maximizing their rewards for mining.

1. *Self-interest Transactions:* Miners have a vested interest in a transaction, where the miners themselves are a party to the transaction, i.e., a sender or a receiver of bitcoins. Miners may have an incentive to selfishly accelerate the commitment of such transactions in the blocks mined by themselves.

2. *Scam-payment Transactions:* Bitcoins are increasingly being used to launch a variety of ransomware and scam attacks (Lee Mathews, 2017; Sheera Frenkel and Nathaniel Popper and Kate Conger and David E. Sanger, 2020; Sheera Frenkel, Mark Scott and Paul Mozur, 2017). A scam attack involved using hijacked Twitter accounts of celebrities to encourage their followers to send bitcoins to a specific Bitcoin wallet address (Sheera Frenkel and Nathaniel Popper and Kate Conger and David E. Sanger, 2020). Given the timely and widespread coverage of this attack in popular press and other similar attacks on crowdsourced websites for reporting scam transactions (BitcoinAbuse, 2021; Whale Alert, 2021), and with governments trying to blacklist wallet addresses of entities suspected of illegal activities (Andrew Hinkes and Joe Ciccolo, 2021; Nikhilesh De, 2021), we hypothesize that some miners might decelerate or even absolutely exclude the commitment of scam-payment transactions out of fear or ethical concerns.

3. *Dark-fee Transactions:* Recently, some mining pool operators have started offering transaction acceleration services (AntPool, 2022; BTC.com, 2022; F2Pool, 2022; Poolin, 2022; ViaBTC, 2022), where anyone wanting to prioritize their transactions can pay an additional fee to a specific mining pool via a side-channel (often, the MPO's website or via a private-channel (Strehle and Ante, 2020)). Such transaction fees are "dark" or opaque to other mining pools and the public, and we hypothesize that some committed low-fee transactions might have been accelerated by using such services.

To detect whether a mining pool has accelerated or decelerated the above types of transactions, we first design a robust statistical test. Later, we report our findings from applying the test on the three types of transactions.

3.3.1 Statistical test for differential prioritization

Our goal here is to propose a robust statistical test for detecting whether a given mining pool m is prioritizing a given set of committed transactions c *differently* than all other miners. The basic idea behind the statistical test is as follows. Suppose a mining pool is accelerating (decelerating) transactions in set c. In that case, these transactions will have

a disproportionately high (low) chance of being included in blocks mined by this mining pool compared to the mining pool's hashing power (or rate).

Test for differential transaction acceleration

Consider a miner m with normalized hash rate $h = \theta_0$ (estimated as fraction of blocks mined by m). Assume that we are given a set of transactions, denoted as c-transactions (for committed transactions), for which we wish to test whether miner m is treating them preferentially.

To test whether m is prioritizing c-transactions, we look at all blocks that include at least one c-transaction, call them c-blocks. Suppose that there are y such blocks. If m is not prioritizing c-transactions, then a fraction θ_0 of all c-blocks should be m-blocks (i.e., mined by m); if m is prioritizing c-transactions (compared to other miners) then the fraction will be higher. We want to test whether the true fraction θ is indeed θ_0 or is higher. We formalize this as follows: We assume that each c-block has a probability θ to be an m-block and do the following test.

$$H_0 : \theta = \theta_0$$
$$H_1 : \theta > \theta_0.$$

Assuming that the observed number of c-blocks that are mined by m is x, the p-value of the test is

$$p = Pr(B \geq x),$$

where B is a binomial distribution of parameter θ_0 and y, that is

$$p = \sum_{k=x}^{y} \binom{y}{k} \theta_0^k (1 - \theta_0)^{(y-k)}.$$

We may fix the size of the test (i.e., the maximal probability of type I error that corresponds to rejecting H_0 when H_0 is true) to $\alpha = 0.01$. Then H_0 should be rejected whenever $p < \alpha$. The smaller p, the higher the confidence in rejecting H_0, that is declaring that m prioritizes c-transactions.

The above test is relative in the sense that we can only detect if a miner treats c-transactions more preferentially than the rest of the miners. This test cannot conclude on whether it is the miner accelerating the c-transactions (relative to their deserved, i.e., fee rate based, priority) or the rest of the miners are decelerating them. So, we look at additional empirical evidence from the position of the c-transactions within the c-blocks

that include them. Specifically, given the set of c-transactions $\{c_1, c_2, ...c_n\}$ committed by a miner m, we compute a measure that we call ***signed position prediction error (SPPE)*** as the average signed difference between the predicted and observed positions (measured as percentile rank) for all c-transactions within the blocks committed by m. More precisely,

$$SPPE(m) = \frac{\sum_{i=1}^{n}(c_i^p - c_i^o) \cdot 100}{n}$$

where c_i^p and c_i^o are the predicted and the observed (percentile rank) positions, respectively, of transaction c_i within the blocks committed by m.

Test for differential transaction deceleration

While the previous test checks for prioritization (or acceleration), one may also want to test for deceleration. To that end, a symmetric test can be used. Specifically, with the previous notation, the test would be

$$H_0 : \theta = \theta_0$$
$$H_1 : \theta < \theta_0;$$

and its p-value would be

$$p = Pr(B \leq x),$$

where B is a binomial distribution of parameter θ_0 and y, that is

$$p = \sum_{k=0}^{x} \binom{y}{k} \theta_0^k (1 - \theta_0)^{(y-k)}.$$

Scaling the tests

While we did not face them in this thesis, our test may have two limitations when scaling to large time windows and/or large numbers of transactions.

First, it may become difficult to compute the p-value from the binomial distribution for large values of y. In such cases, we can use the following approximation for our analysis: If y is large enough and θ_0 is not close to zero or one (i.e., x and $y - x$ are large enough), the binomial distribution of parameters θ_0 and y is well approximated by the normal distribution with mean $y\theta_0$ and variance $y\theta_0(1 - \theta_0)$. Hence, the p-value for the

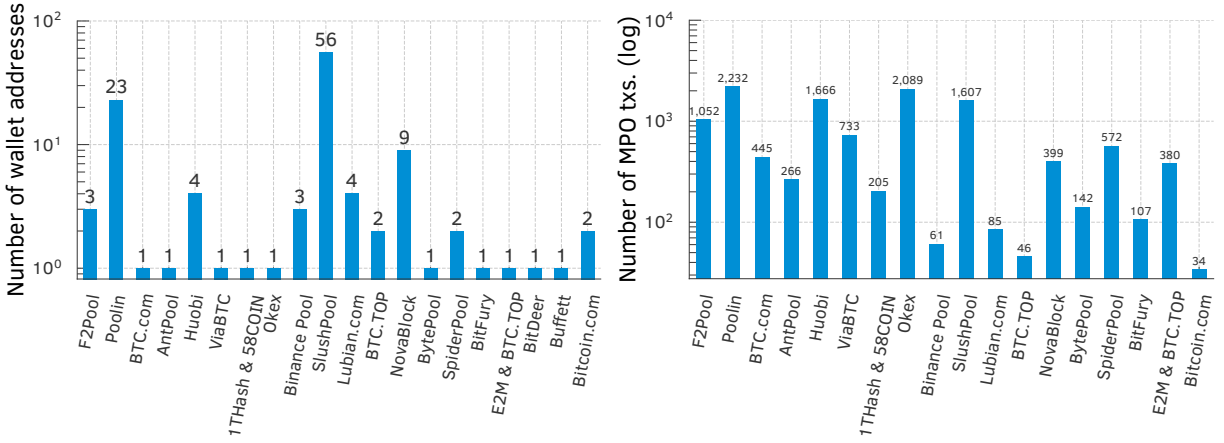

(a) Number of MPOs' wallet addresses **(b)** Number of MPOs' transactions

Figure 3.7: (a) Distribution of the number of wallet addresses in data set $\mathcal{C}$ used by each of the top-20 MPOs to receive its block rewards; SlushPool and Poolin, for instance, used 56 and 23 distinct wallet addresses, respectively. (b) The counts of inferred MPO transactions; in total, 12,121 transactions were inferred as MPOs' transactions, which corresponds to 0.011% of the total issued transactions recorded in the Bitcoin blockchain. Poolin has the majority with 2232 (18.41%), followed by Okex with 2089 (17.24%) and Huobi with 1666 (13.74%) transactions. BitDeer and Buffett have the same wallet address as BTC.com and Lubian.com, respectively. We count the addresses of the former as belonging to the latter.

acceleration test can be computed as,

$$p \simeq \Phi\left(\frac{x - y\theta_0}{\sqrt{y\theta_0(1-\theta_0)}}\right),$$

where Φ is the CDF of a standard normal random variable. A similar approximation can be done for the deceleration test.

Second, the hash rates of miners in our p-value test are assumed to be more or less constant (i.e., θ_0 is a constant), which is not the case (per Figure B.1 and Figure B.2 in §B.3). This assumption is a limitation of our test as, in reality, hash rates of miners may vary over time, particularly over large time windows. In such situations, our test results may be affected, particularly when the arrival times of transactions are not regularly spread over the time window of our analysis. We address this issue by confirming the results of the p-value test through the SPPE-test, which is not affected by variable hash rates. It is possible, however, to alleviate this limitation of our analysis. One natural way is to divide the total time window into multiple windows such that the hash rate is

Table 3.2: *Differential prioritization of self-interest transactions.*

Transactions of ...	*mining pool (m)*	*norm. hash rate (θ_0)*	*x*	*y*	*p-value (accel.)*	*p-value (decel.)*	*% SPPE (m)*
F2Pool	F2Pool	0.1753	466	839	**0.0000**	1.0000	**78.5494**
ViaBTC	ViaBTC	0.0676	412	720	**0.0000**	1.0000	**98.9175**
1THash & 58Coin	ViaBTC	0.0676	34	201	**0.0000**	1.0000	**81.4516**
	1THash & 58Coin	0.0611	39	201	**0.0000**	1.0000	**96.9143**
SlushPool	SlushPool	0.0375	214	1343	**0.0000**	1.0000	**88.3082**
	ViaBTC	0.0676	140	1343	**0.0000**	1.0000	**45.1523**

more or less constant in those shorter time windows; and compute p-values in each time window. We can then combine the obtained p-values using Fisher's method (Fisher, 1992; Mosteller and Fisher, 1948). We leave the investigation of such extended test procedures to future work, when they might be needed.

3.3.2 Self-interest transactions

To identify transactions where a mining pool is a sender or receiver of transactions, we first need to identify Bitcoin wallets (addresses) that belong to mining pools. In Bitcoin, whenever a mining pool discovers a new block, it specifies a wallet address to receive the mining rewards. This mining pool address is included in the Coinbase transaction (refer §2.1) that appears at the start of every block. In our data set $\mathcal{C}$, we gathered all the wallet addresses used by the top-20 mining pools to receive their rewards. For each mining pool, we then retrieved all committed transactions, in which coins were sent from the mining pool's wallet. Figure 3.7 shows the statistics for the mining pool wallets and the transactions spending (sending) coins from (to) the wallets, for each of the top-20 mining pools in data set $\mathcal{C}$. We found hundreds or thousands of self-interest transactions for most of the mining pools.

Acceleration of self-interest transactions

For self-interest transactions belonging to each of the top-20 mining pools, we separately applied our statistical test to check whether any of the top-10 mining pools (that mined at least 4% of all mined blocks in data set $\mathcal{C}$) are preferentially accelerating or decelerating the transactions. In Table 3.2, we report the statistics from our test for mining pools that were found to preferentially treat transactions belonging to their own or other mining pools. Strikingly, Table 3.2 shows that 4 out of the top-10 mining pools namely, F2Pool, ViaBTC, 1THash & 58Coin, and SlushPool *selfishly accelerated* their own transactions, i.e., coin transfers from or to their own accounts (p-value for acceleration test is less

Table 3.3: *Differential prioritization of scam-payment transactions*

mining pool (*m*)	*norm. hash rate* (θ_0)	*x*	*y*	*p-value* (*accel.*)	*p-value* (*decel.*)	*% SPPE* (*m*)
Poolin	0.1528	10	53	0.2856	0.8227	−3.9787
F2Pool	0.1450	10	53	0.2323	0.8629	0.8735
BTC.com	0.1147	9	53	0.1483	0.9233	−2.8333
AntPool	0.1093	4	53	0.8450	0.2989	31.5000
Huobi	0.0955	1	53	0.9951	0.0323	−1.6428
Okex	0.0698	3	53	0.7248	0.4890	−5.0000
1THash & 58COIN	0.0684	8	53	0.0268	0.9907	−0.5000
Binance Pool	0.0590	3	53	0.6120	0.6180	−2.6000
ViaBTC	0.0552	1	53	0.9507	0.2020	−4.0000

than 0.001). Equally, if not more interestingly, Table 3.2 shows collusive behavior among mining pools. Specifically, it shows that transactions issued by 1THash & 58Coin and SlushPool were *collusively accelerated* by ViaBTC (p-value for acceleration test is less than 0.001). That these mining pools were accelerating the transactions is further confirmed by the SPPE measure, which clearly shows that in each of the above cases, the self-interest transactions were also being included within the blocks ahead of other higher fee rate transactions.

3.3.3 Scam-payment transactions

Next, we investigate whether any mining pool attempted to decelerate or exclude scam-payment transactions.

On July 15, 2020, multiple celebrities' accounts on Twitter fell prey to a scam attack. The scammers posted the message that anyone who transferred bitcoins to a specific wallet will receive twice the amount in return (Sheera Frenkel and Nathaniel Popper and Kate Conger and David E. Sanger, 2020). In response, several people sent, in total, 12.87051731 bitcoins—then worth nearly $142{,}000$ (USD)—to the attacker's wallet via 386 transactions, which were confirmed across 53 blocks by 12 miners.

To examine the miners' behavior during this scam attack, we selected all blocks mined from July 14 to August 9, 2020 (i.e., 3697 blocks in total, containing $8{,}318{,}621$ issued transactions as described in §A.6) from our data set $\mathcal{C}$. Once again, we applied our statistical test to check whether any of the top-9 mining pools (that mined at least 5% of all mined blocks from this data) are preferentially accelerating or decelerating the transactions. Table 3.3 shows the test statistics. Interestingly, we find no statistically significant evidence (i.e., p-value less than 0.001) of scam-payment acceleration or deceleration across all top mining pools. Looking at SPPE measure across the mining pools, we find no evidence of mining pools (other than AntPool) preferentially ordering

Table 3.4: *Scam types and their occurrences in wallets and transactions.*

Scam type	*# of txs.*	*% of txs.*	*# of wallets*	*% of wallets*
Sextortion	2,656	40.79	478	86.59
Terrorism	1,093	16.79	1	0.18
Dark Web Shop	1,019	15.65	8	1.45
Fake Giveaway	739	11.35	18	3.26
Fake Exchange	235	3.61	3	0.55
Other	218	3.35	16	2.90
Malware	195	2.99	4	0.72
Fake Investment	164	2.52	6	1.09
Ransomware	139	2.13	17	3.08
Ponzi Scheme	53	0.82	1	0.18
Total	***6,511***	***100***	***552***	***100***

the scam-payment transactions within blocks. In short, our findings show that most mining pool operators today do not distinguish between normal and scam-payment transactions.

Inferring other scam payment transactions from crowded source data

We also want to investigate whether other types of scams payments have been recorded on the Bitcoin blockchain.

To conduct this experiment, we gathered data from the Bitcoin Abuse (Bitcoin Abuse, 2020) crowdsource platform. This platform allows users to report Bitcoin wallet addresses that they suspect are associated with scams. From May 16, 2017, to October 15, 2020, we gathered a total of 186,731 reports from users. These reports identified 54,032 unique wallet addresses that could potentially be linked to scams. The monthly trend of scam reports is illustrated in Figure 3.8a. Notably, there is an increase in the number of reports starting from July 2018. The reports predominantly originated from users in the US, followed by users in the UK (refer Figure 3.8b).

To simplify our analysis, we focused on transactions that occurred during the year 2018.[10] However, out of the total reported wallets, only a small portion, 1169 (2.16%) wallets, were found in our Bitcoin 2018 data set. On average, each wallet address was reported 3.46 times, with a std. of 14.41. Both the minimum and the median number of reports for a wallet address were 1, while the maximum number of reports reached 950.

To ensure accuracy in our analysis, we complemented the data from Bitcoin Abuse with information from another platform called Scam Alert (Scam Alert, 2020). Scam Alert

[10]Due to limitations in data gathering, we considered transactions issued in 2018. This forms a subset of our data set $\mathcal{D}$ introduced in §4.1.

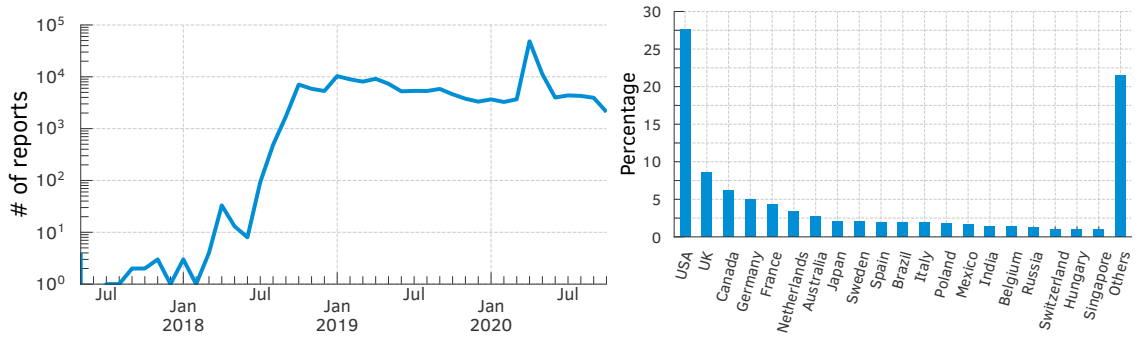

(a) Monthly reports **(b)** Reports per country

Figure 3.8: (a) Distribution of the number of reports per month. (b) USA is the country with more reports accounting for 27.6% of all reports available followed by UK and Canada.

verifies the wallet addresses reported by users to determine if they are indeed associated with scams. Initially, we had a set of 1169 identified wallets reported by Bitcoin Abuse. Upon cross-checking with Scam Alert, we discovered that 100 were not involved with scams. Out of the remaining 1069 wallet addresses, 473 were confirmed to be scams, 79 were labeled as probable scams, and 517 were pending review at the time of our analysis. Thus, for our analysis, we only considered wallet addresses confirmed or labeled as probable scams by Scam Alert. This resulted in a final scam wallet data set containing 552 unique wallet addresses.

Table 3.4 provides information about the number of wallet addresses and the number of transactions that sent coins to at least one of the scam wallets during 2018. We identified several types of scams using the Scam Alert data set, with Sextortion comprising 40.79% of the total, followed by Terrorism at 16.79%, and Dark Web Shop at 15.65%. Together, these three types of scams accounted for 73.23% of the total scams identified in our study. Additionally, Table 3.5 shows the top-20 most frequently used wallet addresses for scam payments. Notably, the wallet 1LaN$\cdots$Qctq has been associated with a terrorist organization using it to collect donations (Huillet, 2020). This specific wallet address was involved in a total of by 1093 transactions.

Moreover, we observed that despite these scam wallets being publicly available, most of the mining pools still included them in their blocks. Figure 3.9 shows the distribution of blocks and transactions associated with scam payments that were included by each of the mining pools. BTC.com included alone 20.09% out of the 6511 scam transactions in 2018.

Table 3.5: *Top-20 most used wallets for scam payments.*

Wallet address	*Scam type*	*# of txs.*
1LaNXgq2ctDEa4fTha6PTo8sucqzieQctq	Terrorism	1,093
167uU5Q3cCPijsfwmmH6ZAQj8yYxQdmzoN	Dark Web Shop	359
17v1cviCPNuGY73wNGvatS3CEZzrcPnXPy	Fake Giveaway	254
1Gs7Aztizk2rNNSE6AbpK4K7yAFTCZKV9a	Dark Web Shop	251
1EU1Ly84tYpTCcjWtvF4tYosRNN2xYYSGF	Dark Web Shop	207
15ESgUNQ9Hgn2h2FDMJi9NwE4g7ZWRAGJE	Fake Giveaway	170
3Lo4nDzH7Bi572T7t8pQGU2Ax9jVymHeC6	Fake Exchange	137
13hjTSbwVJfsDgL3qaQSu3fs2qmHQCHRXT	Sextortion	131
1Hy6BcTtNwrCLQK8ViEP742jRgx8Zpfoja	Dark Web Shop	115
343CXYVBKXT2VgELCdjEeMyPpfiKwkzUNg	Other	103
3L5o1AHLTKUeJDF8U2s5dgQCwoGknVyycn	Fake Giveaway	98
16EegrNMdZ9Rxku6Za5neEFjMW57wkQr1S	Malware	89
1C4SvJQexhAEZzm3f6E6PMQT2xWtJdKKvp	Fake Exchange	80
1HYoMM6mfFiDvkRe5z9RsSo3sugnqaDps3	Fake Investment	77
1B7aczSxaMbRsPJXx22TP1foaHQ6FENwTA	Fake Giveaway	62
3NYKHbX3zRbcZeASxjZmb4bpF8kZytnuvi	Malware	54
1JTtwbvmM7ymByxPYCByVYCwasjH49J3Vj	Sextortion	54
16JL8g7QQthYorTCkjJNE7Yhm7M3DyVyNZ	Ponzi Scheme	53
1GL9JtXPRTPetxgiJ8UcgrEECp12spD4tt	Sextortion	52
122wvcbWhBux5jcf2iyzFLmW7Jex7iSpef	Sextortion	49

3.4 Dark-fee transactions

We refer to transactions that offer additional fees to specific mining pools through an opaque and non-public side-channel payment as dark-fee transactions. Many large mining pool operators allow such side-channel payments on their websites for users wanting to "accelerate" the confirmation of their transactions, especially during periods of congestion. Such private side-channel payments that hide the fees a user pays to miners from others have other benefits for the users (AntPool, 2022; BTC.com, 2022; F2Pool, 2022; Poolin, 2022; SparkPool, 2021). One well-known advantage is, for instance, avoiding the fee rate competition in transaction inclusion, particularly during periods of high Mempool congestion; private side-channel payments would reduce a user's transaction cost volatility and curb front-running risks (Daian et al., 2020; Eskandari et al., 2020; Strehle and Ante, 2020). We use the data set $\mathcal{C}$ to first investigate how such transaction acceleration services work and later propose a simple test for detecting accelerated transactions in the Bitcoin blockchain.

Investigating transaction acceleration services

We examined transaction acceleration services offered by 5 large Bitcoin mining pools namely, BTC.com (BTC.com, 2022), AntPool (AntPool, 2022), ViaBTC (ViaBTC, 2022), F2Pool (F2Pool, 2022), and Poolin (Poolin, 2022). Specifically, we queried BTC.com for

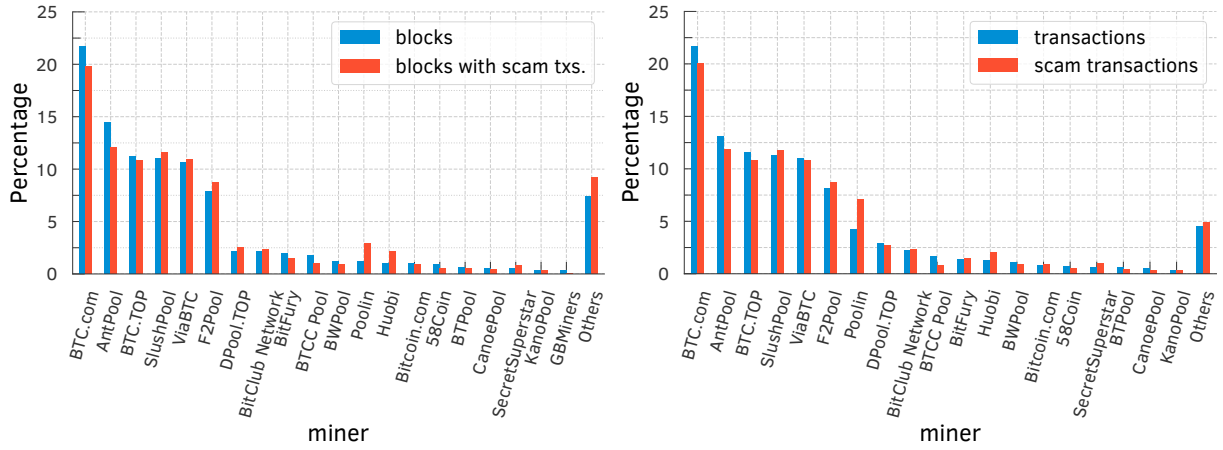

(a) Distribution of blocks

(b) Distribution of transactions

Figure 3.9: Distribution of (a) blocks mined per each mining pool in comparison to the fraction of blocks that contains at least one scam transaction; and (b) transactions included by each mining pool in comparison to their share of scam transaction inclusion. BTC.com included 20.09% out of the 6511 scam transactions in 2018.

the prices of accelerating all transactions in a real-time snapshot of the Mempool in data set $\mathcal{C}$ (see §A.7). We found that the dark fee requested by BTC.com to accelerate each transaction is so high that if it was added to the publicly offered transaction fee, the resulting total fee rate would be higher than the fee rate offered by any other transaction in the Mempool snapshot. Put differently, had users included the requested acceleration fees in the publicly offered fee when issuing the transaction, every miner would have included the transaction with the highest priority.

The above observation raises the following question: *why would rational users offer a dark fee to incentivize a subset of miners to prioritize their transaction rather than publicly announce the fee to incentivize all miners to prioritize their transaction?* One potential explanation could be that as payment senders determine the publicly offered transaction fees, payment receivers might wish to accelerate the transaction confirmation by offering an acceleration fee. Another explanation could be that the user issuing the transaction might want to avoid revealing the true fees they are willing to offer publicly, to avoid a fee rate battle with transactions competing for inclusion in the chain during congestion. Opaque transaction fees can reduce transaction cost volatility, but they may also unfairly bias the level playing field amongst user transactions attempting to front-run one another (Daian et al., 2020; Strehle and Ante, 2020).

On the other hand, every rational mining pool has clear incentives to offer such acceleration services. They receive a very high fee by mining the accelerated transaction.

Table 3.6: *[Data set $\mathcal{C}$] For an SPPE $\geq$ 99%, we observe that 64.98% of BTC.com transactions were accelerated; the fourth column values are derived by dividing the values in the second with those in the third. The number of accelerated transactions decreases to 18.12% for an SPPE $\geq$ 90% and to 1.06% for an SPPE $\geq$ 50%.*

SPPE (≥)	***# txs***	***# acc. txs***	***% acc. txs***
100%	628	464	73.89
99%	1108	720	64.98
90%	5365	972	18.12
50%	95,282	1007	1.06
1%	657,423	1029	0.16

Better still, they keep the offered fee, even if the accelerated transaction were mined by some other miners.

Detecting accelerated transactions

Given the high fees demanded by acceleration services, we anticipate that *accelerated transactions would be included in the blockchain with the highest priority*, i.e., in the first few blocks mined by the accelerating miner and amongst the first few positions within the block. We would also anticipate that *without the acceleration fee, the transaction would not stand a chance of being included in the block based on its publicly offered transaction fee*. The above two observations suggest a potential method for detecting accelerated transactions in the Bitcoin blockchain: An accelerated transaction would have a very high ***signed position prediction error (SPPE)***, as its predicted position based on its public fee would be towards the bottom of the block it is included in, while its actual position would be towards the very top of the block.

To test the effectiveness of our method, we analyzed all 6381 blocks and 13,395,079 transactions mined by BTC.com mining pool in data set $\mathcal{C}$. We then extracted all transactions with SPPE greater or equal than 100%, 99%, 90%, 50%, 1% and checked what fraction of such transactions were accelerated. Given a transaction identifier, BTC.com's acceleration service (BTC.com, 2022) allows anyone to verify whether the transaction has been accelerated. Our results are shown in Table 3.6. We find that more than 64% of the 1108 transactions with SPPE greater or equal than 99% were accelerated, while only 1.06% of transactions with SPPE greater or equal than 50% were accelerated. In comparison, we found no accelerated transactions in a random sample of 1000 transactions drawn from the 13,395,079 transactions mined by BTC.com. Our results show that large values of SPPE for confirmed transactions indicate the potential use of transaction acceleration services. In particular, a transaction with SPPE $\geq$ 99% (i.e., a transaction that

is included in the top 1% of the block positions, when it should have been included in the bottom 1% of the block positions based on their public fee rate) has a high chance of being accelerated.

3.4.1 Layer 2.0 transactions

Bitcoin offers a unique operation code, or simply opcode (Bitcoin Wiki, 2023b), known as *OP_Return*, which allows anyone to write arbitrary data to the Bitcoin blockchain. This opcode was introduced with the release of Bitcoin Core v0.9.0 in 2014. The primary purpose of OP_Return is to enable participants to mark a transaction output as invalid or to store additional data on the blockchain. By using OP_Return, the Bitcoin blockchain can also serve as a Layer 1.0 solution for Layer 2.0 applications like the Omni Layer Protocol (Omni Layer, 2023). However, this usage can lead to a situation where the true value of a transaction transfer (in the case of Bitcoin) might not be directly visible on the blockchain. Instead, the actual value of a transaction can be determined by interpreting the arbitrary data stored within it. To better understand these transactions and their purposes, we aim to parse the data written to the blockchain, using the data specification from (Omni Layer, 2020), and investigate whether transactions with arbitrary data have been accelerated or utilized for specific reasons.

To this end, we considered a 3-year Bitcoin data set named data set $\mathcal{D}$ (refer §4.1), consisting of a total of $313{,}737{,}341$ transactions ($313{,}575{,}387$ issued transactions and $161{,}954$ coinbase transactions), we observed that $42{,}994{,}249$ transactions (13.70%) contained at least one OP_RETURN opcode in their list of transaction output. Focusing solely on the issued transactions, $42{,}832{,}713$ transactions (13.66%) included at least one OP_RETURN opcode. Regarding the coinbase transactions, from the total of $161{,}954$ blocks, a majority of $161{,}536$ coinbase transactions (99.74%) contained at least one OP_RETURN opcode. This indicates that miners have also been actively including arbitrary data in the Bitcoin blockchain.

Moreover, we identified $17{,}993{,}300$ transactions associated with the Omni Layer Protocol (Omni Layer, 2023), averaging 111 transactions per block. These Omni transactions accounted for 5.74% of all Bitcoin issued transactions and 42% of all transactions involving OP_RETURN opcodes. Notably, a significant portion (97.27%) of these Omni transactions were related to the Tether USDT token (Tether, 2023), a stablecoin pegged to the US dollar.

We found 1805 OP_RETURN transactions among a set of $14{,}104$ *accelerated* transactions in data set $\mathcal{D}$, using the methodology discussed in the previous section. Out of these accelerated transactions, 1740 belong to the Omni Layer Protocol, with 1739 of

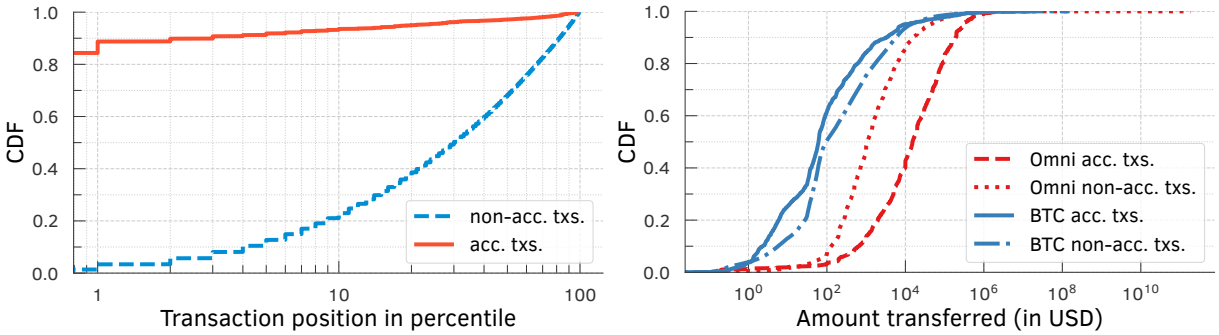

(a) Transaction position percentile

(b) Amount of US dollar transferred

Figure 3.10: Cumulative distribution function for (a) transaction position percentile: 84.30% of Omni transactions were positioned right at the top of their respective blocks. This means they were the very first transactions to be included in those blocks; (b) comparison of Omni transfers and Bitcoin value transfers: The amount transferred in Omni to the corresponding value in Bitcoin for accelerated transactions was 259.97 times higher than the value announced in the Bitcoin blockchain.

them being related to the Tether token, and the remaining one to the Omni token. These Omni accelerated transactions were consistently placed right at the top of each block, on average within the top 3.4%, with a std. of 13.89% and a median of 0%. This indicates that they were the first issued transactions in their respective blocks. In comparison, non-accelerated transactions were usually placed within the top 36.97% of the block, with a std. of 28.64% and a median of 30%. Overall, we observed that 84.30% of all Omni transactions were included at the top of their respective blocks, being the first transactions to be included (refer Figure 3.10a).

Due to the opacity of the true value of Omni transactions, as it requires interpreting the arbitrary data stored in each transaction, we parsed the data using the protocol transaction specification from (Omni Layer, 2020). Then, we compared the values transferred in Bitcoin transactions based on the transaction output value with the values stored in the arbitrary data belonging to Omni. The Bitcoin prices were converted to US dollars, considering the exchange rate at the time the transactions were included in the block, obtained from the Yahoo Finance BTC-USD feed (Yahoo Finance, 2023a). As shown in Figure 3.10b, Omni transactions tend to transfer much higher values among users compared to the values seen in the Bitcoin transaction outputs. For instance, on average, the values transferred in Tether on the Omni layer were 501,600.69 USD with a median of 1053 USD, while the values seen in the Bitcoin blockchain averaged 167,518.43 USD with a median of 84.81 USD. This means that the values transferred in Omni were

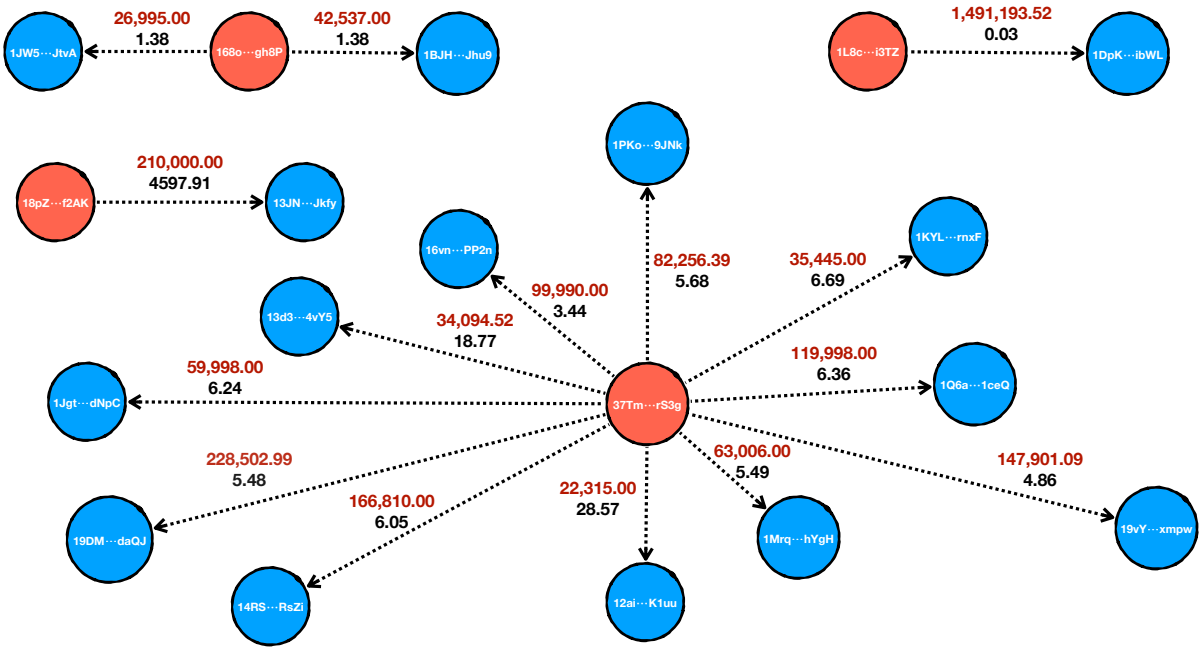

Figure 3.11: Comparison between the value transferred (in USD) in accelerated Omni transactions (shown in the top red color) and the values from the Bitcoin blockchain (in the bottom black color) for transactions included in block 550,912. Each edge in the graph corresponds to a single BTC transaction. Transaction values available in the Bitcoin blockchain appear to be relatively low. However, in contrast, these transactions in the Omni Layer are notably high-value transactions. Nodes in blue indicate receivers, while nodes in red indicate senders.

almost 3 times higher on average and 12.42 times higher in the median compared to Bitcoin transactions.

We also analyzed the results based on acceleration. For half of the evaluated transactions, the accelerated Omni transaction transferred at least 15,274.34 USD, while the value reported in Bitcoin was only 58.76 USD. This indicates that the value transferred through Omni was at least 259.97 times higher for half of the transactions evaluated. In Figure 3.11, we showcase all accelerated transactions included in block 550,912, totaling 15 transactions, which belong to Omni. We also highlight the senders and receivers involved in these transactions. Comparing the value transferred in Bitcoin, we found that the average value was 313.22 USD with a std. of 1185.35 USD and a median of 5.68, ranging from 0.03 to 4597.91 USD. In contrast, on the Omni network, the average value transferred was 188,736.17 USD with a std. of 366,405.31 USD and a median of 82,256.39 USD, ranging from 22,315.00 USD to 1,491,193.52 USD.

3.5 Concluding remarks

In this chapter, we conducted an extensive empirical audit of the miners' behavior to check whether they adhere to the established norms. At a high level, our findings reveal that transactions are primarily prioritized based on assumed norms. However, our analysis also uncovers evidence of a significant number of confirmed transactions that violate these priority norms. An in-depth investigation of these norm violations uncovered many highly troubling misbehavior by miners. Our results demonstrate that large values of SPPE for confirmed transactions indicate the potential use of transaction acceleration services. In particular, a transaction with SPPE $\geq 99\%$ has a high chance of being accelerated.

Strikingly, we show that 4 out of the top-10 mining pools namely, F2Pool, ViaBTC, 1THash & 58Coin, and SlushPool *selfishly accelerated* their own transactions. Furthermore, we uncover instances of collusive behavior between mining pools. Our results are supported by the SPPE metric, which indicates that self-interest transactions were also being included in the blocks ahead of other higher fee rate transactions.

In summary, our findings strongly suggest that several large mining pools tend to give special treatment to transactions that benefit them directly. This included transactions involving payments to or from wallets owned by the mining pool. Some even *collude* with other large mining pools to prioritize their transactions. Additionally, a number of significant large mining pools accept additional *dark (opaque) fees* to accelerate transactions via non-public side-channels (e.g., their websites). This practice of dark-fee transactions contradicts a fundamental, albeit unstated, assumption in blockchain systems: *that the confirmation fees offered by transactions are transparent and equal to all miners*.

In the following chapter, we will explore the implications of the lack of transparency into both the prioritization of transactions and the content of transactions.

Chapter 4

Transaction Prioritization and Contention Transparency

In this chapter, we discuss the implications of our findings regarding the lack of transparency in transaction contention and prioritization. We also argue why our findings and implications would be relevant even in the face of recent changes to blockchain protocols, e.g., Ethereum Improvement Protocol (EIP) 1559 (Buterin et al., 2019a) and the Ethereum Paris Network Upgrade (a.k.a. the Merge) (Ethereum Foundation, 2022b).

The lack of transparency in both transaction contention and prioritization has not been thoroughly explored in the literature, resulting in a limited understanding of its implications. This thesis aims to address this gap by investigating the following research questions.

► **RQ 1**: *To what extent are private relay networks prevalent in facilitating transactions prioritization?* Given the rise of transaction attacks like frontrunning and sandwich attacks, it is reasonable to consider that transactions issuers may prefer sending their transactions privately to the miners to avoid such attacks. This research question aims to explore the current prevalence or widespread adoption of private relay networks as a means for issuers to achieve their goal of protecting their transactions. However, we also consider the potential downsides of private transaction inclusion, particularly in terms of the lack of transparency in transactions contention and prioritization. Thus, we investigate this research question to assess the overall benefits and drawbacks of these private relay networks to the broader blockchain ecosystem.

► **RQ 2**: *Are private transactions preferentially treated by miners?* This research question aims to investigate if miners provide preferential treatment to private transactions. Private transactions offer guaranteed payments (or fees), whereas fees for publicly issued transactions are available to any miners willing to include them. We hypothesize that miners would likely offer preferential treatment for private transactions due to their guaranteed payment nature. To address this research question, we conducted an active

experiment to assess whether miners exhibit preferential treatment towards private transactions in the context of Ethereum blockchain.

► **RQ 3**: *To what extent do transaction bundling practices occur? Do they include public transactions?* This research question focuses on exploring the frequency and characteristic of the transaction bundling practices, particularly the inclusion of public transactions, in order to exploit MEV opportunities. Arbitrageurs may have an incentive to create Flashbots (Flashbots, 2022a) bundles that combine both private and public transactions to capture the financial opportunity derived from the execution of public transactions. It is worth noting that the fees associated with private transactions remain private to the relay and the miner until the transactions within the bundle are included in a block. Hence, we investigate the types of public transactions included in these bundles and the specific contracts they call. Additionally, we also investigate the revenues earned by miners from accepting these bundles. This analysis is crucial for advancing the transparency goals of blockchain systems.

► **RQ 4**: *Has there been collusion among miners to prioritize transaction inclusion?* This research question focuses on examining whether miners engage in collusion to prioritize the inclusion of transactions. For instance, if an issuer sends a transaction to a particular miner, we aim to investigate whether other miners share this transaction to accelerate or prioritize its commitment. If such collaboration exist, it suggests that miners may cooperate not only to accelerate transactions but could also potentially censor specific transactions if they choose to do so. To address this research question, an active experiment was conducted in the context of Bitcoin.

These research questions are key to investigating the impact of the lack of transaction contention and prioritization in both Bitcoin and Ethereum blockchains. They allow us to explore the prevalence and extent of private relay networks or acceleration services currently in use within these blockchains. Furthermore, by examining whether miners collude to prioritize transactions, we gain insight into potential trust issues within blockchains. For example, this collusion can undermine the trust in the blockchain system, as miners could also censor transactions if they choose to. Next, we discuss our methodology, our findings, and the implications.

Relevant publication

The results presented in this chapter have been published in (Messias et al., 2023a).

Table 4.1: *Bitcoin and Ethereum data sets ($\mathcal{D}$ and $\mathcal{E}$) used to evaluate the lack of contention and prioritization transparency.*

Attributes	*Data set $\mathcal{D}$*	*Data set $\mathcal{E}$*
Time span	Jan. 1st, 2018 – Dec. 31st, 2020	Sept. 8th, 2021 – Jun. 30th, 2022
Block height / number	501,951 – 663,904	13,183,000 – 15,049,999
Number of blocks	161,954	1,867,000
Count of transactions issued	313,575,387	347,629,393
Percentage of CPFP-transactions	21.02%	—
Count of empty-blocks	992	43,069

4.1 Methodology

In this section, we outline our methodology for evaluating the lack of transparency in transaction contention and prioritization. First, we provide an overview of our Bitcoin and Ethereum data sets utilized in our analysis. Subsequently, we offer detailed information on the active experiments employed to assess preferential treatment of transactions and the aggregated power of colluding miners.

4.1.1 Data set collection

Data set $\mathcal{D}$. To identify accelerated transactions, we gathered all Bitcoin blocks mined from January 1st 2018 to December 31st 2020. In total, per Table 4.1, there are 161,954 blocks from block height 501,951 to 663,904, and 313,575,387 transactions. In Bitcoin, mining pools may indicate their ownership of the block by including a *signature* or *marker* in the *Coinbase* transaction (i.e., the first transaction of every block). We used such markers for identifying the mining pool (owner) of each block following techniques from prior work (Judmayer et al., 2017; Messias et al., 2021; Romiti et al., 2019). We failed to identify, however, the owners of 4911 blocks (approximately 3% of the blocks) and grouped these blocks under the label "Unknown." Figure 4.1a shows the distribution of the count of blocks mined and transactions confirmed by the top-20 mining pools. We further removed 65,902,514 (21.02%) *child-pays-for-parent (CPFP)* transactions from our acceleration analyses. If we rank the MPOs by the number of blocks (B) mined (which essentially approximate the hashing capacity h of the MPOs), the top five MPOs are BTC.com (B: 27,534; h: 17.00%), F2Pool (B: 20,665; h: 12.76%), AntPool (B: 20,188; h: 12.47%), Poolin (B: 15,096; h: 9.32%), and ViaBTC (B: 13,419; h: 8.29%). But, unsurprisingly, the MPOs' hash rates significantly varied over the years (refer Figure B.1 in

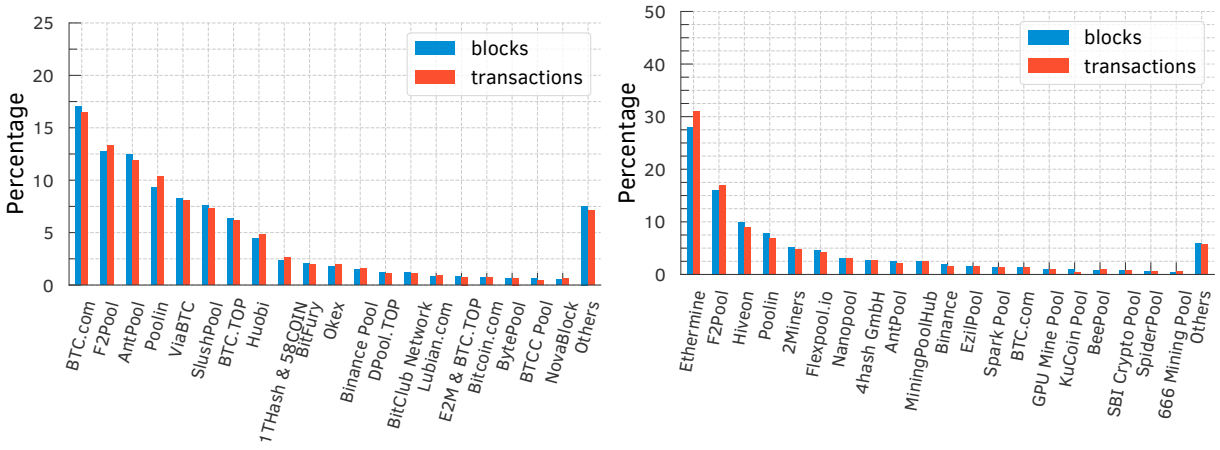

(a) Bitcoin distribution **(b)** Ethereum distribution

Figure 4.1: Blocks mined and transactions confirmed in (a) Bitcoin and (b) Ethereum by the top-20 mining pools; "Others" consolidates the remaining mining pools.

§B.3). Furthermore, we rely on our SPPE metric to infer whether a transaction was likely accelerated.

Data set $\mathcal{E}$. We gathered all Ethereum blocks mined over a 9-month time period—from September 8^{th}, 2021 to June 30^{th}, 2022—to investigate the behavior of Ethereum mining pools (refer to Table 4.1). This data set contains 347,629,393 issued transactions and 1,867,000 blocks (from block number 13,183,000 to 15,049,999). We used miners' wallet addresses to infer the block owners, but we failed to identify the owners of 46,895 blocks (or 2.51% of the total); we grouped the latter into one category, "Unknown." Figure 4.1b shows the distribution of blocks and transactions mined in Ethereum by the top-20 mining pools. If we rank the mining pools by the number of blocks (B) mined (which approximate the hashing capacity h of the mining pools), the top five mining pools are Ethermine (B: 523,633; h: 28.05%), F2Pool (B: 299,418; h: 16.04%), Hiveon (B: 185,495; h: 9.94%), Poolin (B: 147,983; h: 7.93%), and 2Miners (B: 97,308; h: 5.21%) that together account for 67.17% of all blocks mined. We also report the weekly Ethereum's hash rate in Figure B.2 in §B.3. Hash rates of mining pools in Ethereum across the study period did not vary as much as in Bitcoin (refer to Figure B.1 in §B.3).

Additionally, we also retrieved 6,937,292 transactions (2% of all issued transactions in Ethereum) contained in 3,284,886 bundles from Flashbots; these are transactions sent privately to miners.

4.1.2 Bundling public transactions

To identify bundles likely sent through the public P2P network, we use a simple heuristic. We concentrate on transaction bundles of sizes 2 and 3, seeking transactions that probably contributed to a publicly transmitted transaction being bundled and signs of sandwich attacks (Qin et al., 2022).

The underlying idea is that miners have no incentive to include transactions offering zero fees, as there is no reward for mining such transactions—unless they receive additional payment via Flashbots coinbase transfer. Consequently, transactions with a non-zero max-priority fee likely underwent public sending.

We discuss the details and our results in §4.3.1.

4.1.3 Aggregated power of colluding miners

Collusion among mining pools directly challenges the fundamental principle of truly decentralized blockchains. For instance, when powerful mining pools collude to give preference to certain transactions, and they collectively have a hash rate surpassing 50% of the network, there is no barrier preventing them from also censoring the validation of other transactions. As a result, they could potentially gain substantial control over which transactions are included in the blockchain, creating a significant risk of centralization.

To assess the real-world occurrence of mining pool collusion, we conducted an active experiment within the Bitcoin network. In this experiment, we paid a single mining pool to accelerate a set of public-low-fee-rate transactions. This was done during periods of high congestion in the Mempool. Without this acceleration, these transactions would have faced long delays before being included in the blockchain, due to their public low fees. However, although we paid just one mining pool to accelerate these transactions, our findings unveiled a concerning revelation: *these accelerated transactions were also included and prioritized within a block by other powerful mining pools*. Collectively, these mining pools possess a hash rate exceeding 50%. This situation raises significant concerns that reverberate across the entire blockchain ecosystem.

In §4.3.2, we delve into a comprehensive discussion of our findings and implications of the mining pool collusion.

4.2 On contention transparency

In this section, we show that contention transparency does not hold in practice as transaction relay networks become popular in Ethereum. This allows miners to include

transactions privately and therefore not every miner or even transaction issuers have the full view of all available transactions pending for inclusion.

4.2.1 The rise of private relay networks

With the lucrative market of Decentralized Finance (DeFi) in Ethereum, today, bots engage in predatory front-running behaviors such as sandwich attacks and transaction-replay attacks (Daian et al., 2020; Kiffer et al., 2017; Qin et al., 2022, 2021; Torres et al., 2021; Weintraub et al., 2022; Zhou et al., 2021). Relay networks help users to counter such attacks: They provide users with a private channel for communicating with miners, who have to prove their identity to participate in the relay. Relay networks help users completely bypass the P2P network: Users send their transactions to the relay network, which in turn relays them to its participant miners. The relay network and its participants claim (a) not to front-run these transactions; and (b) to keep them private until they are included in a block (Flashbots, 2022a). These transactions, hence, by construction, experience no front-running issues. Relay networks are centralized; if miners misbehave, they may lose their network membership and forfeit their future profits. Multiple relay networks (e.g., bloXroute (BloXroute Labs, 2022), Taichi Network (SparkPool, 2021), and others (Eden Network, 2022; Ethermine, 2022)) exist today, but we focus on Flashbots (Flashbots, 2022c), the largest relay network for Ethereum.

Flashbots private relay network

As discussed previously, at the time of our analysis, Flashbots is the most popular private relay network in Ethereum. Flashbots's users *bundle* one or more transactions in some specific order (Flashbots, 2022b). Miners are expected to mine the entire bundle (retaining the ordering of transactions within the bundle) and place it at the top of their blocks. The miners receive a fee (paid via a direct transfer to their wallets) for including the bundle in addition to the (traditional) fees associated with the transactions in that bundle. If there are two competing bundles—capturing the same financial opportunity, e.g., liquidations—miners will choose the one with the highest reward (i.e., maximizing financial incentives). The other bundle is *discarded* (since the financial opportunity no longer exists after having been captured by the included bundle), albeit its transactions do *not* expend *any* gas. Therefore, except for a network base fee introduced in EIP-1559,[11] arbitrageurs and liquidators can participate without having any balance in their wallet: If they successfully capture a financial opportunity, they pay the miner from the profit

[11]The EIP-1559 went live in the Ethereum's London hard fork upgrade on August 5th, 2021, at block number 12,965,000.

secured and pocket the rest (Flashbots, 2022c). Flashbots is a *free* to use relay network, and they allow anyone to query whether a transaction used their relay network and the private fees paid to the miner (after it has been committed in a block). We use this publicly available data for analyzing the transactions issued (privately) on Flashbots. Flashbots, however, does not list the discarded bundles (or its transactions): we have access, hence, only to committed transactions.

4.2.2 Characterizing private relay networks

Flashbots labels its bundles (and constituent transactions) into one of three categories: (i) *flashbots*, which represent those sent through their private relay; (ii) *rogue*, referring to those delivered to a (Flashbots) miner, but via a different relay network; and (iii) *miner payout*, indicating a bundle containing payouts to users of a mining pool (Weintraub et al., 2022). We found 58.82%, 27.93%, and 13.25% of transactions belonging to the flashbots, miner payout, and rogue categories, respectively. We also noticed that $70{,}260$ (1.01%) of all Flashbots transactions failed to execute after inclusion in a block. A small fraction of transactions is, hence, not successfully executed despite using private relays.

Flashbots also claims to have $\approx 85\%$ of the total Ethereum hash rate (Flashbots, 2022c). Per our analyses, however, the majority of the mining pools (47 out of 48—barring EthPool) use Flashbots, accounting for 99.99% of the total Ethereum hash rate, A recent work also corroborates our findings (Weintraub et al., 2022).

Some of the most powerful mining pools like Spark Pool[12] (which cooperates with Taichi Network (SparkPool, 2021)), Ethermine (Ethermine, 2022), and F2Pool (part of Eden Network (Eden Network, 2022)) offer their own relay networks. As these networks allow transaction issuers to send transactions exclusively to a specific miner, we hypothesize that miners would prefer (or prioritize) these transactions to those sent via the public P2P network. Crucially, payments from these private transactions are guaranteed, while those from publicly issued transactions are not—they are available to any miner willing to commit them. *Miners, hence, would likely offer preferential treatment for private transactions.*

4.2.3 On preferential treatment of private transactions

We substantiate our hypothesis of preferential treatment for private transactions via an active experiment conducted on September 8^{th}, 2021. We issued 8 transactions, where 4 were sent privately via the Taichi Network, powered by Spark Pool, and 4 through the

[12]Spark Pool suspended their mining services on September 30^{th}, 2021, due to regulatory requirements introduced by Chinese authorities (Helen Partz, 2021).

public Ethereum network (refer Table B.1 in §B.1). We spent 100 Euros for running this experiment.

While running the experiment, we checked if the popular Ethereum blockchain explorers (i.e., Etherscan (Etherscan, 2023b), Blockchain.com (Blockchain.com, 2021), and Blockchair (Blockchair, 2023)) observed any of our private transactions; if they did, it would imply that the Taichi Network leaked the transactions to the public. While the public transactions appeared in these blockchain explorers, right after we sent them through the public P2P network, the private transactions were not observed by any of them until the transactions were included in a block. More importantly, our private transactions were *not* flagged by Etherscan (which relies on Flashbots API (Flashbots, 2022a) and more recently on EigenPhi (EigenPhi, 2022)) as private, *even after inclusion in a block*. Measuring the prevalence of private transactions is, hence, challenging; it is likely that our estimates of the volume of private transactions based on such tools represent, hence, a lower bound.

Our results show that Babel Pool included 2 out of our 4 private transactions. Spark Pool technically supports this mining pool, implying that they "collaborate" in committing private transactions sent over the Taichi network (Babel Finance, 2021). Our transactions were included, however, in the appropriate position in the block based on their fees. We delve into the prioritization of transactions in the next section.

We also characterize the prevalence of private transactions in Ethereum and indicate that mining pools can each have a distinct set of private transactions in their Mempool. Users, as a result, can no longer rely on the public Mempool alone to estimate their transaction fee. Given the absence of other data, they are highly likely to end up with a false estimate of the "appropriate" transaction fees for their transactions.

4.3 On prioritization transparency

In this section, we delve into our analysis of prioritization transparency within the Ethereum and Bitcoin blockchains. We show that the current assumptions about transparency in blockchains do not hold in practice. Then, we show that transaction relay networks are becoming more popular in Ethereum, with miners creating their own transaction relay networks for private transactions. We also show that miners have different measures for utility of mining a transaction than just the offered fee rate or gas price of a transaction. For instance, transactions issuers that pay miners via a direct transfer to their wallet address or through off-chain fee receive a higher prioritization than their corresponding transaction fee rates would suggest.

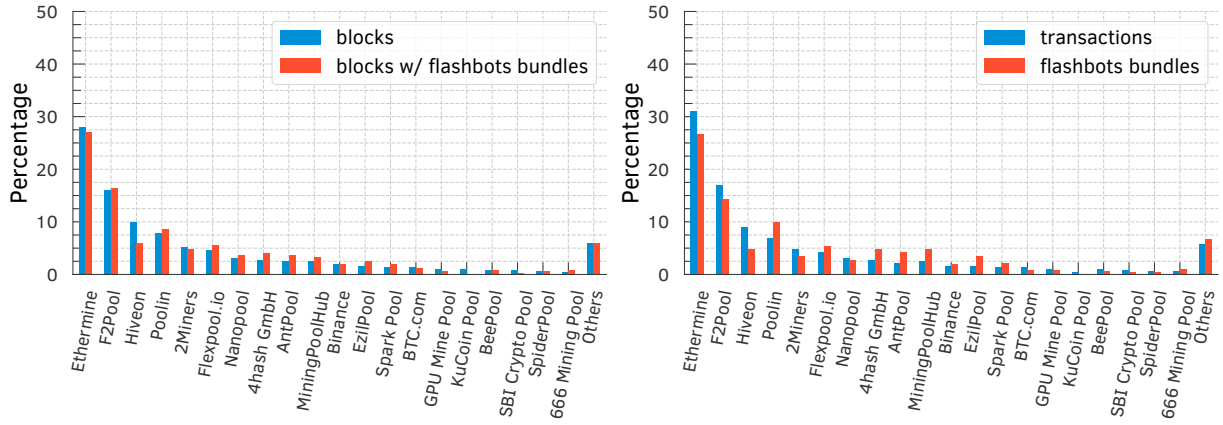

(a) Distribution of Flashbots blocks

(b) Distribution of Flashbots bundles

Figure 4.2: Distribution of (a) blocks with at least one Flashbots bundle; and (b) bundle of transactions per block, per mining pool. Ethermine included 27.05% of all blocks with a Flashbot bundle and 26.63% of all Flashbots bundles, while mining around 28.05% and 31.11% of all blocks and transactions, respectively.

4.3.1 Prevalence of transaction bundling

In this section, we use the Flashbots data set outlined in §4.1, which has 6,937,292 transactions (2% of all issued transactions in Ethereum) contained in 3,284,886 bundles from Flashbots. These bundles constitute transactions privately sent to miners. For instance, among all blocks in the data set $\mathcal{E}$, 972,911 (52.11%) of blocks have at least one such Flashbots transaction: *Private transactions are becoming quite common across most of the powerful mining pools in Ethereum.*

Flashbots bundles are quite prevalent in Ethereum, representing 99.99% of the total Ethereum hash rate (refer §4.2.2). Our analysis shows that each Flashbots bundle contains at least 1 transaction and at most 631 transactions; on average they contain 2.11 transactions, with a median of 1 and a std. of 6.47. We noticed that Ethermine alone included more than a quarter (26.63%) of all 3,284,886 bundles (refer Figure 4.2). Also, blocks contain at most 40 bundles, with an average of 3.38, a median of 3, and a std. of 2.64 bundles.

Miner Incentives in Incorporating Flashbots Bundles

Flashbots allows users to bundle together a set of transactions, thereby specifying the order in which they are executed. The bundles can also include public transactions, propagated over the public P2P network. A public transaction that buys a coin on a Decentralized Exchange (DEX) can, for example, lead to an arbitrage opportunity (Qin

et al., 2021). A user can include this transaction in a bundle along with one of their own to capture this arbitrage opportunity. The last transaction in the bundle usually pays the miner (based on the profit made) in *Ether*[13] via a direct transfer (i.e., *coinbase transfer*) to their wallet addresses (Flashbots, 2022c). This essentially means that miners are being offered different prices for mining the same transaction. In other words, miners have a financial incentive for including transactions that are in a bundle at the top of a block, even though the public fee offered through gas price in the transaction data is very low (refer Figure 4.3). Hence, each transaction in the bundle has a normal gas price and a *bundle gas price*, which is calculated using the total gas used by all transactions in the bundle and the total miner reward for mining the bundle.

Bundling public transactions

To identify bundles with transactions that were probably sent through the public P2P network, we rely on a simple heuristic. Specifically, we focus on transaction bundles of size 2 and 3, and search for transactions that have likely resulted in a publicly sent transaction being bundled. Then, we find bundles issued from different issuers that include a zero and non-zero *max-priority fee*[14] transactions. The intuition is that miners have no incentive to include transactions that offer a zero max-priority fee, as they receive no rewards for mining these transactions. Unless they receive extra payment (through Flashbots coinbase transfer). Hence, transactions that have a non-zero max-priority fee were likely sent publicly.

For transaction bundles of size 2, we look for transactions whose issuers are not the same. Furthermore, we look for cases where the first transaction offers a non-zero max-priority fee, with no coinbase transfer to the miner, and the second transaction offers a 0 max-priority fee and a non-zero coinbase transfer to the miner.

For transaction bundles of size 3, we look for signs of sandwich attacks (Qin et al., 2022). We look for bundles where the first and last transactions have the same issuer, but the second transaction has a different issuer. Additionally, we check that the first and third transactions offer a 0 max-priority fee, meaning that the miner receives no reward from the gas price for mining these transactions. Then, we ensure that the second transaction offers miners a non-zero max-priority fee, while the third offers miners a fee through direct coinbase transfer. This scenario might be a classic sandwich attack, where

[13]Ether (ETH) is the cryptocurrency used in the Ethereum blockchain to incentivize and reward its participants. Its smallest denomination is called Wei, representing a fraction of 10^{-18} ETH. When referencing gas prices or fees, the notation GWei is utilized, equating to a value of 10^{-9} ETH.

[14]The *max-priority fee* was introduced in EIP-1559 as the unique financial incentive miners get for including publicly announced transactions. The other fees are burned.

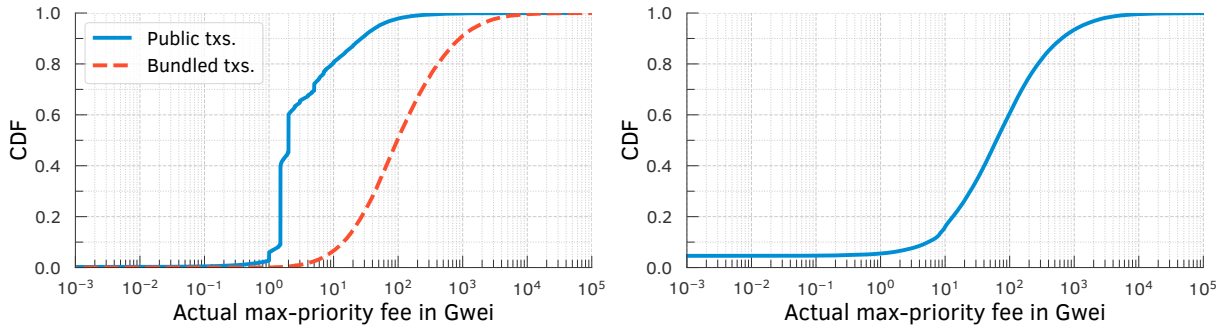

(a) Public vs bundled's actual fee

(b) Difference in actual max-priority fee

Figure 4.3: Difference between the actual max-priority fee of public transactions and Flashbots bundles; bundles typically offer a larger *effective* fee to the miners.

public transactions are bundled between two private transactions, sent by the same issuer, and the miner gets paid via a coinbase transfer from the third transaction (Qin et al., 2022).

We found 853,394 transactions in 426,697 bundles of length 2, and 1,231,695 transactions in 410,565 bundles of length 3. From those, we found that 110,401 (25.87%) and 37,447 (9.12%) bundles, of lengths 2 and 3, respectively, fit our heuristic. We then calculate the *actual max-priority fee* for these bundles, as the total gas used by all transactions in the bundle divided by the total miner reward (from gas usage and coinbase transfer). Figure 4.3 shows the price difference miners get for including publicly and bundled transactions. Note that around 40% of transactions differ in the actual max-priority fee by 100 gwei-per-units-of-gas. Flashbots bundles offers much higher gas prices in comparison to the public announced max-priority fee alone.

Towards liquidations through bundling

Lending protocols rely on *over-collateralization* of assets: In order to borrow assets from these protocols, a user has to deposit a collateral of at least 150% of the borrowed amount. To borrow 1 USDC on AAVE, for example, a user would have to collateralize at least 1.5 USDC worth of another asset (e.g., in ETH or BTC). If the ratio of the collateral asset versus the borrowed asset falls below 1.5, the user's position can be liquidated by any other participant until the ratio stabilizes to 1.5 again. The liquidator then pays back a portion of the user's debt to receive the collateral asset at a discount. In order to assess an asset's on-chain value, lending protocols rely on oracle services, e.g., Chainlink Data Feeds (Breidenbach et al., 2021; Chainlink, 2022). In the case of the two largest lending platforms, AAVE V2 (AAVE, 2022) and Compound (Compound, 2022), for instance, Chainlink provides the price of each asset in ETH and USD, respectively.

We found 16,418 liquidations in AAVE and 6387 liquidations in Compound. Out of these, there were 4863 AAVE liquidations and 2036 Compound liquidations that were sent privately through Flashbots. In AAVE, the three largest collateral assets that were liquidated were WETH (57.58%), LINK (11.84%), and WBTC (8.99%). The debt assets paid for, i.e., the assets borrowed by the users, were USDC (33.77%), USDT (22.27%), DAI (19.39%), and GUSD (5.12%), all of which are stablecoins and account for over 80% of the assets repaid by liquidators. In Compound, the three largest collateral assets that were liquidated were WETH (69.7%), WBTC (10.31%), and UNI (5.5%). The debt assets were USDC (38.9%), DAI (30.45%), USDT (23.38%), and TUSD (2.7%), all of which are stablecoins and account for over 90% of the assets repaid by liquidators.

Liquidation with bundled oracle updates

To check the adverse effect of bundling oracle updates, we looked at bundles with Chainlink (Chainlink, 2022) oracle updates as they are a key part of liquidations. We identified 1165 AAVE liquidations distributed within 1154 bundles (2662 transactions including 1301 oracle updates) that contained at least one oracle update. In Compound, we found 648 liquidations distributed within 641 bundles (1457 transactions including 751 oracle updates) that contained oracle updates. In AAVE, out of 1154 bundles, there were 994 (86.14%) bundles that contained an oracle update followed by a liquidation, and 52 (4.51%) with two oracle updates followed by liquidations. In Compound, out of 641 bundles, there were 548 (85.49%) bundles that contained an oracle update followed by a liquidation, and 39 (6.08%) with two oracle updates followed by liquidations. For details on the specific liquidations for both AAVE and Compound, please refer §B.2 in the appendix. Out of the total 1813 liquidations in AAVE and Compound we found that only 24 were possible in the previous block. Almost 98.68% of such liquidations were, hence, only possible because of the Chainlink updates in that block.

In order to calculate the profit made by the liquidators, we get the amount of debt that was repaid and the amount of the underlying collateral that was received by the liquidator. We calculate the price of each token at the time of liquidation by looking at the on-chain oracle price from Chainlink at the same block number, where the liquidation took place. For AAVE and Compound, we specifically use the Chainlink on-chain price used by AAVE and Compound in their respective protocols. AAVE uses the price in ETH as a reference for its tokens, whereas Compound's price oracles are denominated in USD. For AAVE, in order to calculate the profit made by each liquidation, we calculate the profit in ETH, and then multiply the profit by the current Chainlink on-chain price of ETH in USD. Per Figure 4.4, liquidations that are bundled with a Chainlink update also

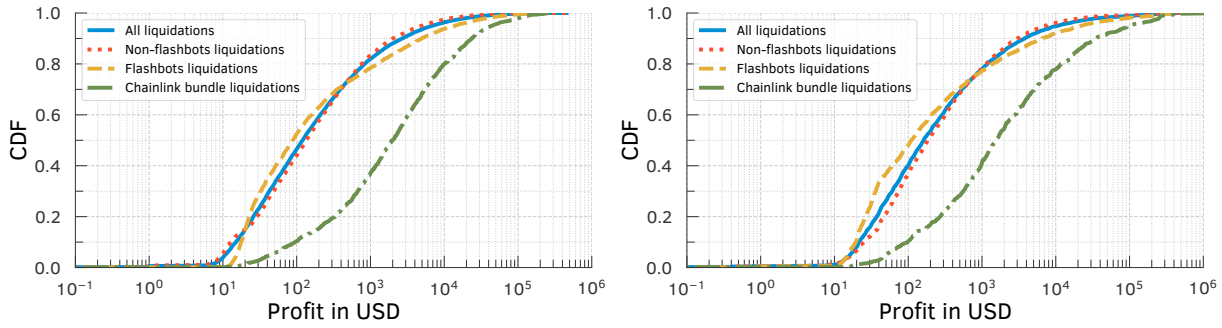

(a) Liquidations profit in AAVE

(b) Liquidations profit in Compound

Figure 4.4: Profits of liquidators in (a) AAVE and in (b) Compound. Liquidations bundled with Chainlink updates generally provide higher profits.

Table 4.2: *There are* 2,231,051 *(*67.92%*) unique Flashbots bundles, and* 3,076,760 *(*44.35%*) transactions, that called the following decentralized exchange contracts in Ethereum: 0x Protocol, Balancer, Bancor, Curve, SushiSwap, Uniswap V1, or V3. Note that a single transaction or bundle might call one or more contracts.*

	Balancer	*Bancor*	*Curve v1 & v2*	*Uniswap v2 & Sushiswap*	*Uniswap v3*	*0x Protocol v1, v2 & v3*	*Total*
# of bundles	85,422	96,122	53,296	1,710,985	1,337,715	28,753	2,231,051
	3.83%	4.31%	2.39%	76.69%	59.96%	1.29%	67.92%
# of transactions	87,865	99,040	58,188	2,533,084	1,692,485	29,100	3,076,760
	2.86%	3.22%	1.89%	82.33%	55.01%	0.95%	44.35%

have larger profits for liquidators, which implies that the lucrative liquidations are more likely to be bundled together with a Chainlink update.

Characterizing transaction bundling

To investigate which DEXes protocols are called within Flashbots bundles, we focus on the following contract calls: 0x Protocol (0x Protocol, 2022), Balancer (Balancerl, 2022), Bancor (Bancor, 2022), Curve (Curve, 2022), SushiSwap (SushiSwap, 2022), and Uniswap V1 and V3 (Uniswap, 2022). In our set of 3,284,886 Flashbots bundles, we find that 2,231,051 (67.92%) unique Flashbots bundles (and 3,076,760 transactions) called at least one of these contracts. Table 4.2 shows the distribution of the number of transactions and the number of bundles for each of these contracts. We see that Uniswap and SushiSwap are the most bundled DEXes protocols in Flashbots.

4.3.2 Side channel (dark-fee) payments and transaction acceleration

In this section, we focus on the Bitcoin blockchain, with a particular emphasis on the data set $\mathcal{D}$. Our goal is to build upon our earlier discussion in §3.4 regarding dark fees transactions.

Prevalence of transaction acceleration

As previously discussed in §3.4, dark-fee transactions (or accelerated transactions) are transactions that offer additional fees to specific mining pools via an opaque and non-public side-channel payment. In Bitcoin, the top 5 mining pools named BTC.com (BTC.com, 2022), AntPool (AntPool, 2022), ViaBTC (ViaBTC, 2022), F2Pool (F2Pool, 2022), and Poolin (Poolin, 2022), deploy transaction acceleration services, which enables users to "accelerate" the confirmation of their transactions by offering mining pools dark-fees.

These (dark-)fees are paid in fiat currency through a direct bank transfer or via other crypto coins to the mining pool. They are, therefore, opaque or dark to other participants. Strangely enough, these fees are also non-refundable as the miner receives them regardless of whether they include the transaction in a block or not—a guaranteed payment. The fees paid by the transaction issuer are, furthermore, not made public: only the user and the miner knows the actual fee paid by the transaction inclusion. Since transaction issuers pay the fees off-chain, miners have an incentive for prioritizing these transactions despite the low fee rate offered on-chain. It also implies that the transaction issuer offers a miner a different fee compared to that offered to other miners for including their transaction in a block. Miners do not disclose such private fees paid by issuers. This behavior is different from that of Flashbots in Ethereum: The latter discloses the final dark-fee after the transaction is committed (see §4.3.1).

Characterizing transaction acceleration

In order to detect accelerated transactions, we proposed two metrics called *signed position prediction error (SPPE)* and *position prediction error (PPE)* that are described in §3.2.2.

To estimate the prevalence of accelerated transactions in blocks mined by different mining pools, we compute the fraction of blocks mined by the top-15 mining pools, based on their hash rates in our 3-year data set $\mathcal{D}$ (refer to §B.3 and Figure 4.1a), that contained transactions with SPPE $\geq 99\%$. Per Figure 4.5, we find that many large mining pools such as BTC.com, F2Pool, and ViaBTC are likely including accelerated transactions in a sizeable fraction of their mined blocks, with ViaBTC including it in over 40% of their blocks.

Table 4.3: *Accelerated transactions have fewer delays and are included at the top of the block, i.e., at higher positions compared to non-accelerated transactions.*

metrics	*delay in # of blocks*		*perc. position in a block*	
	acc.	*non-acc.*	*acc.*	*non-acc.*
minimum	1	9	0.07	17.47
25-perc	1	148	0.08	75.88
median	2	191	0.09	87.92
75-perc	2	247	0.20	95.00
maximum	3	326	4.39	99.95
average	1.8	198.5	0.79	84.46

If we consider all mining pools' transactions with an SPPE $\geq 50\%$ (1,869,043 transactions, in total), from 2018 to 2020, users transferred in total 11,631,217 BTC (or ≈ 223.55 billion USD[15]). The accelerated transactions accounted for 240,226 BTC (or ≈ 4.62 billion USD), corresponding to approximately 2.07%.

Aggregated power of colluding miners

In order to check the impact of transactions acceleration services on commit time of transaction, we ran active real-world experiments. Specifically, we paid ViaBTC (ViaBTC, 2022) to accelerate selected transactions (see Table B.2 in §B.4) during periods of high congestion between November 26th and December 1st, 2020. From 10 Mempool snapshots during this period, we selected transactions that offered a very low fee rate (i.e., 1–2 sat-per-byte) for acceleration. To keep our acceleration costs low, we selected transactions with the smallest size (which was 110 bytes) within this set. For each of the 10 snapshots, we had multiple transactions with such low fee rates and small size, for a total of 212 transactions across all the snapshots. We randomly selected one transaction from each snapshot (i.e., 10 transactions) and paid ViaBTC 205 EUR to accelerate them.

We then compare the priority with which the accelerated transactions and the 202 ($= 212 - 10$) non-accelerated transactions with similar fee rates and sizes were included in the Bitcoin blockchain. The impact of acceleration was strikingly apparent as shown in Table 4.3. All 10 accelerated transactions were included within 1–3 blocks after their acceleration, with an average delay of 1.8 blocks. In contrast, the minimum delay for the 202 non-accelerated transactions of comparable fee rates and sizes was 9 blocks, with an average delay of 198.5 blocks. Interestingly, 38 of the non-accelerated transactions were yet to be included in the blockchain by December 4th, 2020. Similarly, the accelerated

[15]Based on the Bitcoin exchange rate on October 19th 2022, 1 BTC = 19,219.90 USD

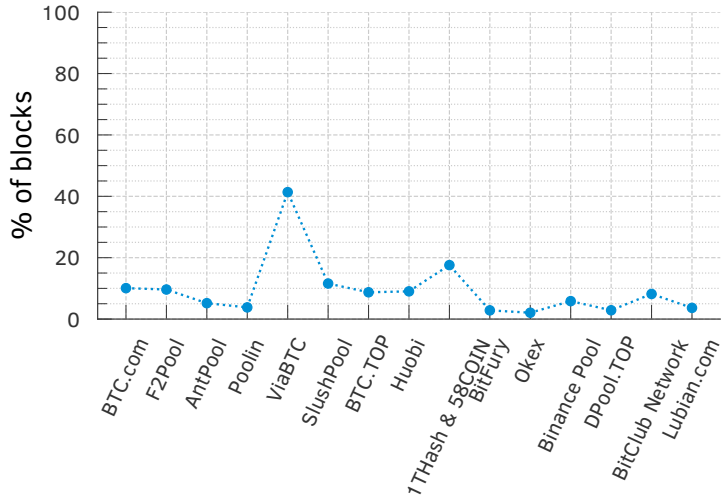

Figure 4.5: Blocks with accelerated transactions (with SPPE $\geq 99\%$) are quite common among the top 15 mining pools. In Bitcoin, the mining pools with a high percentage of such blocks are ViaBTC (41.36%), 1THash & 58COIN (17.58%), SlushPool (11.58%), BTC.com (10.03%), and F2Pool (9.63%).

transactions were included in top 0.07–4.39 percentile positions, with an average 0.79 percentile position, while the non-accelerated transactions were included in the beyond top 17.47–99.95 percentile positions, with an average 84.46 percentile position. From the above observations, it is clear that the transactions we accelerated were included with high priority, meaning Bitcoin mining pools take off-chain fees into account when prioritizing transactions.

Although, we accelerated our transactions using ViaBTC mining pool, our 10 transactions were included by 5 different mining pools, namely F2Pool, AntPool, Binance, Huobi, and ViaBTC. As we accelerated transaction during time of high congestion in Bitcoin, no mining pool would have included a transaction offering 1–2 sat-per-byte, unless they were accelerated. Since we only paid the ViaBTC mining pool, this implies that ViaBTC is colluding with other mining pools to accelerate transactions that offer off-chain fees. Except for Binance, all these colluding pools rank amongst the top-8 mining pools in terms of their hash rates at the time of our experiments. Table 4.4 shows the individual as well as the combined hash rates of these 5 colluding mining pools over the last day, last week, and last month before the conclusion of our experiment on December 1^{st}, 2020. The most striking and the most worrisome fact is that **the combined hash rates of these colluding mining pools exceeds 55% of the total Bitcoin hash rate**. For more details, refer to Figures B.3 and B.4 in §B.4 in the appendix. Additionally, if mining pools are colluding to include accelerated transactions, then they might also potentially collude in malicious ways.

Table 4.4: *If we rank the miners who confirmed the accelerated transactions based on their daily, weekly, and monthly hash rate power, at the time these experiments were conducted, the combined hash power of these mining pools exceeds 55% of the Bitcoin's total hashing power.*

Mining Pool	*Hash-rate*		
	last 24h	*last week*	*last month*
F2Pool	19.9%	18.7%	19.9%
AntPool	12.5%	10.6%	10.2%
Binance	9.6%	10.3%	10.0%
Huobi	8.1%	9.3%	9.8%
ViaBTC	5.1%	7.1%	7.7%
Total	55.2%	56%	57.6%

Furthermore, due to the lack of transparency into their queue, miners can charge higher prices for their acceleration services when colluding. It means that they can overcharge the transaction issuers for including their transactions.

4.4 Concluding remarks

In this section, we present the findings derived from our analysis of private relayed transactions, along with the results obtained from our active experiments conducted on Bitcoin and Ethereum blockchains. The main objective of these experiments was to evaluate the lack of transparency in transaction contention and prioritization.

In summary, our findings indicate that private transactions and private relay networks are quite prevalent in both Ethereum and Bitcoin blockchains. Flashbots, in particular, is extensively used in Ethereum, accounting for a significant portion of 99.99% of the total Ethereum hash rate. It also enables arbitrageurs to exploit MEV opportunities by bundling their private transactions with public transactions like oracle updates or taking advantage of sandwich attacks. Similarly, in Bitcoin, miners offer transaction acceleration services, allowing users to privately offer a dark-fee to incentivize miners for a faster commit time. Through active experiments, we show that miners highly prioritize these transactions, on average including them in 1.8 blocks, with a range of 1 to 3 blocks. Worrisome, we uncover evidence of collusion among miners with a combined hash rate exceeding 50% to ensure the inclusion of these dark-fee transactions.

In the following chapter, we delve into the voting power distribution for amending smart contracts.

CHAPTER 5

Decentralized Governance

In this chapter, we present our research questions, methodology, and discuss the implications of our findings regarding the level of decentralization in governance protocols. To investigate this, we focus on the Compound governance protocol as a case study. Our analyses reveal that the distribution of voting power in Compound is highly concentrated among a small number of participants, which can significantly hinder the achievement of a fair and decentralized governance system.

The concentration of voting power poses a challenge to achieving truly decentralization in governance protocols. For example, when a small group of participants holds a majority of the tokens, they can make decisions that benefit themselves at the expense of others. Therefore, ensuring a fair distribution of tokens becomes crucial to foster decentralization in these protocols. In this chapter, we aim to analyze transaction data associated with Compound in order to assess the level of decentralization in Compound's voting power. To guide our analysis, we propose the following research questions.

► **RQ 1**: *How frequently are amendments proposed and voted on in the Compound protocol?* This research question aims to investigate the activity level of the Compound protocol and its community engagement. For instance, by examining the frequency with which proposals are amended or voted, we can assess the level of participation and the extent to which the community is actively contributing to improving the protocol.

► **RQ 2**: *What is the distribution of Compound tokens among its participants? How small or large is the set of voters who determine the outcomes for the amendments?* This research question aims to investigate the distribution of Compound tokens among its participants. Hence, we can assess to which extent the Compound tokens is truly decentralized. Understanding this distribution is crucial for proposing fairness to the protocol token's distribution.

► **RQ 3**: *What is the cost associated with casting a vote in the Compound protocol?* Voting in on-chain governance protocols, where the entire voting process happens on the blockchain, requires the payment of transaction fees that vary depending on the network

congestion. These voting costs can disproportionately affect small token holders, potentially limiting their participation in the decision-making process. This research question aims to investigate the impact of voting costs on voter participation in the Compound protocol. It provides insights into the fairness of the decision-making process and shed light on potential barriers that may discourage certain participants from exercising their voting rights.

► **RQ 4**: *What are the voting patterns of delegates, and do voters form coalitions?* This research question aims to analyze the voting patterns of delegates in the Compound protocol and investigate whether voters form coalitions, where they align their votes as a collective group. The formation of coalitions among voters can lead to the marginalization of certain voters, as they consistently find themselves in a minority group. This undermines the core principle of decentralization and has the potential to compromise the security and effectiveness of the governance protocol. Specifically, instead of expressing their individual opinions on a proposal, voters may choose to mimic the voting behavior of their peers. Therefore, exploring the presence and impact of coalition formation can provide valuable insights into the decision-making dynamics among voters and help mitigate the concentration of tokens (or voting power) within the system.

Addressing these research questions is key for improving the protocol's fairness and achieving more decentralization in the distribution of voting power. In the following section, we discuss our methodology for gathering the necessary data related to Compound protocol from our Ethereum archive node.

Relevant publication

The results presented in this chapter have been submitted, and we are currently awaiting a decision (Messias et al., 2023b).

5.1 Methodology

To analyze the voting power concentration among Compound token holders (or voters), we adopt a data-driven approach. Our methodology involves collecting events triggered by transaction executions when voters cast votes, create proposals, cancel proposals, transfer tokens, or delegate their voting rights to another address. We cover events from the inception of the Compound token and Compound governance protocol. To address the possibility of a single entity owning multiple addresses, we have developed a methodology to infer address ownership and group them accordingly. This approach allows us to identify and consolidate addresses that are likely controlled by the same

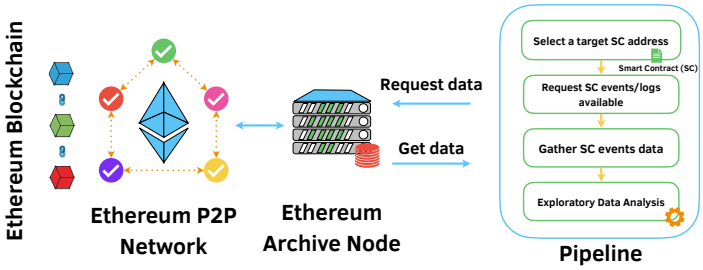

Figure 5.1: Overview of the data collection methodology and analysis.

entity. This process utilizes data from well-known blockchain explorers and publicly disclosed information regarding address ownership. For those addresses that we were not able to infer their ownership we renamed them to their specific wallet addresses.

To gather the data, we deployed an Ethereum *archive* node on a server with 64 cores (with a base clock frequency of $2.25\,\mathrm{GHz}$ that can be boosted to $3.4\,\mathrm{GHz}$), $256\,\mathrm{MB}$ L3 cache, $252\,\mathrm{GB}$ of RAM, and $21\,\mathrm{TB}$ of NVMe-based storage. The archive node took about 4 weeks—a relatively long time, though not unexpected—to fully synchronize with the Ethereum blockchain. We used Web3.py (web3.py team, 2022), a Python library for interacting with Ethereum nodes, to query and retrieve the information that we need from the archive node. Figure 5.1 summarizes our methodology for the Ethereum data gathering.

5.1.1 Smart contract events.

Smart contracts in Ethereum can generate and dispatch *events* for signaling various types of activities (e.g., ERC-20 token transfers or state changes) within the contract. We can subscribe to these events, or analyze them later since Ethereum persists the events in the blockchain via the "logs" field of the transaction receipt attribute. In this thesis, we leveraged these logs to filter transactions that triggered specific events, e.g., sending, receiving, or swapping tokens. We also filtered and analyzed transactions that triggered events related to governance protocols to track the evolution of each proposal, including when it was created, when users started voting, and when it was executed or canceled.

5.1.2 Data set collection

We gathered various details on Compound tokens and Compound governance contracts between March 3, 2020 (block #$9{,}600{,}000$) and November 7, 2022 (block #$15{,}917{,}000$) from our Ethereum archive node. This 32-month study period includes Compound's entire lifetime (from its inception). We illustrate our methodology and data-analysis pipeline

Table 5.1: *Summary of events related to the Compound (COMP) token that we gathered from the Ethereum blockchain.*

Event name	*# of events*	*Description*
Approval	213,220	Standard ERC-20 approval event.
DelegateChanged	12,095	Emitted when an account changes its delegate. This means that the delegatee will receive voting power from the sender. Users can only delegate to one address at a time, and the number of votes added to the delegatee's vote count is equal to the user's balance. The delegation of votes will take effect from the current block until the sender either delegates to a different address or transfers their tokens.
DelegateVotesChanged	75,820	Emitted when a delegate account's vote balance changes.
Transfer	1,886,618	Emitted when users/holders transfer their tokens to another address.

Table 5.2: *Summary of events related to the Compound Governor contracts recorded on the Ethereum blockchain.*

Event name	*# of events*	*Description*
ProposalCanceled	17	Emitted when a proposal is canceled.
ProposalCreated	133	Emitted when a new proposal is created.
ProposalExecuted	101	Emitted when a proposal is executed in the TimeLock.
ProposalQueued	105	Emitted when a proposal is added to the queue in the TimeLock.
VoteCast	9500	Emitted when a vote is cast on a proposal: 0 for against, 1 for in-favor, and 2 for abstain.

in Figure 5.1. We obtained 213,220 *Approval* events, 12,095 *DelegateChanged* events, 75,820 *DelegateVotesChanged* events, and 1,886,618 *Transfer* events for Compound tokens (refer to Table 5.1). We also collected various events (refer to Table 5.2) related to the Compound Governance contract for analyzing various aspects of the proposal creation and voting processes.

5.1.3 Inferring wallet address ownership.

Since an entity can control multiple wallet addresses in the blockchain, identifying the ownership of these wallets helps in grouping together the accounts that are owned by the same entity. However, this task of wallet-address ownership determination is challenging due to the inherent anonymity of blockchains (Antonopoulos, 2014; Antonopoulos and Wood, 2018). This task is further complicated because owners are only identifiable if they choose to voluntarily make their identities public. To address this challenge, we combine wallet ownership information from two widely used data sources: Etherscan (Etherscan, 2023b) and Sybil-List (Sybil, 2023b). The former is a blockchain explorer that helps in identifying the top holders of various cryptocurrencies, and the latter, a Uniswap governance tool for discovering delegates addresses (Sybil, 2023a). It uses cryptographic

proofs for verifying wallet addresses voluntarily disclosed by the wallet owner. From these two data sources, we gathered the owners of 3191 public wallet addresses. We used these addresses to infer the owners of 17 (51.52%) of the 33 unique addresses associated with proposal creation, 114 (3.42%) out of 3335 proposal voters, and 265 (0.13%) out of 210,598 token holders. By analyzing the top 10 most influential voters for each proposal, determined by the number of delegated tokens they possessed when casting their vote, we were able to infer the ownership of 67 (50.37%) of these 133 unique addresses. Finally, as an entity can control more than one address, we grouped the addresses we identified belonging to the same entity together to conduct our analysis.

5.2 Attacks on governance

A potential issue in the governance of blockchain networks is the concentration of governance tokens in the hands of a few participants, which can pose a threat to the protocol (Mike Dalton, 2022). This issue manifested in Balancer (Balancer.fi, 2023), a decentralized exchange (DEX) running on top of Ethereum, where a user with large amount of governance tokens voted for decisions that were beneficial for the user but detrimental for the protocol (Haig, 2022). When a minority holds a large portion of the tokens, decision-making power can become centralized, which conflicts with the goal of decentralization of governance protocols.

Yet another issue concerns many centralized exchanges that *hold* their users' tokens; they could potentially use these tokens for voting *without* their users' knowledge, compromising the integrity of the voting process (Francisco Rodrigues, 2022; Sam Kessler, 2022). Alameda Research, a former cryptocurrency trading firm, which was affiliated with FTX, for example, voted on 8 proposals and even initiated three proposals (#13, #14, and #16) on Compound (Research, 2020a,b,c). Eventually, one of the proposals was executed. Their goal was to raise the collateral of WBTC from 0% to 40%, which allowed WBTC to be utilized for borrowing other assets (Zack Voell and William Foxley, 2020a). This change may have been beneficial to Alameda Research as they were one of the biggest WBTC minters and held highly leveraged positions (i.e., borrowed money to invest even more) (Jeff Kauflin and Emily Mason, 2022; Zack Voell and William Foxley, 2020b). To alleviate these concerns, centralized exchanges typically promise that they will not use their users' tokens to vote on their behalf (Shaurya Malwa, 2022). While there is no guarantee that they will keep their promise, we can monitor their public wallet addresses to check if the exchange has delegated these governance tokens to another address, or whether they used the tokens for voting while they were stored on that exchange.

Governance protocols intend to eliminate (or at least minimize) centralized decision-making in blockchains. Their effectiveness in achieving that goal can, however, be compromised depending on how the tokens (i.e., voting power) are distributed. This thesis evaluates whether governance protocols uphold their promise of decentralized governance of smart contracts, and, if they do not, investigates exactly how they renege on that promise.

5.3 Compound's governance

Compound (Leshner and Hayes, 2019) is a decentralized lending protocol that allows users to lend and borrow tokens or assets via smart contracts. Lenders earn interest (*yield*) by supplying liquidity to the protocol, while borrowers obtain tokens from the protocol and pay interest on the borrowed tokens.

Compound protocol has two versions of its governance contract: *Alpha* and *Bravo*. *Compound Governor Alpha*, the first version of the governance contract, was deployed on March 4, 2020 (block number 9,601,459) and was active until March 28, 2021 (block number 12,126,254).[16] The improved version, *Compound Governor Bravo*, was deployed on March 9, 2021 (block number 12,006,099) and has been active since April 14, 2021 (block number 12,235,671).[17] Brave introduced several improvements such as smart-contract upgradability (through proxies), a new option for voters to abstain from voting, and the ability for voters to state the reasons behind their voting choices through text comments attached to on-chain votes. The Bravo contract was proposed in proposal #42, and it received 1,438,679.86 votes from 59 voters—all but one vote were in favor of its implementation (Labs, 2021).

5.3.1 Control of governance tokens

The voting power of a user in Compound is proportional to the amount of (delegated) tokens held by that user—one token equals one vote. Below, we examine how these tokens are distributed over time among Compound participants.

[16]The Compound Governor Alpha was deployed at the Ethereum smart contract address 0xc0dA01a04C3f3E0be433606045bB7017A7323E38.

[17]The Compound Governor Bravo was deployed at the Ethereum smart contract address 0xc0Da02939E1441F497fd74F78cE7Decb17B66529.

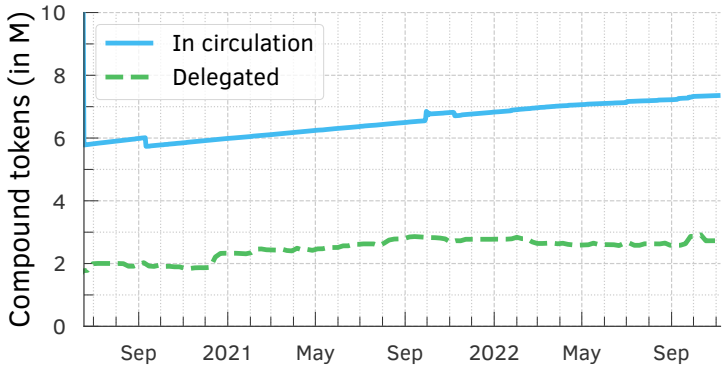

Figure 5.2: Amount of COMP tokens (in millions) in circulation and delegated overtime. Compound tokens have been released to the public since June 15, 2020.

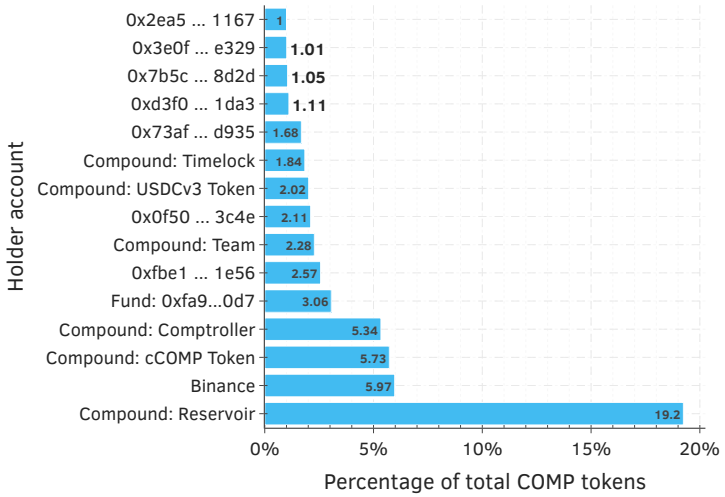

Figure 5.3: Distribution of the top 15 COMP tokens holders. Together, these accounts hold 56.02% (5.6 million) out of 10 million COMP tokens.

Distribution of token holding

Initially, 42.15% of the total Compound supply (10 million COMP tokens) was allocated to liquidity mining,[18] 23.95% to shareholders, 22.46% to the founders and the Compound team, 7.73% to the community, and 3.71% to future team members (CoinGecko, 2023). The public release of COMP tokens started only after proposal #7 was executed on June 15, 2020 (Labs, 2020). This proposal enabled the continued distribution of COMP tokens to the protocol users over time (see Figure 5.2). At the time of our analysis (November 7, 2022), the 10 million COMP tokens were distributed among 210,573 accounts. The largest holder is *Compound Reservoir* with 19.24% (1,924,344.52) of the tokens followed by Binance (5.97% or 397,289.78 tokens) and cComp (5.73% or 572,723.77 tokens) as shown

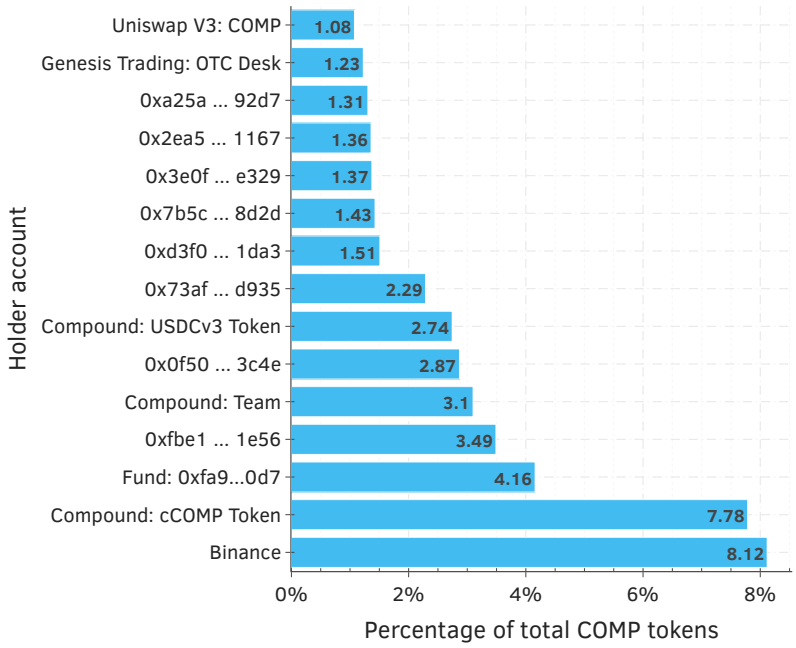

Figure 5.4: Distribution of the top 15 COMP tokens holders (in circulation). These accounts hold 43.83% (3.2 million) out of 7.3 million COMP tokens in circulation.

in Figure 5.3. The Compound Team holds 2.28% (228,061.62) and Compound Timelock 1.84% (184,258.39) of the tokens.

Of the total supply, only 7.3 million COMP tokens are, however, in circulation (Figure 5.2), and we characterize their distribution among a few top token holders in Figure 5.4. In calculating the tokens in circulation, we only included tokens that can be traded or exchanged between users. We excluded *locked* tokens from the Compound Reservoir, Comptroller, and Timelock contracts from our analysis (Compound Labs, Inc., 2022a; kybx86, 2020), which are *not* in circulation. These locked tokens require a governance proposal to be released, although some of them are released daily through the Comptroller as an incentive for users to use the protocol, by lending or borrowing these tokens.

We plot the cumulative distributions of all available COMP tokens along with the locked, delegated, and in-circulation tokens, i.e., the tokens available for users to buy, trade, or sell, in Figure 5.5. The top-15 accounts (in terms of the amount of tokens held) together account for 43.83% of all tokens in circulation (Figure 5.4). Binance (Binance, 2023), a popular centralized cryptocurrency exchange, leads this ranking with 8.12% of the available tokens. It is technically feasible for them to delegate these tokens to

[18]Liquidity mining is a process where users provide liquidity (i.e., tokens) to a protocol in exchange for rewards or interest.

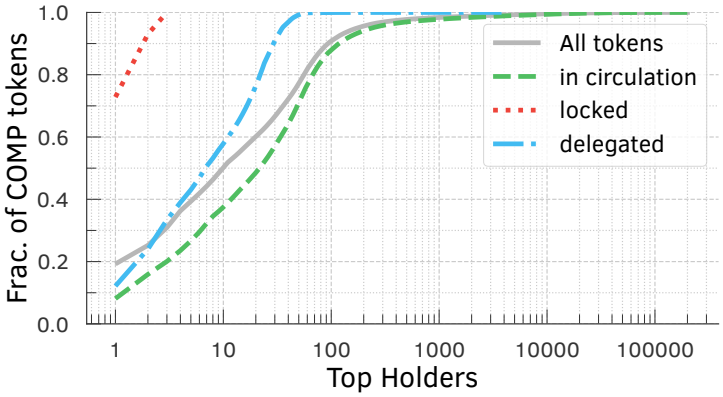

Figure 5.5: Cumulative distribution of the fraction of COMP tokens held per account. The 10 million tokens available are shared among 210,573 accounts (in grey). The dashed green line shows the distribution of the fraction of 7.3 million (73.57%) COMP tokens in circulation held by 210,570 accounts. The 2.6 million locked tokens are held by 3 accounts (in dotted red). Finally, the dash-dotted blue line shows the delegated tokens' distribution where 10 out of 4186 accounts have 57.86% of all delegated COMP tokens available.

themselves to vote or propose changes to the protocol (refer §5.2), but Binance stated that it will not use these tokens to vote on behalf of its users (Shaurya Malwa, 2022).

Takeaway: A significant number of tokens were released at the start, and the amount of unlocked tokens continues to increase over time. A small number of token holders hold the vast majority of all tokens in Compound.

Distribution of token delegation

Delegation is a prerequisite for voting (refer §2.4.1), and Compound allows its participants to delegate their voting rights to others. This ability enables users to delegate their voting power to individuals who share their interests, and allows participants with less voting power to pool their votes together and have a significant voting impact. Users, however, can only delegate *all*, not a fraction, of their tokens. The protocol, nevertheless, enforces this limitation at the wallet address level. Users can own multiple wallet addresses and divide their tokens into them, thereby allowing them to delegate a subset of their tokens to others (Amico, 2023; Fritsch et al., 2022). To determine if delegated tokens are held by a few voters, we group together all inferred addresses (as discussed in §5.1.3) that belong to the same entity and then count the total number of delegated tokens held by each group. We observe, per Figure 5.5, that delegated tokens are concentrated among few voters, and we show the distribution of delegated tokens across several top token holder accounts in Figure 5.6. Out of 4186 COMP delegatee accounts (or accounts

with voting rights), the top 50 (1.19%) hold 99.23% of all delegated tokens, giving them significant decision-making power when voting on proposals. On November 7, 2022, Polychain Capital held the most delegated tokens, with 12.15% (330,986.09) followed by Bain Capital Ventures with 11.85% (322,763.87) and a16z with 9.40% (256,046.13). These three addresses together held 33.41% (909,796.10) of all the 2,723,123.73 delegated tokens in our analysis.

We note that only approximately half of the tokens in circulation are delegated (Figure 5.2). If we investigate token delegation among the top token holders in Figure 5.4, we observe that many of them are crypto exchanges (e.g., Binance and Uniswap V3:COMP) that do *not* delegate their tokens. This observation assuages concerns that crypto exchanges that hold their users token could abuse their users' trust (§5.2). Binance publicly stated that they will not abuse their users' voting rights by voting on behalf of them, and our empirical observations, so far, lend credence to that claim.

5.3.2 Voting on governance proposals

To propose changes to the Compound protocol, an address must have at least 25,000 COMP tokens delegated to it to create a proposal.[19] However, as of September 18, 2021, proposal #60 introduced an exception to this rule, allowing also whitelisted-addresses to create proposals even if they do not have 25,000 delegated tokens (Capital, 2021).

Per Figure 1.2, when a proposal is created, there is an approximately 2-day voting delay period (or 13,140 blocks) that is used to allow the community to discuss the proposal before the voting period begins. During the approximately 3-day voting period (or 19,710 blocks),[20] voters can cast their votes. In order for a proposal to be executed, it needs to meet two requirements. Firstly, it must receive a minimum of 400,000 votes in favor of the proposal. This number corresponds to 4% of the total supply and is known as the *quorum*. Secondly, the majority of the votes cast must be in favor of the proposal. The number of votes each voter has is determined by the number of delegated tokens they held in the block before the voting period began. This prevents voters from changing their delegated tokens after the voting has begun, which could potentially lead to sudden changes in the outcome of the election. After a proposal is approved, it is placed in the *TimeLock* for a minimum period of 2 days before it can be implemented (or executed) (Compound Labs, Inc., 2022a). A proposal can be cancelled at any time by the

[19] Prior proposal 89, an address should have at least 65,000 delegated tokens to create proposals (at Berkeley, 2021).

[20] The duration of the voting period is determined by the number of blocks added to the Ethereum blockchain (specifically, 19,710 blocks). The actual length of the voting period may be slightly longer than 3 days.

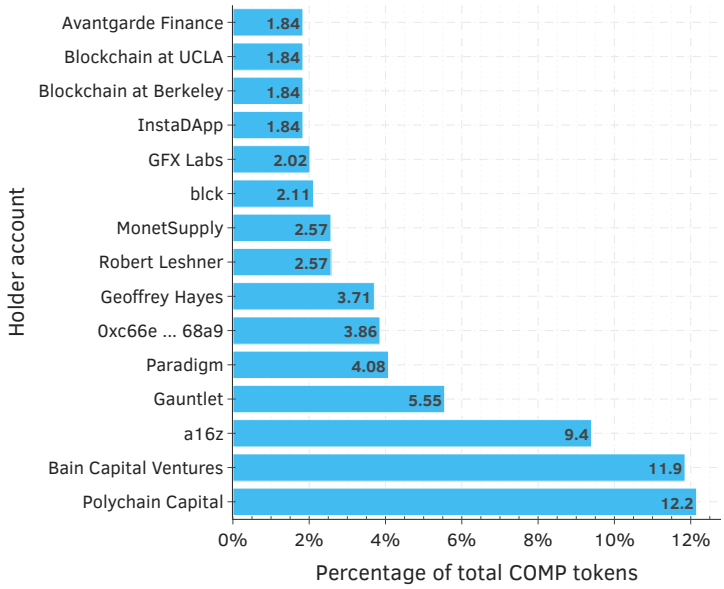

Figure 5.6: Distribution of the top 15 delegated COMP tokens per accounts on November 7, 2022. These addresses have 63.56% of all 2.7 million delegated tokens.

proposer prior to its execution, or by anyone if the proposer fails to maintain at least 25,000 delegated tokens.

In total, 3335 voters cast their votes through 9500 transactions with 8769 (92.31%) for in-favor votes, 644 (6.78%) for against votes, and finally 87 (0.91%) for abstained votes. The majority of voters (51.36%) only voted for 1 proposal. 1% of participants voted for at least 26.66 proposals. On average, participants voted on 2.85 proposals with a *standard deviation (std.)* of 5.23. The address *0x84e3···5a95* voted on the maximum number of proposals (100), followed by *MonetSupply* and *blck* who voted on 96 and 88 proposals, respectively.

Creation of proposals

In total, 33 proposers created the 133 proposals. Of these proposers, 16 (48.48%) created one proposal, while 10% of the proposers created at least 8 proposals. The average number of proposals created per proposer is 4.03 proposals, with a std. of 5.27 and a median of 2. The highest number of proposals was created by *Gauntlet*, who created 24 proposals, followed by *blck*, who created 20 proposals.

The maximum number of proposals were created in March 2022, 11 proposals created (from #86 to #96). However, of those, only 5 were executed, as 1 was defeated and 5 were cancelled (see Figure 5.7). Proposals were submitted, on average, every 6.95 days (std. of 6.41), with a median of 5.08 days. This may be because the proposal lifecycle lasts 7 days, and the voters might not want to consider multiple active proposals at once. The shortest

and longest interval between proposals was 0 and 31.14 days, respectively. Additionally, proposals typically take 1.64 days (std. of 0.72 days) to reach the quorum, as depicted in Figure C.2 in §C.6.

Takeaway: Compound is actively and regularly used: It received a constant stream of proposals over the course of our study period.

Participation in voting

Next, we computed the voting participation per proposal (see Figure 5.8). This metric is calculated by dividing the number of votes (or delegated tokens) cast on a proposal by the total number of delegated tokens eligible to vote on that proposal at the start of the voting period. This is a crucial measurement as it shows the proportion of all delegated tokens that are used in the governance election process by the voters on proposals. Also, protocols with low voter turnout are more susceptible to vote-buying, as non-voting users may sell their voting rights to others (Daian et al., 2018). Our results show that the average Compound voter turnout is 33.25% (with a std. of 17.61%), the median is 32.10%, and the maximum 80.80%. Based on Figure 5.8, we observe higher voting participation for early proposals compared to recent ones, likely due to the limited availability of tokens to a select few in the beginning.

On average, the 133 proposals had 71.43 voters participating in their election, with a standard deviation of 98.97 voters. 50% of the proposals received votes from 38 voters, while the numbers of voters varied between 0 (when proposals are cancelled before the voting period begins) and a maximum of 619, as seen in proposal #111. This particular proposal received a total of 686,289.04 votes from 615 voters in favor, 3 against, and 1 abstention. The next proposals with higher number of voters are proposals #115 and #105 that received votes from 579 and 404 voters, respectively.

Each time a voter casts a vote in the Compound governance protocol by issuing a transaction, an event is triggered, as described in §5.1.1. We analyzed 9500 transactions with events triggered by voters during the voting process. Of these events, 1732 (18.23%) were votes cast by voters who did not have any delegated tokens available, resulting in zero voting power or *useless vote*. Although this is allowed by the protocol, it does not count for or against a proposal. However, it shows support for the proposal, as these voters still participate in the election despite not having any delegated tokens available. Therefore, the average number of votes cast (or tokens used to vote) was 10,961.73, with a std. of 39,212.17 and a median of 0.1. The range of votes cast was from 0 to 345,067.49 as shown in Figure 5.9. This indicates that most of the voters are small players (or accounts with a low amount of delegated tokens).

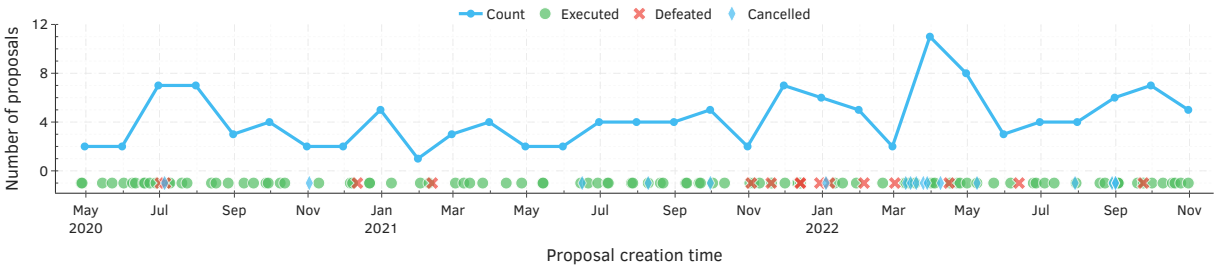

Figure 5.7: Monthly number of Compound proposals created overtime and their respective outcomes (executed, defeated, or cancelled). Proposals are created, on average, every 6.95 days.

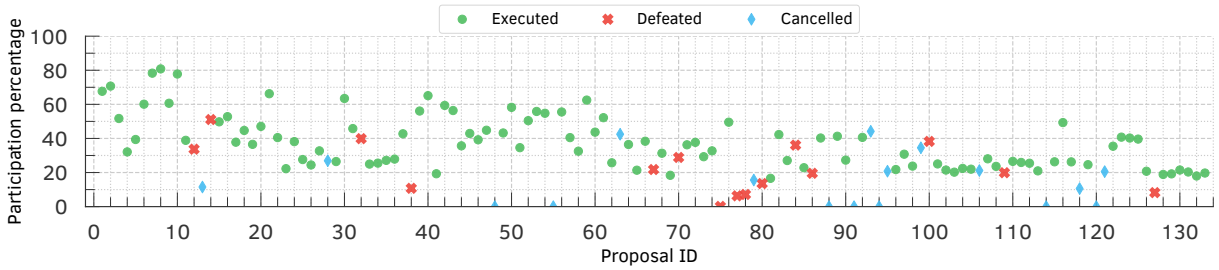

Figure 5.8: Compound's voting participation per proposal in terms of delegated tokens used from all delegated tokens available. Proposals are indicated either as executed (in green), defeated (in red), or cancelled (in blue).

In addition, when voting in Compound, there is a financial cost involved due to the on-chain transactions required to cast votes. To determine these costs, we collected the relevant transactions from the Ethereum blockchain and analyzed the fees paid by voters to issue the transactions and cast their votes. We report the voting cost in US dollars, using the ETH-USD Yahoo Finance data feed (Yahoo Finance, 2023b) to compute the exchange rate at the time the transaction was included in a block. In total, voting for the 133 Compound proposals, voters paid \$74,865.74. The average voting cost per proposal is \$7.88 with a std. of \$22.29. The median voting cost is \$1.48 with a range from \$0.03 to \$294.02. Figure 5.10 shows the voting cost distribution per proposal. We also computed these metrics at proposal level, on average, each proposal costed \$594.17 with a std. of \$745.62 and a median of \$291.92. The cost ranges from \$2.39 to \$4247.25.

Voting on proposals can, hence, present a significant cost barrier, especially for voters with relatively few tokens. In such cases, the cost per token vote (or vote unit) may be too high compared to those with a higher number of delegated tokens. To better understand this, we normalized the cost of casting a vote by the number of votes cast (measured by the total number of delegated tokens available to voters' addresses). For this analysis, we focused on voters who cast at least 10^{-6} votes in any proposal. As shown in Figure 5.11,

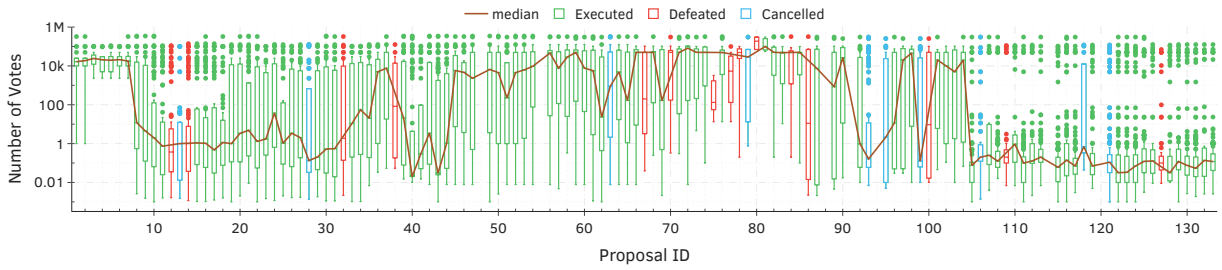

Figure 5.9: Compound's distribution of voting power by voter per proposal. For better illustration, we consider a cutoff of 0.001 votes.

some voters faced prohibitively high costs per vote unit. For example, the cost per vote unit has a mean of $\$358.54$ and a std. of $\$9334.73$, indicating a highly skewed distribution. However, half of the voters faced a cost per vote unit of only $\$6.69$. The cost per vote unit ranged from $\$3.79 \times 10^{-7}$ to $\$725{,}248.10$.

Additionally, we analyzed the number of voters required for all 101 (75.94%) executed proposals in our data set to reach the quorum and pass. Our results show that, for 99 proposals, the average number of voters required for a proposal to reach the quorum and pass was 3.25, with a std. of 1.65. The median number of voters required was 2, and the range of voters required varied from 2 to 8. This sheds light on how centralized these delegated tokens are distributed among a few participants, where for half of the proposals only 2 voter casting their votes would be enough to pass (or execute) a proposal.

Furthermore, we analyzed the number of voters needed for proposals to reach 50% of the total votes cast. Out of 133 proposals, we excluded 7 proposals that were cancelled before the voting period, leaving us with 126 (94.74%) proposals for analysis. On average, those proposals required 2.84 voters with a std. of 0.97 and a median of 3 voters. The minimum and maximum number of voters were 1 and 5, respectively. This again suggests that the token distribution is concentrated among few voters who hold a high voting power. We present the cumulative voting power for the top 10 most powerful voters for each of these 126 proposals in our data set in §C.5.

Margin of victory/defeat

During the analyzed period from March 3, 2020 (block number $9{,}600{,}000$) to November 7, 2022 (block number $15{,}917{,}000$), 133 proposals have been created. Of these, 17 (12.78%) were cancelled and 15 (11.28%) defeated, leaving 101 (75.94%) executed proposals. Figure 5.12 shows the percentage of in-favor, against, and abstain votes for each proposal. The majority of the proposals received significant support from the voters. On average,

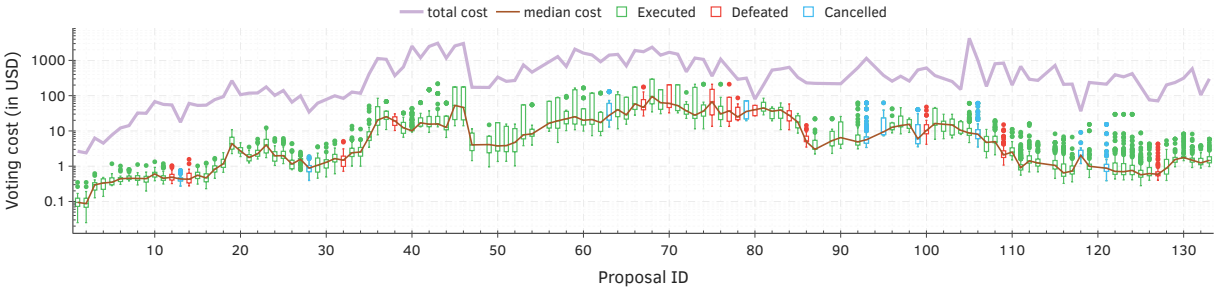

Figure 5.10: Voting cost distribution per proposal. On average, casting a vote costs $7.88 with a std. of $22.29.

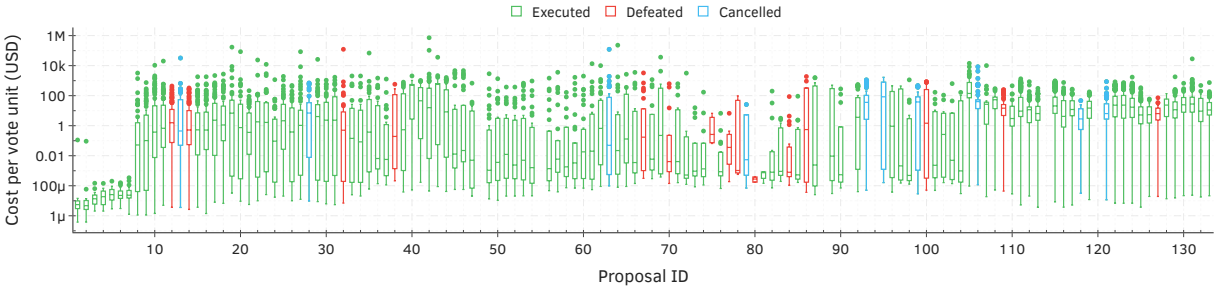

Figure 5.11: Voting cost distribution *normalized* per the voting power. We consider a cutoff of 10^{-6} votes for better illustration.

proposals received 89.39% of the votes in favor, with a std. of 23.98% votes and a median of 99.99%. We highlight the proposals' outcome at each stage of their lifecycle in Figure 5.13. Our analyses show that 7 (5.26%) out of 133 proposals were cancelled right after they were created and, therefore, they had not reached the *Voting Period* meaning they were not available for voting. Next, 4 proposals were cancelled before the *Voting Ends* stage, meaning they were pulled out before the election finished. 2 were also cancelled after they succeeded in the election (after the *Voting Ends* stage) but before they were queued in the *Timelock*. Further, 4 proposals were cancelled when in the *Timelock*. These proposals account for 6 cancelled proposals after they successfully passed, which could indicate a lack of community consensus (Sharma et al., 2023). Finally, 101 (75.94)% proposals were successfully executed. We gathered data from Messari (Messari, 2023) to categorize these executed proposals and report their importance level in §C.1.

Temporal dynamics of voting

Compound Governor does not allow voters to change their votes once they have been cast. This means that voters can only vote once on each proposal. Nevertheless, voters can view all votes that have been cast on-chain in real-time. Thus, understanding how

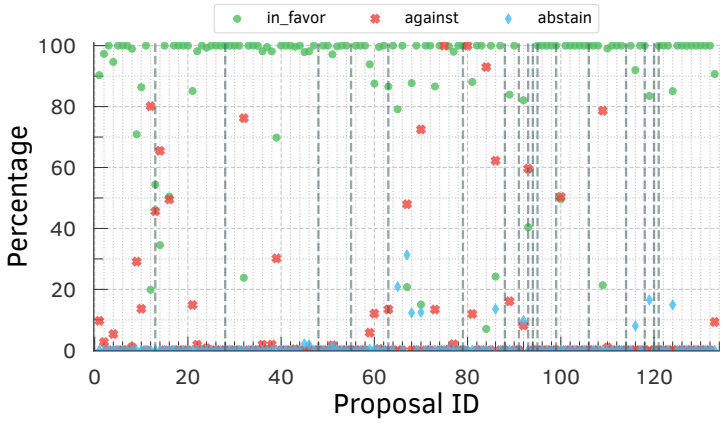

Figure 5.12: Percentage of in-favor (in green), against (in red), and abstain (in blue) votes for each proposal. A total of 15 (11.28%) proposals were defeated, and vertical lines represent 17 (12.78%) cancelled proposals.

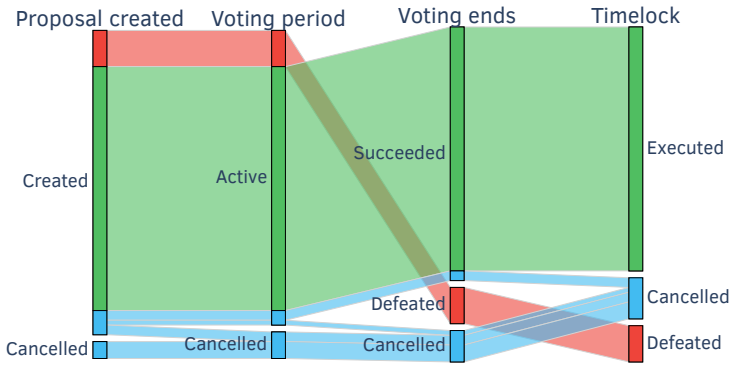

Figure 5.13: Summary of the outcome of 133 Compound proposals at each stage of their lifecycle. There are 101 proposals executed (in green), 15 defeated (in red), and 17 cancelled (in blue).

long it takes voters to cast their votes is interesting because it can shed light on whether they want to wait until the last minute to cast their votes.

According to our analysis, voters take an average of 1.4 days (with a std. of 0.95 and a median of 1.34 days) to cast their votes after the voting period began. The shortest and longest recorded delays in our data set are 0 and 3.39 days, respectively. Figure 5.14 shows the distribution of the time it takes voters to cast their votes for each proposal. We also highlight the voting delays for all votes cast per proposal in §C.5.

When examining voting delay behavior, voters typically take longer to cast votes against proposals (1.58 days on average) in comparison to all other votes (see Figure 5.15a). Considering only executed proposals, voters take longer to abstain but are faster to vote against executed proposal (Figure 5.15b). For defeated proposals, on the contrary, they abstain faster and take longer to vote against defeated proposal (Fig-

ure 5.15c). Even for cancelled proposals, Figure 5.15d shows that voters take longer to vote against these proposals. We believe that the executed proposals must have been better discussed prior to the voting period, and therefore voters were more likely to vote for the proposal with high approval rates (Figure 5.12). Similarly, voters were more likely to vote against proposals that were defeated.

5.3.3 Real-world decision-making using Compound governance

Interestingly, Compound has also been utilized for real-world decision-making purposes, such as allocating grants to contributors (Gauntlet, 2021) or hiring an audit company to review the governance protocol through the Compound code (Sukernik, 2021). For instance, on September 29, a bug was introduced in the Comptroller of the Compound Protocol through proposal #62 that allowed users to claim more COMP tokens than they were entitled to, resulting in a loss of $50 million worth of COMP tokens (Loewen, 2021a; Nick Martitsch, 2021). The Compound community sought to hire, through the Compound governance protocol, a smart contract auditor to audit the protocol (Sukernik, 2021). Three companies, ChainSecurity, OpenZeppelin, and Trail of Bits, posted their business plans for discussion and then created proposals via the Compound Governor. Voters were able to vote for their preferred proposal, and the winning proposal was eventually implemented. The losing proposals would have been cancelled by the community's multi-signature mechanism after the voting period ended, ensuring only one could pass.

OpenZeppelin was the only proposal to reach quorum and get the majority of votes to be implemented. They audited the Compound code, assisted proposers, participated in community discussions, and reviewed any new proposals formally created by the Compound community (OpenZeppelin, 2023).

Takeaway: We believe that these governance protocols will be used even more in the future for transparent decision-making in real-world applications like the ones mentioned above. This will have a positive impact on the use of governance protocols in the everyday life of society.

5.3.4 Voting patterns of delegates

In this section, we analyze the formation of coalitions among voters, where they cast their votes as a group. This analysis is crucial because such behavior may compromise the security of the governance protocol. Specifically, instead of expressing their individual opinions on a proposal, voters may choose to mimic the votes of their peers. The transparency of the Ethereum blockchain used for voting in Compound allows anyone to view the addresses of voters and their corresponding votes (e.g., their voting power and

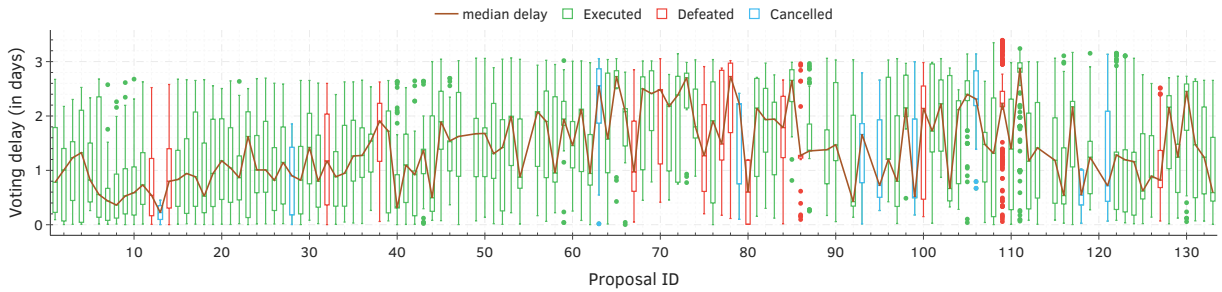

Figure 5.14: Distribution of the number of days it takes voters to cast their votes.

voting preference) during the election process, potentially facilitating this behavior. As a result, exploring the possibility of coalition formation could provide valuable insights into the decision-making patterns of voters. Figure 5.16 shows a heatmap of how each of the top 15 voters cast their votes across all 133 proposals in our data set.

Further, we use cosine similarity to quantify how similar the voting patterns of different voters are. Cosine similarity calculates the similarity between two vectors by determining the cosine of the angle between them (Scikit Learn, 2023; Xia et al., 2015). It is useful in the context of voting because it allows us to compare voting patterns and determine whether and which voters vote for the same proposals. The cosine similarity value ranges from -1 to 1, with a value of 1 indicating a high degree of similarity.

Our analysis shows that the top 3 voters (i.e., 0x84e3$\cdots$5a95, MonetSupply, and blck) have a strong cosine similarity in their voting behavior when casting a vote in favor of a proposal, meaning that they cast their votes similarly (see Figure 5.17). Moreover, Gauntlet, Dakeshi, Robert Leshner, and Arr00 also show a strong similarity with 0x84e3$\cdots$5a95. We also analyzed the voting similarity when voters cast a vote against a proposal. However, we cannot make definitive conclusions regarding abstained votes as they are infrequent: only 87 (0.91%) out of 9500 votes. Regarding votes against proposals, Blockchain at Michigan and Blockchain at Berkeley have the highest cosine similarity with 0.73 followed by blck and Dakeshi with 0.67. These results suggest that these voters have similar voting patterns when indicating their opposition to a proposal.

5.4 Concluding remarks

In this chapter, we analyzed data from the Ethereum blockchain related to Compound, a widely used smart contract. Our analysis is centered on the decentralized governance of Compound, with a particular focus on amendments to the smart contract. We found that the Compound contract is being actively amended—token holders continuously propose amendments that are then voted on by other token holders. We observed a

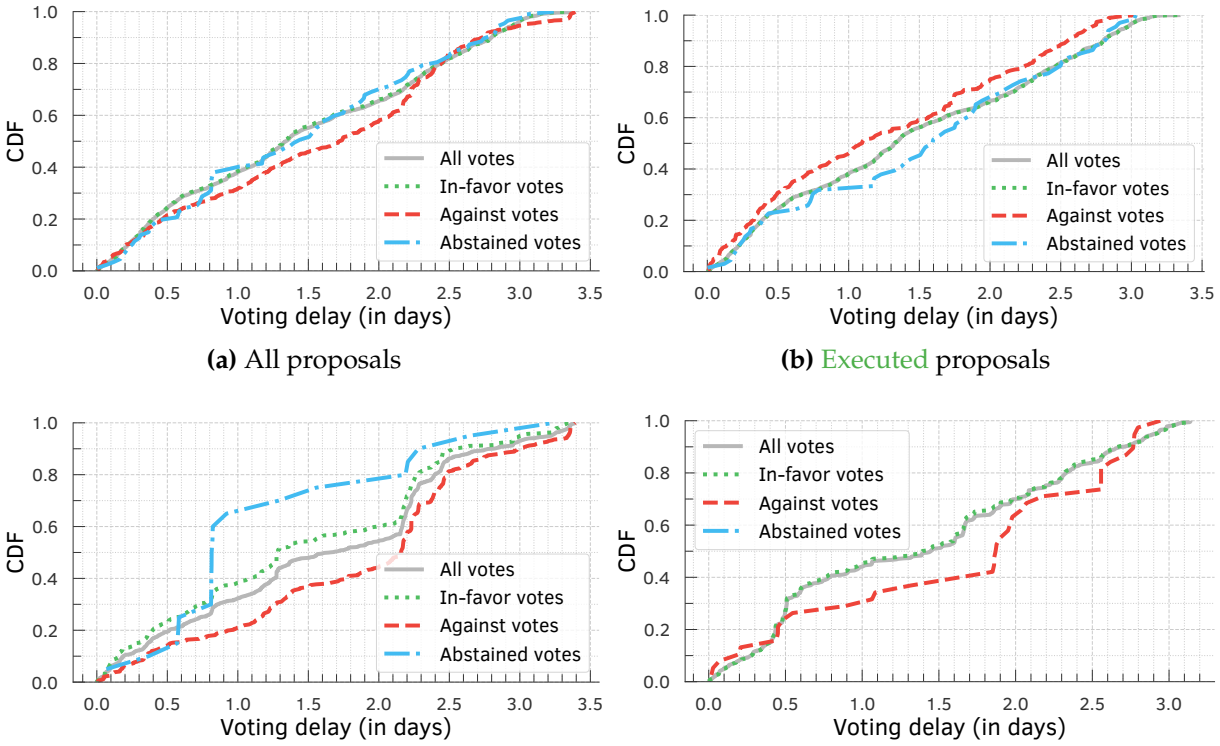

(a) All proposals

(b) Executed proposals

(c) Defeated proposals

(d) Cancelled proposals

Figure 5.15: Cumulative distribution function of the time it takes voters to cast their votes since the voting period began considering: (a) All proposals; (b) Executed proposals; (c) Defeated proposals; and (d) Cancelled proposals.

striking concentration of tokens (be it in terms of their ownership, their delegation, or their voting participation) in the hands of a few participants, which raises serious concerns about the extent to which governance is decentralized in practice. For instance, our analysis shows that, on average, only 3.25 voters were needed for the proposals to reach *quorum* and pass, and only 2.84 voters were needed to reach 50% of the total votes. Our analysis also highlights issues with the Compound use of on-chain voting—in particular, the transaction fees voters must pay to cast an on-chain vote can make it prohibitively expensive for voters with fewer tokens. These costs have implications for voting participation and can affect how voters, proposers, and other stakeholders interact with these protocols.

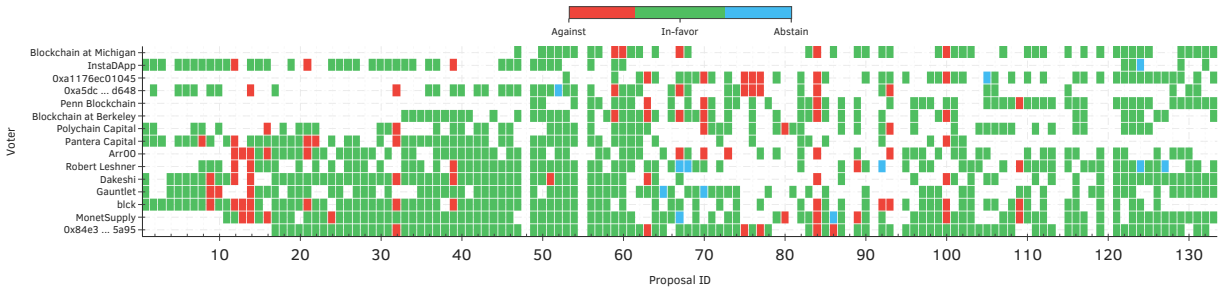

Figure 5.16: Votes cast by the top-15 voters. In-favor votes are in green, against in red, and abstain in blue color.

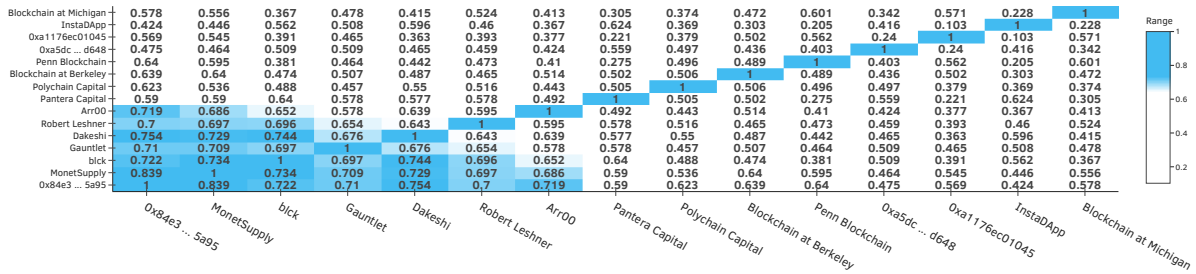

Figure 5.17: Cosine similarity of the top-15 voters voting in-favor a proposal.

Chapter 6

Related Work

In this chapter, we examine the literature relevant to this thesis. We explore three main topics: (i) transaction prioritization norms; (ii) transaction prioritization and contention transparency; and (iii) decentralized governance. The latter encompasses works that explore the distribution of decision-making power for blockchain governance.

6.1 Transaction prioritization norms

A few recent papers proposed solutions to enforce that transaction ordering follows a certain norm, mostly based on statistical tests of potential deviations (Asayag et al., 2018; Lev-Ari et al., 2020; Orda and Rottenstreich, 2019). These works were, however, mostly of theoretical nature in that they did not contain empirical evidence of deviation by miners, but rather assumed that miners might deviate. Prior efforts also proposed consensus algorithms to guarantee fair-transaction selection (Baird, 2016; Kelkar et al., 2020; Kursawe, 2020). Kelkar *et al.* (Kelkar et al., 2020) proposed a consensus property called *transaction order-fairness* and a new class of consensus protocols called *Aequitas* to establish fair-transaction ordering in addition to also providing consistency and liveness. A number of prior work focused on enabling miners to select transactions. For instance, SmartPool (Luu et al., 2017) gave transaction selection from mining pools back to the miners. Similarly, an improvement of Stratum, a well-used mining protocol, allows miners to select their desired transaction set through negotiation with a mining pool (Braiins, 2021b). All these prior work are, again, mostly of theoretical nature. In contrast, this thesis provides empirical evidence of deviation from the norm by miners in the current Bitcoin system.

Additionally, fairness issues have been studied in blockchain from the point of view of miners. Pass *et al.* (Pass and Shi, 2017) proposed a fair blockchain where transaction fees and block rewards are distributed fairly among miners, decreasing the variance of mining rewards. Other studies focused on the security issues showing that miners

should not mine more blocks than their "fair share" (Eyal and Sirer, 2018) and that mining rewards payout is centralized in mining pools and therefore unfairly distributed among their miners (Romiti et al., 2019). Chen *et al.* (Chen et al., 2019) studied the allocation of block rewards on blockchains showing that Bitcoin's allocation rule satisfies some properties. It does not, however, hold when miners are not risk-neutral, which is the case for Bitcoin. In contrast to these prior works, this thesis touches upon fairness issues from the viewpoint of transaction issuers and not miners.

There is a vast literature on incentives in mining. Most of it, however, considers only block rewards (Chen et al., 2019; Eyal and Sirer, 2018; Fiat et al., 2019; Goren and Spiegelman, 2019; Kiayias et al., 2016; Noda et al., 2020; Pass et al., 2017; Romiti et al., 2019; Sompolinsky and Zohar, 2015; Zhang and Preneel, 2019). As the block reward halves every four years in the Bitcoin blockchain, some recent work focused on analyzing how the incentives will change when transaction fees dominate the rewards. Carlsen *et al.* (Carlsten et al., 2016) showed that having only transaction fees as incentives will create instability. Tsabary and Eyal (Tsabary and Eyal, 2018) extended this result to more general cases including both block rewards and transaction fees. Easley *et al.* (Easley et al., 2019) proposed a general economic analysis of the system and its welfare with various types of rewards. Those prior works, however, assume that miners follow a certain norm for transaction selection and ordering (mostly the fee rate norm) and look at miners' incentives in terms of how much compute power to exert and when (or some equivalent metric). There are also prior studies on the security issues of having transaction fees as the prime miners' incentive (Carlsten et al., 2016; Li et al., 2018); and a vast literature on the security of blockchains more generally (e.g., (Gencer et al., 2018; Karame, 2016; Vasek et al., 2014)). Again, however, these studies focus on miners' incentives to mine and not on transaction ordering; for the latter, they assume that miners follow a norm. These prior studies are, hence, somewhat orthogonal to this thesis.

Only a few recent works touched upon the issue of how miners select and order transactions, and how this is interlaced with how the fees are set. Lavi *et al.* (Lavi et al., 2019) and Basu *et al.* (Basu et al., 2019) highlighted the inefficiencies in the existing transaction fee-setting mechanisms and proposed alternatives. They showed that miners might not be trustworthy, but without providing empirical evidence. Siddiqui *et al.* (Siddiqui et al., 2020) showed through simulations that, with transaction fees only as incentives, miners would have to select transactions greedily, increasing the latency for most of the transactions. They proposed an alternative selection mechanism and performed numerical simulations on it. This thesis takes a complementary approach: We analyze empirical evidence of miners deviations from the transaction ordering norm

in the current ecosystem. We also empirically analyze existing collusion at the level of transaction inclusion.

To the best of our knowledge, our study is the first of its kind—showing empirical evidence of norm violations in Bitcoin—and our results help motivate the theoretical studies mentioned above.

6.2 Transaction prioritization and contention transparency

As previously mentioned, recent work analyzed the implications of relying on transaction fees separately (Carlsten et al., 2016) and in conjunction with block rewards (Tsabary and Eyal, 2018), as well as the relationship between such incentives and transaction waiting times (Easley et al., 2019). These prior works assume that transactions are broadcast to all miners and the fees offered is uniform across miners. None of them acknowledge the issue of transparency. Prior work also analyzed the Ethereum fee (i.e., gas price) mechanism to determine the gas price for a given transaction (Antonio Pierro et al., 2020; Liu et al., 2020; Mars et al., 2021; Turksonmez et al., 2021). However, the fee estimation and fee-based prioritization schemes in these studies do not take into account dark-fees or private mining.

Many transaction-accelerator, or front-running as a service (FRaaS), platforms exist for both Bitcoin (BTC.com, 2022; ViaBTC, 2022) and Ethereum (Eskandari et al., 2020; Flashbots, 2022b; SparkPool, 2021). Transaction issuers might resort to such acceleration or off-chain payment channels to hide their true fee from competitors and avoid being front-run (Daian et al., 2020; Strehle and Ante, 2020). Tim Roughgarden (Roughgarden, 2021) discussed the incentives for off-chain agreements (such as dark-fees) between miners and users for first-price auctions and different deviations of the new Ethereum fee mechanism *EIP-1559 protocol* (Buterin et al., 2019a). Roughgarden showed that miners and users cannot strictly increase their joint utility through off-chain payments under EIP-1559 because on-chain bids can be easily replaced by the off-chain bids. However, utility here is only based on the revenue of bidding for block space. The author did not take into account that utility might depend on other factors, such as transaction issuers wanting to keep their actual bids for block space hidden through off-chain payments, which strictly increases their chances of prioritization, as other bidders cannot counter bid, as they are unaware of the bid itself.

There are two work that analyze private mining. Strehe and Ante (Strehle and Ante, 2020) investigated *exclusive mining* (or private mining), where transactions issuers and miners collude to include transactions that have been sent through a private network. In this case, the transactions are not publicly disclosed until they have been included

in a block; besides, the fees can remain opaque to everyone forever, as such off-chain agreements may use fiat currencies. Weintraub *et al.* (Weintraub et al., 2022) measured the popularity of *Flashbots*, the most used private relay network for Ethereum. This thesis, in contrast, extensively investigates private transactions and dark-fees in the context of Bitcoin and Ethereum blockchains. Through active measurements, we empirically show that Bitcoin miners collude and highlight the colluding mining pools. We show that Flashbots bundles are quite prevalent in Ethereum and are mainly used for calling Decentralized Exchanges (DEX) contracts to take advantage of *Maximal Extractable Value (MEV)* opportunities. Finally, we discuss why our findings are still valid after "The Merge"—an Ethereum hard fork deployed on September 15th, 2022 (Ethereum Foundation, 2022a,b).

6.3 Decentralized governance

There is rich literature on decentralized governance and social contracts, decentralized autonomous organizations (DAOs), and on-chain governance protocols. Below, we review prior efforts that is most relevant to this thesis.

6.3.1 Decentralized governance and social contracts

Prior work have studied the potential of blockchain-based (decentralized) governance for replacing centralization in traditional applications and services. Atzori *et al.* discussed, for instance, the extent to which blockchain-based governance can mitigate or replace the centralized and hierarchical societal structures and authorities (Atzori, 2017). Reijers *et al.* examined the relationship between blockchain governance and social contract theory (Reijers et al., 2016). They analyzed the political implications of the blockchain technology and how it follows or deviates from the governance principles established by philosophers such as Thomas Hobbes (Hobbes, 1651), Jean-Jacques Rousseau (Rousseau, 1920), and John Rawls (Chapman, 1971). Chen *et al.* presented the trade-offs between decentralization and performance (Chen et al., 2021). Arruñada and Garicano suggested new forms of "soft" decentralized governance to surpass traditional centralized governance structures (Arruñada and Garicano, 2018). Zwitter and Hazenberg conducted a comprehensive review of governance theory and proposed a re-conceptualization of the term governance that is tailored to DAOs (Zwitter and Hazenberg, 2020). These prior work provide valuable insights into decentralized governance structures, albeit they neither confirm the extent to which their (theoretical) observations hold in real-world implementations nor characterize the behavior of governance protocols deployed today.

6.3.2 Decentralized Autonomous Organizations (DAOs)

Several prior studies analyzed the governance structures of DAOs (Beck et al., 2018; Hassan and De Filippi, 2021; Rikken et al., 2019). Hassan and De Filippi analyzed what DAOs constitute and discuss their key traits (Hassan and De Filippi, 2021). Rikken *et al.* identified various political challenges in governance of blockchains (Rikken et al., 2019). Beck *et al.* (Beck et al., 2018) presented a case study of a DAO in Swarm City (City, 2023), a decentralized commerce platform. A recent work categorized the governance of several blockchains such as Bitcoin, Ethereum, Tezos, Polkadot, and some governance protocols like Uniswap (Adams et al., 2021), MakerDAO (MakerDAO, 2023) and Compound (Leshner and Hayes, 2019) into different types (Kiayias and Lazos, 2023). These invaluable prior work do not, nevertheless, empirically examine the data on existing DAOs to characterize how users interact with on-chain governance smart contracts.

There are three works closely related to ours (Feichtinger et al., 2023; Fritsch et al., 2022; Sharma et al., 2023). Their findings agree with our own, e.g., they too found a high concentration of token delegation among a small number of users. Similarly, they also showed that the largest token holders are more active in voting, further exacerbating the centralization problem. However, while they analyzed voting participation and the cost of voting on the blockchain for more than 10 DAOs, our study presents a comprehensive and in-depth analysis focused on Compound. Specifically, our analysis reveals the complete life cycles of proposals, highlighting how voting behavior evolves over time for different proposals. We also examine token ownership in detail revealing among which entities the tokens held are concentrated as well as how delegations (by individual entities) affect the concentration of tokens. Finally, we discover a vast inequality in voting costs among the token holders and present its implications for decentralized governance.

CHAPTER 7

Discussion, Limitations & Future Work

In this chapter, we discuss some consequential points that follow from the prior chapters, mention the limitations of our work, and explore avenues for future work.

7.1 Transaction ordering

Our findings have significant implications for both bitcoin users and miners. Bitcoin users (using their wallet software) typically assume complete transparency regarding the fees associated with competing transactions when setting fees for their own transactions. However, our results challenge this common assumption. Similarly, the practice of transactions having different confirmation fees for different miners raises notable fairness concerns.

Furthermore, our findings also call for a community-wide debate on defining transaction prioritization norms and enforcing them transparently. Specifically, we highlight three challenging questions that need to be addressed for the future.

What are the desired transaction prioritization norms in public proof-of-work blockchains? What aspects of transactions besides fee rate should miners be allowed to consider when ordering them? For instance, should the waiting time of transactions also be considered to avoid indefinitely delaying some transactions? Should the transaction value (i.e., amount of bitcoins transferred between different accounts) be a factor in ordering, as fee rate based ordering favors larger value over smaller value transactions? Similarly, while we did not find evidence of miners decelerating or censoring (i.e., refusing to mine) transactions, the current protocols do not disallow such discriminatory behaviors by miners. Should prioritization norms also explicitly disallow discriminating transactions based on certain transaction features like sending or receiving wallet addresses? Such norms would be analogous to *network neutrality* norms for Internet Service Providers (ISPs) that disallow flows from being treated differently based on their source/destination addresses or payload.

How can we ensure that the distributed miners are adhering to desired and defined norms? Miners in public proof-of-work blockchains, such as Bitcoin and Ethereum, operate in a distributed manner, over a P2P network. This model of operation results in different miners potentially having distinct, typically different, views of the state of the system (e.g., set of outstanding transactions). Given these differences, are there mechanisms (say, based on statistical tests (Asayag et al., 2018; Lev-Ari et al., 2020; Orda and Rottenstreich, 2019)) that any third-party observer could use to verify that a miner adheres to the established norm(s)?

How can we model and analyze the impact of selfish, non-transparent, collusive behaviors of miners? While the above themes align well with a long-term vision of defining and enforcing well-defined ordering norms in blockchains, in the short term one could focus on examining the implications of the norm violations in today's blockchains. Specifically, how can we characterize the ordering that would result from different miners following different prioritization norms, especially given an estimate of miners' hashing or mining powers (i.e., their likelihood of mining a block). Such a characterization has crucial implications, for example, for Bitcoin users.

7.2 Transaction transparency

In this section, we discuss the implications of transactions prioritization and contention transparency in blockchains. Initially, we highlight the importance of incorporating these aspects into blockchain design to fulfill the overarching goal of transparency. Subsequently, we explore the implications for publicly mined transactions. Then, we delve into the implications for privately mined transactions. Lastly, we emphasize that our implications hold both for before or after the introduction of two major improvements to blockchains: EIP-1559 and the Merge.

Our results show that with private mining and accelerated transactions, the promise of the public decentralized blockchain does not hold. First, through the Bitcoin active experiment, we show that mining pools with combined hash rates of over 50% are colluding with each other, showing a centralization in the system. Further, these accelerated transactions are highly prioritized by the miners and included mostly on top of their blocks. This enables miners to also censor certain transactions, breaking the ethos of decentralized public blockchains with no central authority. Second, it breaks the assumption that all activities in the blockchain are transparent. Although this is true for transactions included in the blockchain, prioritization of transactions is becoming more opaque with the rise of private mining and off-chain fees. Hence, we make the case that to fulfill the transparency promise of public blockchains, prioritization of transactions

should be transparent as well. Third, with private mining in Ethereum, Flashbots is increasingly being used for malicious and predatory activities such as sandwich attacks, which essentially levies a tax on users interacting with financial institutions on the blockchain (e.g., in DEX). These concerns need to be addressed if public blockchains are going to live up to their promises.

Implications for publicly mined transactions. Most wallet software and crypto-exchanges today rely on reconstructing the current public Mempool state in order to suggest a suitable fee to transaction issuers. With the lack of contention and prioritization transparency, transaction issuers can no longer accurately recreate the current Mempool state for different miners. Consequently, they cannot reliably estimate the fees transactions need to pay for their desired prioritization. Worse, as the fraction of privately mined and accelerated transactions keeps rising, the transaction fees will become less (reliably) predictable in the future.

Implications for privately mined transactions. The problem of reliable fee estimation for a desired level of prioritization is even worse for privately mined transactions that are announced on private relay networks. When transaction issuers announce on a private relay network today, they are often unsure what fraction of total network power is controlled by the miners listening to the private relay network. Hence, it is important to estimate the network power controlled by private mining pools to estimate the commit (waiting) times for transactions. Furthermore, transaction issuers on private relay networks are completely blind to other competing transactions. This opacity allows miners offering private mining and transaction acceleration services to overcharge and demand exorbitant fees to commit transactions. For example, in the Ethereum blockchain, users are observed to be overcharged by miners for having their transactions confirmed with high priority through Flashbots bundles (Weintraub et al., 2022).

Relevance of findings in light of EIP-1559 and the Ethereum Merge. Our observations about the lack of transparency and their implications are fundamental to the current blockchain architectures and hold both before and after the recent major improvements to blockchains, e.g., EIP-1559 and the Ethereum Merge. While EIP-1559 attempts to improve the estimation of transaction fees that need to be offered, it does not address the problems associated with the lack of transaction contention and prioritization transparency. Similarly, after the Ethereum Merge, *validators* that stake a certain amount of Ether (ETH) rather than *miners* would be responsible for selecting and validating transactions to include in the next block (Ethereum Foundation, 2022a). Our observations about private mining would still hold for private validation and the implications would still be valid after the Merge.

7.3 Voting power distribution to amend smart contracts

An inherent concern in the governance of blockchain networks revolves around the concentration of governance tokens among a select group of participants. This situation can potentially pose a threat to the protocol and compromise its integrity, especially if the voting power or authority to make important changes is proportional to the amount of tokens held by each participant. This issue was highlighted in the case of Balancer, a decentralized exchange (DEX) built on top of Ethereum. In this example, a user with a significant amount of governance tokens voted for decisions that were beneficial to the user but detrimental to the protocol (Haig, 2022). Therefore, this scenario of a minority holding a significant amount of tokens can lead to a centralization of decision-making power, which is contrary to the goal of decentralizing governance protocols.

While governance protocols in blockchains aim to eliminate (or at least minimize) centralized decision-making, our work reveals that Compound is not effectively achieving its intended goal. The distribution of tokens, which corresponds to voting power, plays a crucial role in determining the level of decentralization in a protocol. Our work highlights the importance of measuring and analyzing governance protocols to ensure that they are working as intended. In addition, this work motivates further research in this area. For example, our empirical evidence supports recent proposals to redefine voting power based on social rewards, such as a voter's reputation or contributions to the protocol (Guidi et al., 2021; Liu et al., 2022; Sharma et al., 2023), or the use of a quadratic voting scheme, where voting power is calculated as the square root of the number of tokens held by voters (Buterin et al., 2019b; Lalley and Weyl, 2018).

In light of our findings, we argue for integrating these insights into the design of future governance protocols. There, we can effectively increase fairness and decentralization within these protocols. In addition, it would also be interesting to analyze other widely used governance protocols, such as Uniswap, to ensure that these governance protocols are truly decentralized.

Conclusion

In this thesis, we adopted a data-driven approach to examine fairness within blockchain contexts, focusing on three key aspects: (i) Fairness in ordering; (ii) Fairness in transparency; and (iii) Fairness in voting power to amend smart contract applications.

Our findings reveal a discrepancy between assumed prioritization norms and actual practices within the blockchain community. In particular, miners often deviate from these norms by prioritizing transactions that serve their own interests or friendly miners. This contradicts the principle of exclusively fee-based prioritization.

Through active experiments, we have uncovered instances of miner collusion involving dark-fee transactions. These transactions provide miners with off-chain incentives in a non-transparent manner, contributing to a lack of transparency in the ecosystem. These fees are kept private between the miner and the issuer of a particular transaction, even after the transaction is confirmed on the blockchain. This exacerbates the challenge of accurately estimating fees. As a result, transaction issuers struggle to determine appropriate fees because they do not have a complete view of all transaction fees being offered.

In addition, blockchain applications, or smart contracts, are often amended by governance protocols. These protocols aim to distribute decision-making power among participants. However, we show that the concentration of voting power based on token ownership skews the dynamics of decision-making. A small subset of participants with a significant token stake wields disproportionate influence, allowing them to shape proposals and votes in line with their self-interest. This practice undermines the true decentralization of decision-making power in the blockchain ecosystem.

We believe that our findings provide valuable insights for designing new and more fair blockchains. Additionally, to ensure the reproducibility of our results, we have made the code and data sets used in this thesis publicly available (Messias, 2023a,b).

Appendices

APPENDIX A

Additional Analysis of Transactions Prioritization Norms

A.1 Congestion in Mempool of data set $\mathcal{B}$

Congestion in Mempool is typical not only in $\mathcal{A}$ (as discussed in §3.2.1), but also in $\mathcal{B}$. Indeed, Figure A.1 reveals a huge variance in Mempool congestion, much higher than that observed in $\mathcal{A}$. Mempool size fluctuations in $\mathcal{B}$ are, for instance, approximately three times higher than that in $\mathcal{A}$ (with the size of unconfirmed transactions at one point in time exceeding almost 50 times the maximum block size). Around June 22^{nd}, there was a surge in Bitcoin price following the announcements of Facebook's Libra[21] and another surge around June 25^{th} after the news of US dollar depreciation (Paul R. La Monica, 2019). These price surges significantly increased the number of transaction issued, which in turn introduced delays. As a consequence, at times, Mempool in $\mathcal{B}$ takes much longer duration than in $\mathcal{A}$ to be drained of all transactions.

A.2 Significance of transaction fees

Table A.1 shows the contribution of transaction fees towards miners' revenue across all blocks mined from 2016 to 2020. In 2018, fees accounted for an average of 3.19% of miners' total revenue per block; in 2019 and 2020 were 2.75% and 6.29%, respectively. However, if we consider only blocks mined from May 2020 (i.e., blocks with a mining reward of 6.25 BTC), the fees account for, on average, 8.90% with an std. of 6.54% in total. Therefore, revenue from transaction fees is increasing (Easley et al., 2017), and it tends to continue.

[21] On June 18^{th}, Facebook announced its cryptocurrency, Libra, which was later renamed to Diem. https://www.diem.com

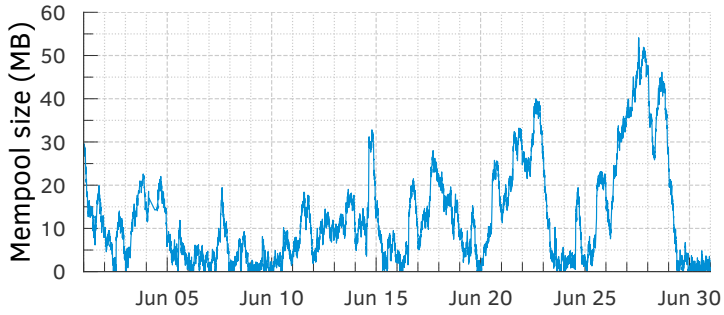

Figure A.1: Mempool size from $\mathcal{B}$ as a function of time.

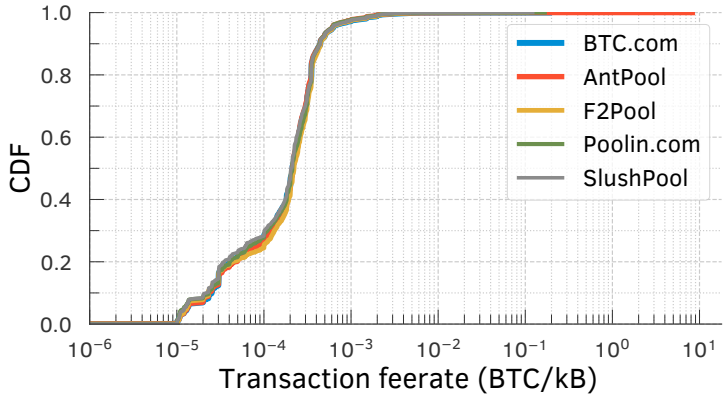

Figure A.2: Distributions of fee rates for transactions committed by the top-5 mining pools in data set $\mathcal{A}$.

Table A.1: *Miners' relative revenue from transaction fees (expressed as a percentage of the total revenue) across all blocks mined from 2016 until the end of 2020.*

Year	*# of blocks*	*mean*	*std*	*min*	*25-perc*	*median*	*75-perc*	*max*
2016	54,851	2.48	2.12	0	0.87	1.78	3.84	92.10
2017	55,928	11.77	7.73	0	6.33	10.49	15.58	86.44
2018	54,498	3.19	5.85	0	0.52	1.22	2.60	44.19
2019	54,232	2.75	2.77	0	0.80	1.81	3.70	24.32
2020	53,211	6.29	6.34	0	1.37	4.00	9.71	39.46

A.3 Transaction fee rates across mining pools

Transaction fee rate of committed transactions in both data sets $\mathcal{A}$ and $\mathcal{B}$ exhibits a wide range, from 10^{-6} to beyond $1\,\mathrm{BTC/kB}$. A comparison of the fee rates of transactions in $\mathcal{A}$ committed by the top five mining pool operators (in a rank ordering of mining pool operators based on the number of blocks mined), in Figure A.2, shows no major differences in fee rate distributions across the different MPOs. Around 70% of the transactions offer from 10^{-4} to 10^{-3} $\mathrm{BTC/kB}$ that is one to two orders of magnitude more

than the recommended minimum of 10^{-5} BTC/kB. We hypothesize that users increase the fee rates offered during high Mempool congestion—they assume that higher the fee rate implies lower the transaction delay or commit time.

A.4 On fee rates and congestion

In Figure A.3, we show the fee rates of transactions observed in 4 different bins or congestion levels in data set $\mathcal{B}$. Each bin in the plot corresponds to a specific level of congestion identified by the Mempool size: lower than 1 MB (*no congestion*), in $(1, 2]$ MB (*lowest congestion*), in $(2, 4]$ MB, and higher than 4 MB (*highest congestion*). Fee rates at high congestion levels are strictly higher (in distribution, and hence also on average) than those at low congestion levels. Users, therefore, increase transaction fees to mitigate the delays incurred during congestion.

Figure A.4 shows that users' strategy of increasing fee rates to combat congestion seems to work well in practice—higher the fee rate, lower the transaction commit delay. Here, we compare the CDF of commit delays of transactions with low (i.e., less than 10^{-4} BTC/kB), high (i.e., between 10^{-4} and 10^{-3} BTC/kB), and exorbitant (i.e., more than 10^{-3}) fee rates, in data set $\mathcal{B}$. The commit delays for transactions with high fee rates (i.e., greater than 10^{-3} BTC/kB) are significantly smaller than those with low fee rates (i.e., lesser than 10^{-4} BTC/kB).

A.5 Child-Pays-For-Parent (CPFP) Transactions

Given any block B_i that contains a set of issued transactions $T = \{t_0, t_1, \cdots, t_n\}$, where each transaction has at least one transaction input identifier $V = \{v_0, v_1, \cdots, v_m\}$, the transaction $t_j \in T$ is said to be a *child-pays-for-parent transaction (CPFP-tx)* if and only if there exists at least one input $v_k \in V$ that belongs to T. In other words, a transaction is a CPFP transaction if and only if it spends from any previous transaction that was also included in the same block B_i.

A.6 Miners' behavior during the scam

To examine the miners' behavior during the Twitter scam attack from July 14$^{\text{th}}$ to August 9$^{\text{th}}$, 2020, we selected all blocks mined (3697 in total, containing $8{,}318{,}621$ issued transactions) during this time period from our data set $\mathcal{C}$. If we rank the MPOs responsible for these blocks by the number of blocks (B) mined (or, essentially, the approximate hashing

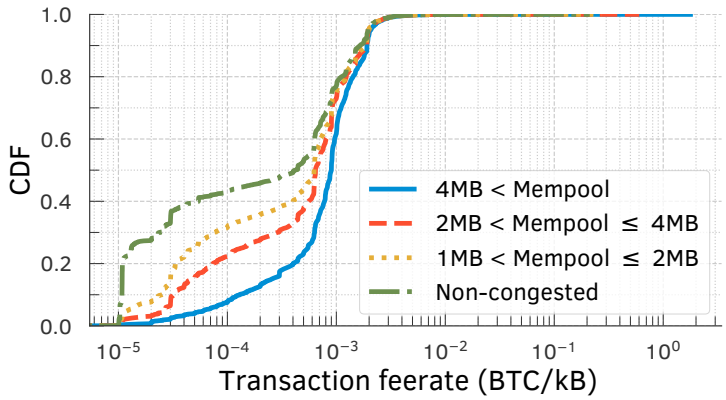

Figure A.3: Distribution of fee rates for transactions in data set $\mathcal{B}$ issued at different congestion levels clearly indicate that users incentivize miners through transaction fees.

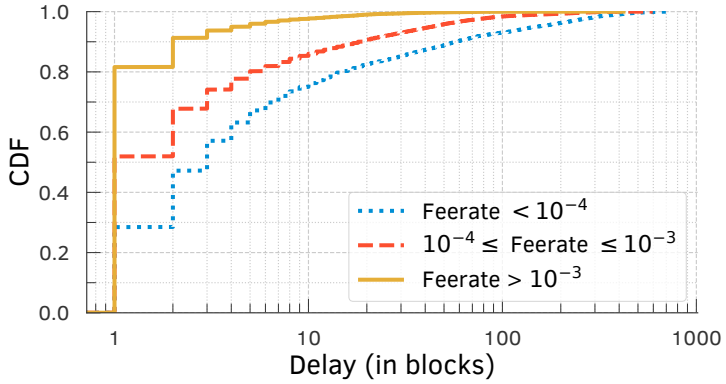

Figure A.4: Distributions of transaction-commit delays in data set $\mathcal{B}$ for different transaction fee rates.

capacity h), the top five MPOs (refer Figure A.5) turn out to be Poolin (B: 565; h: 15.28%), F2Pool (B: 536; h: 14.5%), BTC.com (B: 424; h: 11.47%), AntPool (B: 404; h: 10.93%), and Huobi (B: 353; h: 9.55%).

A.7 Transaction-acceleration fees

In this experiment, we compare the transaction-acceleration fee with the typical transaction fees in Bitcoin. To this end, we retrieved a snapshot containing 26,332 unconfirmed transactions from our node's Mempool on November 24$^{\text{th}}$ 2020 at 10:08:41 UTC. Then, for each transaction, we searched its respective transaction accelerator price (or acceleration fee) via the acceleration service provided by BTC.com (BTC.com, 2022). We inferred the acceleration fees for 23,341 (88.64%) out of the 26,332 unconfirmed transactions. Figure A.6 shows the CDF of both the Bitcoin transaction fees and the acceleration fees

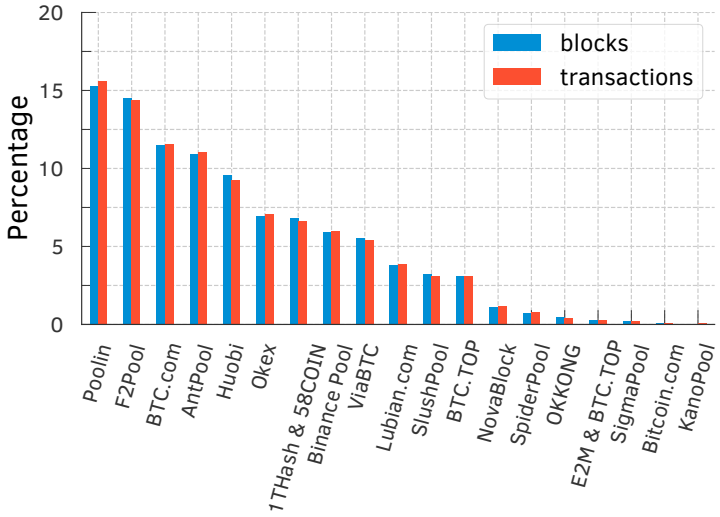

Figure A.5: Distribution of blocks mined and transactions confirmed by different MPOs during the Twitter Scam attack from July 14th to August 9th, 2020.

provided by BTC.com. Acceleration fee is on average 566.3 times higher (4734.67 of std.) and on median 116.64 times higher than the Bitcoin transaction fees. At the time of this experiment, 1 BTC was worth 18,875.10 USD.

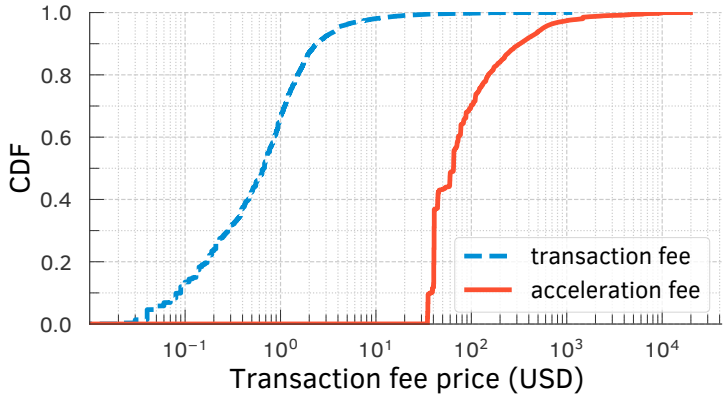

Figure A.6: Fee price comparison between the transaction fee and the acceleration services from an snapshot of our Mempool on November 24th, 2020. Acceleration service provided by BTC.com is on average 566.3 times higher (4734.67 of std.) and on median 116.64 times higher than the Bitcoin transaction fees. The minimum is 0.54, the 25-perc is 51.64, and the 75-perc and the maximum are 351.8 and 428,800, respectively.

Appendix B

Additional analysis of transactions prioritization and contention transparency

B.1 Ethereum private transaction experiment

We conducted 4 active experiments where we issued 8 Ethereum transactions; half issued publicly and the other half privately through a private-channel network known as Taichi Network (SparkPool, 2021). Table B.1 summarizes the transactions in our experiment. Spark Pool and Babel Pool included all private transactions (2 transactions each) sent directly to these miners through Taichi Network.

B.2 Liquidation with Chainlink oracle updates

In AAVE, of 1154 bundles, 994 (86.14%) include one Chainlink oracle update followed by a liquidation. There are 52 (4.51%) with two oracle updates followed by liquidations. Out of 1301 oracle updates bundled with liquidations, 282 (21.68%) are USDC-ETH, 203 (15.60%) are USDT-ETH, 169 (12.99%) are DAI-ETH, 70 (5.38%) are SUSD-ETH, and 60 (4.61%) are LINK-ETH. In Compound, of 641 bundles, 548 (85.49%) contain one Chainlink oracle update followed by one liquidation, while 39 (6.08%) include two oracle updates followed by liquidations. Out of 751 oracle updates bundled with liquidations, 311 (41.41%) are ETH-USD, 128 (17.04%) are BTC-USD, and 53 (7.06%) are UNI-USD.

Table B.1: We conducted 4 active experiments in Ethereum by simultaneously accelerating transactions privately and publicly via Taichi Network. Private transactions were included only by Spark Pool and Babel Pool. If we rank these mining pools according to their hash-rate, they account for 27.72% of the total Ethereum hash-rate.

#	type	tx hash	block number	miner	tx. position per # of txs.	block delay (in blocks)	fee paid (in Ether)	base fee (Gwei)	max fee (Gwei)	max priority fee (Gwei)	gas price (Gwei)	block timestamp in UTC
1	public	bbe88e···a4f000	13,183,516	Nanopool	305 / 336	1	0.00190489	88.98082939	116.52835749	1.72836605	90.70919543	2021-09-08 06:39:18
	private	c46b75···ead538	13,183,520	Babel Pool	29 / 39	5	0.00225209	105.51391459	120.56586232	1.72836605	107.24228063	2021-09-08 06:40:29
2	public	6d994f···c1aadd	13,183,561	Binance	209 / 213	2	0.00244137	114.95482846	137.64014705	1.30100683	116.25583529	2021-09-08 06:49:26
	private	a4d4ae···42ebf5	13,183,565	Spark Pool	294 / 296	6	0.00240978	113.45059961	137.64014705	1.30100683	114.75160643	2021-09-08 06:50:12
3	public	725743···0a6c45	13,183,634	Unknown	124 / 126	2	0.00263298	123.27216185	135.21393222	2.10805685	125.38021870	2021-09-08 07:06:31
	private	f2beec···15cdf1	13,183,635	Spark Pool	321 / 340	3	0.00257468	120.49562077	135.21393222	2.10805685	122.60367762	2021-09-08 07:06:44
4	public	e21695···2c1574	13,183,679	Ethermine	280 / 302	13	0.00223433	104.69510748	108.95262574	1.70164453	106.39675202	2021-09-08 07:18:37
	private	4c482b···87c76f	13,183,690	Babel Pool	150 / 212	24	0.00179917	83.97323655	108.95262574	1.70164453	85.67488108	2021-09-08 07:20:12

B.3 Hashing rates of mining pools

Per Figure B.1, the hash rates of Bitcoin mining pools such as BTC.com, F2Pool, and AntPool alone accounted for almost half the total hash rate of the network around May 2018, and roughly a year later, i.e., from March 2019, together with Poolin the four mining pools alone represent more than 50% of the total network hash rate. At the end of 2020, new MPOs, e.g., Lubian.com and Binance Pool, started mining Bitcoin, which help improve the decentralization of Bitcoin. However, BTC.com, F2Pool, AntPool, and Poolin still account for almost half of the hash rates showing that a few mining pools control a considerable portion of the Bitcoin hash rate.

Hash rates of Ethereum mining pools, in contrast to Bitcoin, do *not* show a high variance (see Figure B.2). We also observed that Spark Pool, the second-largest Ethereum mining pool, suspended their mining services on September 30, 2021, due to regulatory requirements in response to Chinese authorities (Helen Partz, 2021).

B.4 Bitcoin transaction acceleration experiment

We ran an active Bitcoin transaction acceleration experiment where we paid 205 EUR to ViaBTC (ViaBTC, 2022) to accelerated 10 transactions from 10 different snapshots of our Mempool. To select these transactions, we checked whether the Mempool was congested (i.e., having more transactions waiting for inclusion than the next block would be able to include), with its size being at least $8\,\mathrm{MB}$. Then, we considered only transactions with low fee rates—less than or equal to 2 sat-per-byte—to ensure that these transactions would be highly unlikely to be included soon in a subsequent block. Next, we sorted the remaining transactions by size to limit the experiment cost as the acceleration-service costs grow proportional to the transaction size. Finally, we select the transaction with the smallest size in bytes for our active experiment.

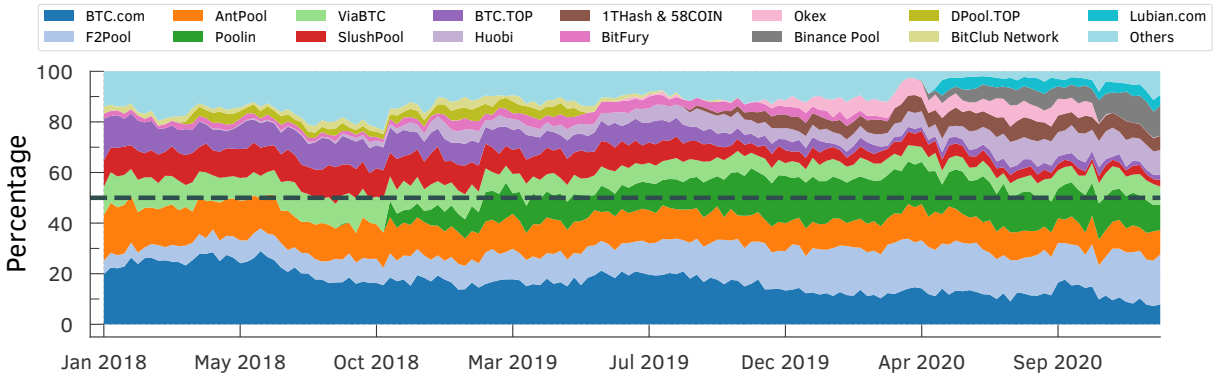

Figure B.1: Monthly Bitcoin hash rate over the 3-year period.

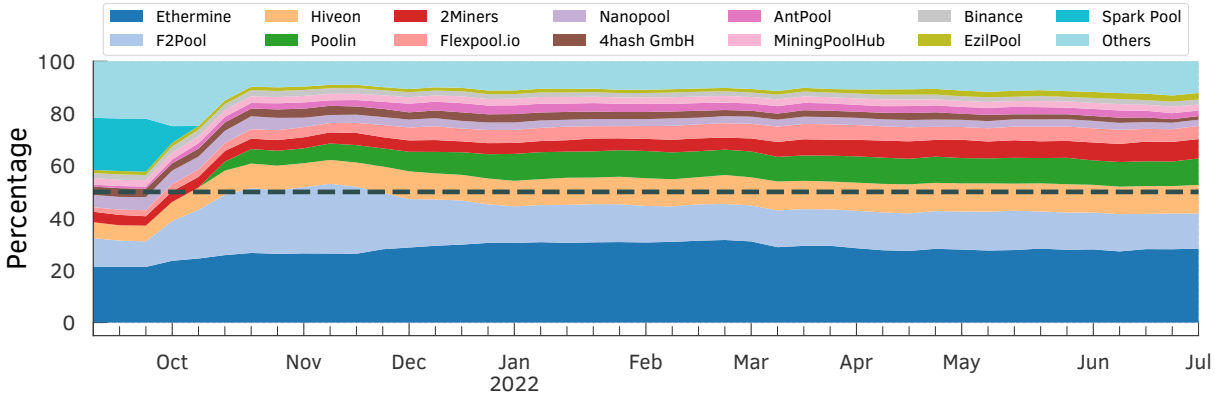

Figure B.2: Weekly Ethereum hash rate from Sept 8^{th}, 2021, to Jun 30^{th}, 2022.

Most of these 10 accelerated transactions were included nearly in the next block, demonstrating the acceleration efficiency. Also, these transactions were wrongly positioned in the block: They appeared, for instance, at the top of the block, i.e., higher than the non-accelerated transactions, showing that miners indeed prioritized them (see Table 4.3). Further, we observed that although we had only accelerated transactions via ViaBTC, other top mining pools were also involved in confirming the accelerated transactions.

Table B.2 shows the transactions used in our experiments. At the time we conducted our experiments, if we rank the miners whose included these transactions based on their daily hash-rate power as (D) and weekly hash-rate power as (W), we would have Huobi (D: 8.1%; W: 9.3%), Binance (D: 9.6%; W: 10.3%), F2Pool (D: 19.9%; W: 18.7%), AntPool (D: 12.5%; W: 10.6%), ViaBTC (D: 5.1%; W: 7.1%). Together these mining pools corresponds to a hash-rate power of (D: 55.2%; W: 56%). Figures B.3 and B.4 show the hash-rate of mining pools in the active experiment and considering the passive experiment (inferred to be accelerated by BTC.com API), respectively.

Table B.2: We conduct 10 transaction acceleration experiments in Bitcoin. If we rank the miners whose included these transactions based on their daily hash-rate power as (D) and weekly hash-rate power as (W), together these mining pools corresponds to a hash-rate power of (D: 55.2%; W: 56%).

txid	*block height*	*miner*	*tx. position*	*delay (in blocks)*	*acc. cost (BTC)*	*vsize (byte)*	*fee rate sat-per-vsize*	*Mempool # of txs.*	*Mempool vsize (MB)*	*timestamp in UTC*
35b18e···52dbc1	658,805	Huobi	2nd	2	0.001254	110	2	36,644	44.63	2020-11-26 19:10
65765c···baede2	658,898	F2Pool	73rd	1	0.001254	110	2	20,998	32.55	2020-11-27 11:06
0c2098···29fbf0	658,912	AntPool	2nd	2	0.001254	110	1	30,126	38.01	2020-11-27 13:38
1515a7···179af3	658,971	Binance	2nd	3	0.001254	110	1	25,922	37.89	2020-11-27 21:55
48a0a5···0ddaec	659,335	ViaBTC	3rd	1	0.001045	110	1	15,605	9.82	2020-11-30 10:09
9a17cf···f3734c	659,341	Huobi	2nd	2	0.001045	110	1	14,945	9.41	2020-11-30 10:28
831b24···95d421	659,351	AntPool	2nd	1	0.001045	110	1	10,990	8.66	2020-11-30 12:22
1f59bf···47096c	659,355	F2Pool	111th	3	0.001045	110	1	17,093	11.40	2020-11-30 12:58
6942e0···8c06c3	659,362	Huobi	2nd	2	0.001045	110	1	30,836	19.06	2020-11-30 14:49
8e49e2···ae825f	659,481	ViaBTC	6th	1	0.001254	110	2	30,935	22.59	2020-12-01 10:40

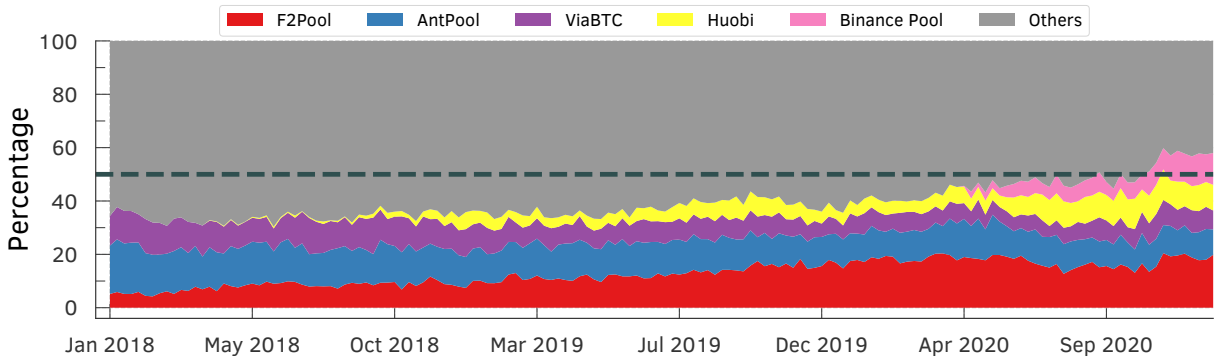

Figure B.3: Active vs. others experiment: Bitcoin mining pools in the active experiment (i.e., mining pools that included transactions accelerated by ourselves) increased their hash rate in 2020. Together, they accounted for more than 55% of the overall hash rate. The plot shows the weekly average percentage of the mining pool's hash-rate over 3 years.

Furthermore, BTC.com (BTC.com, 2022), one of the leading Bitcoin mining pools, provides transaction acceleration services and allows users to verify if transactions have been accelerated through their platform or partner services. From our dataset, we selected those with a SPPE greater than or equal to 1% (12,983,282 transactions in total) and checked if they were said to be accelerated by BTC.com's API. Of these transactions, 14,104 were found to have been accelerated. Our findings also show that transaction acceleration services are becoming quite common among Bitcoin mining pools (as shown in Figure B.5). Between 2018 and April 2019, only BTC.com and F2Pool alone accounted for most of the accelerated transactions. However, as of December 2020, we see that BTC.com accounts for a very small fraction of accelerated transactions, with AntPool, Huobi, and F2Pool accounting for most of the accelerated transactions.

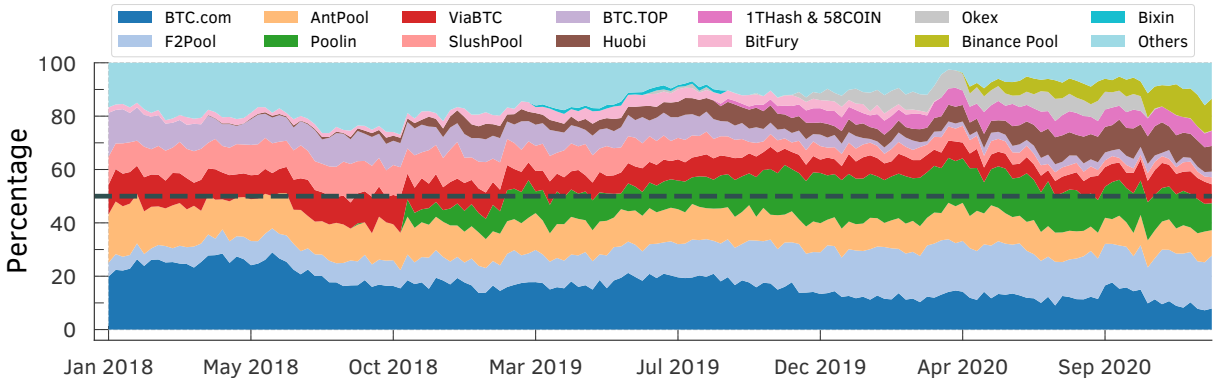

Figure B.4: Passive + active vs. others experiment: Bitcoin mining pools in the active experiment (i.e., mining pools that included transactions accelerated by ourselves) and passive experiment (mining pools that included transactions inferred to be accelerated using the BTC.com API) increased their hash rate in 2020. The plot shows the weekly average percentage of the mining pool's hash-rate over 3 years.

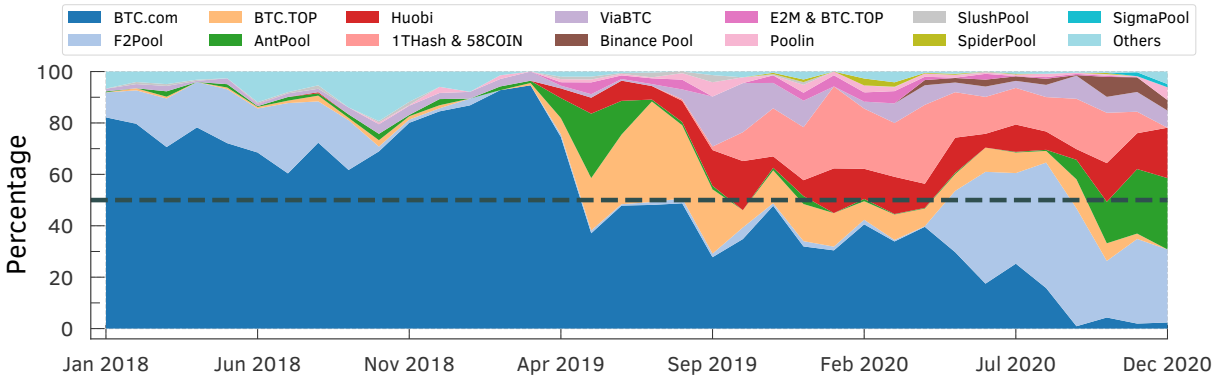

Figure B.5: The plot shows the monthly average percentage of accelerated Bitcoin transactions inclusion by each mining pool over 3 years. Transaction acceleration services or simply Front-running as a Services (FRaaS) are becoming popular across all mining pools.

Appendix C

Additional Analysis of Distribution of Voting Power

C.1 Compound proposals categorization

We gathered data from Messari (Messari, 2023) to determine the categories, subcategories, and the level of importance associated with each Compound proposal. Figure C.1 shows the distribution of 101 executed Compound proposals across different categories and subcategories. We show the degree of importance for each proposal according to Messari divided into "low", "medium", "high", and "very high". As a result, a few proposals categorized as "Parameter Change" and "Security" demonstrate a high level of importance. Furthermore, proposals with the highest level of importance are found within the "Security" category, specifically within the "Mining and Validation" subcategory. This refers to the proposal 64 that was created to fix a bug introduced by proposal #62 (Loewen, 2021a,b).

The majority of the proposals (61 proposals, accounting for 60.4%) are related to "Parameter Change" followed by "Team and Operations" and "Token Supply" accounting for 10 (9.9%) each, and "Governance" with 7 (6.93%) proposals. According to the level of importance reported by Messari, out of the total of 101 executed proposals, 51 proposals (50.5%) are classified as low importance, 46 proposals (45.54%) as medium importance, 3 proposals as high importance, and 1 proposal as very high importance.

C.2 Filtering events to construct our Compound data Set

This section describes the details required to filter and collect transactions data that triggered events of interest from any smart contract on the Ethereum blockchain. Before creating a filter, we need the address of our target contract and its Application Binary

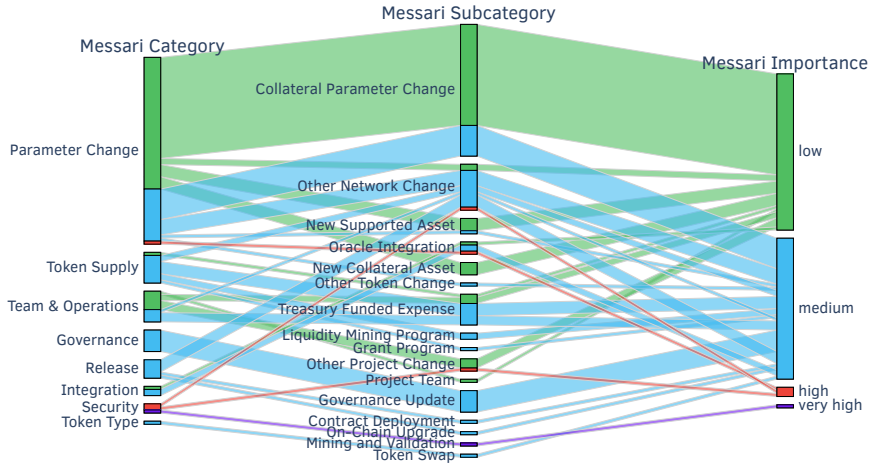

Figure C.1: Categorization of executed proposals. Most of the proposals (60.4%) are related to "Parameter Change". We also show the importance level (low in green, medium in blue, high in red, and very high in purple color) for each proposal according to Messari (Messari, 2023).

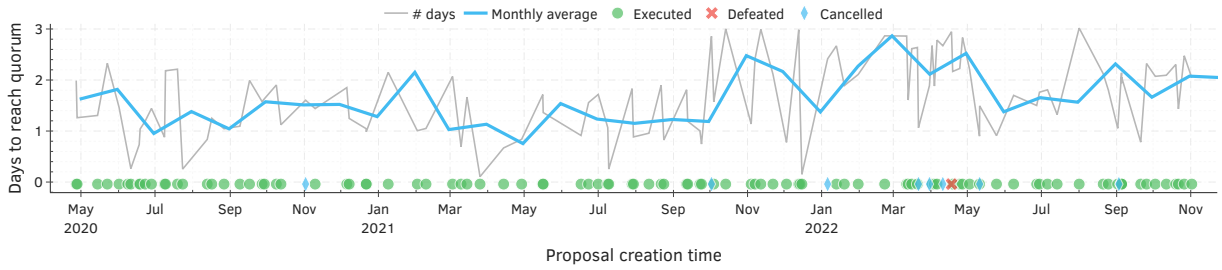

Figure C.2: Compound proposals typically reach the quorum after 1.64 days on average.

Table C.1: A comparison of voting mechanisms in decentralized governance protocols such as AAVE (AAVE, 2023), Balancer (Balancer.fi, 2023), Compound (Leshner and Hayes, 2019), Convex Finance (Convex, 2023a), Curve (Curve, 2023), Maker (MakerDAO, 2023), and Uniswap (Adams et al., 2021). *SC* stands for smart contract.

Protocol	*Type*	*Voting*	*Who can vote?*	*Delegation*	*Voting Aggregation*	*How proposals are implemented*
AAVE	Lending	on-chain	addresses with delegated tokens	yes	on-chain	on-chain via an SC call.
Balancer	DEX	off-chain	stakers with locked tokens	yes (off-chain)	off-chain	via 6-of-11 multisig.
Compound	Lending	on-chain	addresses with delegated tokens	yes	on-chain	on-chain via an SC call.
Convex Finance	Yield Farming	off-chain	stakers with locked tokens	yes (off-chain)	off-chain	via 3-of-5 multisig.
Curve	DEX	on-chain	holders	yes	on-chain	on-chain through an SC call.
Maker Executive	Stablecoin	on-chain	holders	no	on-chain	New Governance Contract requires more MKR staked than previous.
Maker Polling	Stablecoin	on-chain	addresses with delegated tokens	yes	off-chain	Engineers at Maker create the governance contract based on the voting outcome.
Uniswap	DEX	on-chain	addresses with delegated tokens	yes	on-chain	on-chain via an SC call.

Interface (ABI). The ABI is a JSON file that specifies the functions available in the contract, their signatures, and the available events. We can retrieve this information by calling the Etherscan API (Etherscan, 2023a). Once we have the contract address and ABI, we can create a filter to track the contract's activity on the Ethereum blockchain using an important Python library for interacting with Ethereum nodes called Web3.py (web3.py team, 2022) to facilitate the communication with our node's API.

The Web3.py library provides a filtering function called *createFilter*. This function can be used to filter transactions that triggered events of interest from a specific contract within a range of block numbers. We use this function to efficiently collect all transactions that triggered these events from the Compound (Leshner and Hayes, 2019) smart contract.

C.3 Inferring wallet addresses ownership

We aim to identify the ownership of public wallet addresses on the Ethereum blockchain. Due to the inherent anonymity of blockchain addresses, this proves to be a challenging task as we can only know the owners of an address if the owner chooses to disclose it. However, popular blockchain explorers such as Etherscan (Etherscan, 2023b) often provide information on the top holders of specific cryptocurrencies, which allows us to partially overcome this obstacle.

Then, we first obtained the lists of the top 10,000 Ether holders from which there are 290 (2.9%) identified addresses and the top 1000 COMP holders from which there are 82 (8.2%) identified addresses from Etherscan. By comparing these lists to our data set, we were able to identify most of the top COMP holder addresses in our sample. However, this method did not work for the top delegated accounts, as most of them were not included in the list of top COMP holders on Etherscan. This means that most of the delegated accounts does not hold many tokens. Further, we also used the list of top 100 delegated Compound addresses by voting weight available on the Compound website (Compound Labs, Inc., 2022b) from which there are 66 identified addresses.

Furthermore, to extend the available identified addresses in our analysis, we obtained the addresses of 2783 identified users from the Sybil-List (Sybil, 2023b), a project maintained by Uniswap that uses cryptographic proofs to verify wallet addresses ownership. By combining the identified addresses from both sources, we were able to obtain the ownership of 3191 inferred public wallet addresses to use in our analysis. We were able to infer 114 (3.41%) out of the 3341 unique addresses in our data set. Considering the top 10 most powerful voters for each proposal (refer to Figure C.3 in §C.5), we were able to infer 67 (50.37%) of the 133 unique addresses. Overall, our methodology allowed us to partially overcome the anonymity of public wallet addresses on the Ethereum blockchain and shed light on the ownership of these addresses in our data set. Finally, as an entity can control more than one address, we grouped the addresses we identified belonging to the same entity together to conduct our analysis.

C.4 Types of existing governance protocols

There are various smart contract applications that utilize decentralized governance protocols for decision-making, including those for lending, decentralized exchanges (DEXes), and stablecoins, among others. An example of such protocols can be found on the Ethereum blockchain, where a number of these applications are available. We have selected some of the most protocols that use decentralized governance for decision-making. Table C.1 presents 8 protocols, including Maker Executive and Maker Pooling, which are part of the MakerDAO (MakerDAO, 2023) stablecoin protocol responsible for the DAI token. These protocols use decentralized governance mechanisms, and we characterize them based on whether their votes are cast on- or off-chain, the delegation methods they use, how they aggregate the votes, and how the proposal outcome take effect.

C.5 How voters cast their votes

This section examines how each of the top-10 voters of Compound and Uniswap cast their votes. Some proposals may not have received any votes if they were cancelled before the voting period began. See §5.3.2 for details. Figure C.3 shows how each of the top-10 voters cast their votes in each of the 126 (94.74%) out of 133 Compound proposals.

Figure C.4 shows the all votes cast in chronological order per proposal. On average, voters took 1.4 days (with a standard deviation of 0.95 and a median of 1.34 days) to cast their votes after the voting period began.

C.6 Time until reaching the quorum in Compound

For a proposal to pass, it must receive a majority of in favor votes and at least 400,000 (4%) votes in favor from the total supply of Compound tokens. This minimum number of in favor votes is referred to as the *quorum* and is defined by the Compound Governor Bravo contract.

We analyzed how long it takes for these proposals to reach the required quorum. Figure C.2 shows the number of days it took each of the evaluated Compound proposals to reach the quorum. On average, it takes 1.64 days with a standard deviation of 0.72 days for the proposals to reach the quorum. The cumulative distribution function of our results, where 32% take more than 2 days to reach the quorum. The shortest and longest time it took was 0.11 and 3 days, respectively.

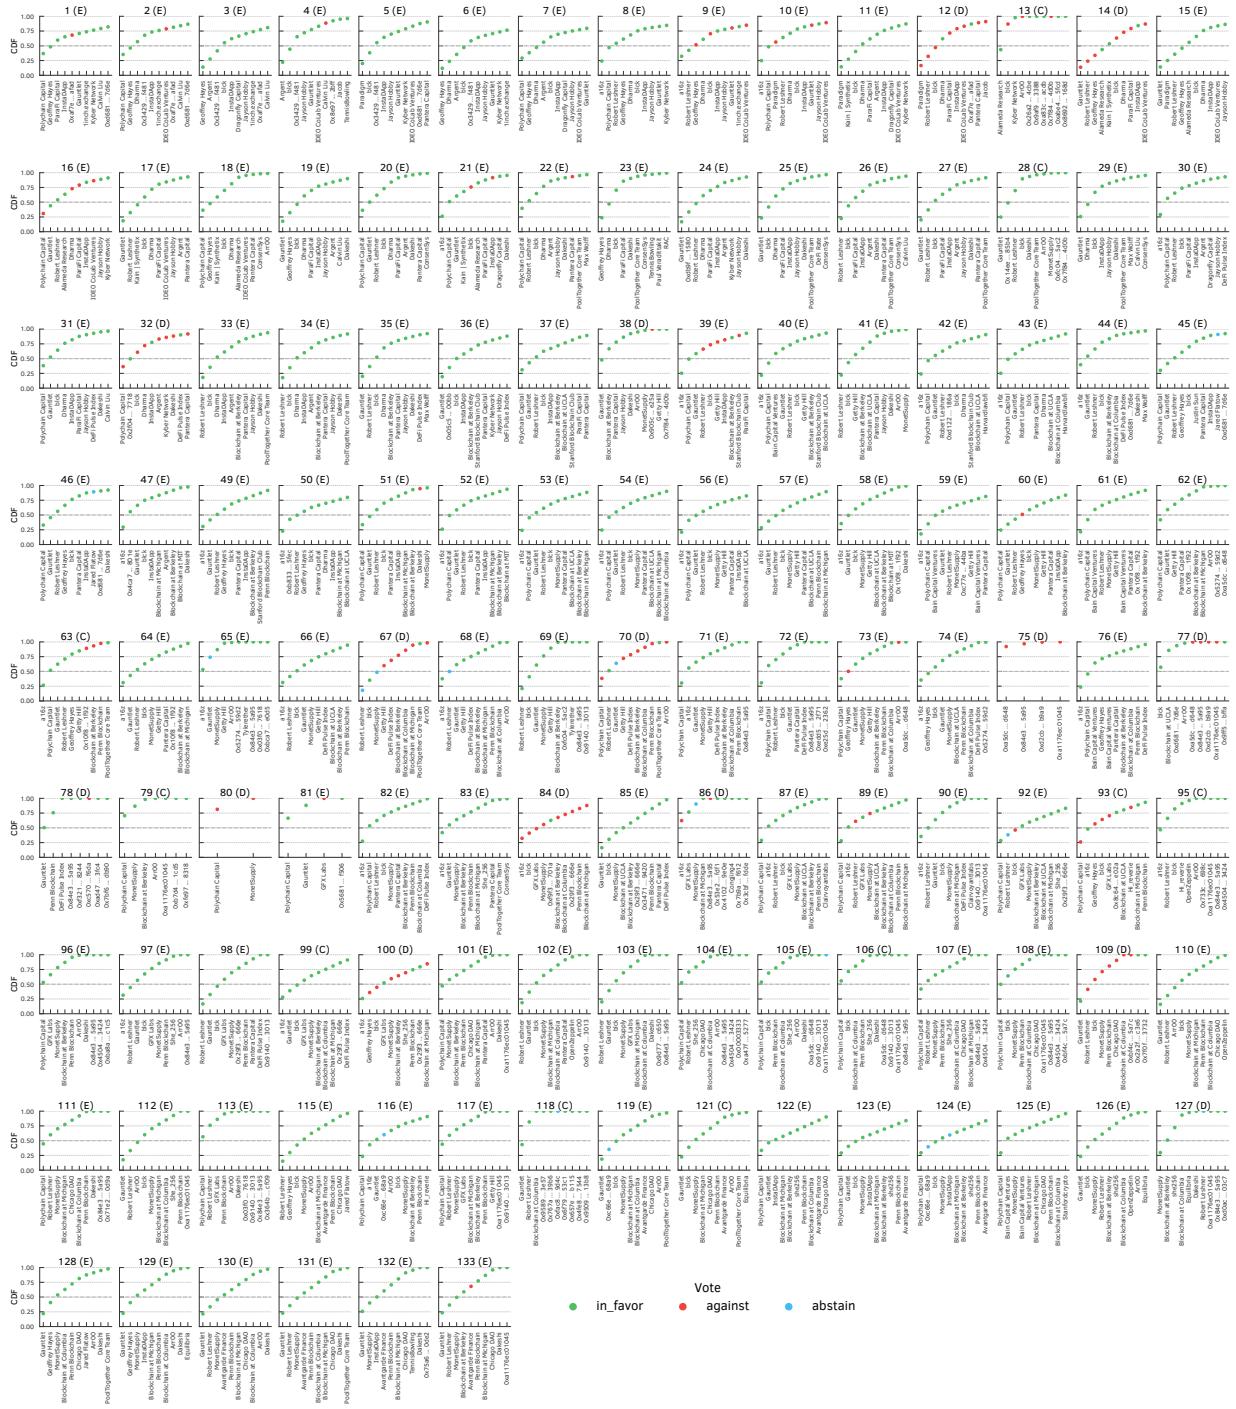

Figure C.3: Cumulative voting power distribution of the top-10 Compound voters per proposal. On average, proposals required 2.84 voters (std. of 0.97) to reach at least 50% of their total votes. The median was 3 voters, with a range of 1 to 5 votes. This indicates a concentrated amount of voting power. The subtitles indicate the proposal ID and outcome ("E" for executed, "D" for defeated, and "C" for cancelled).

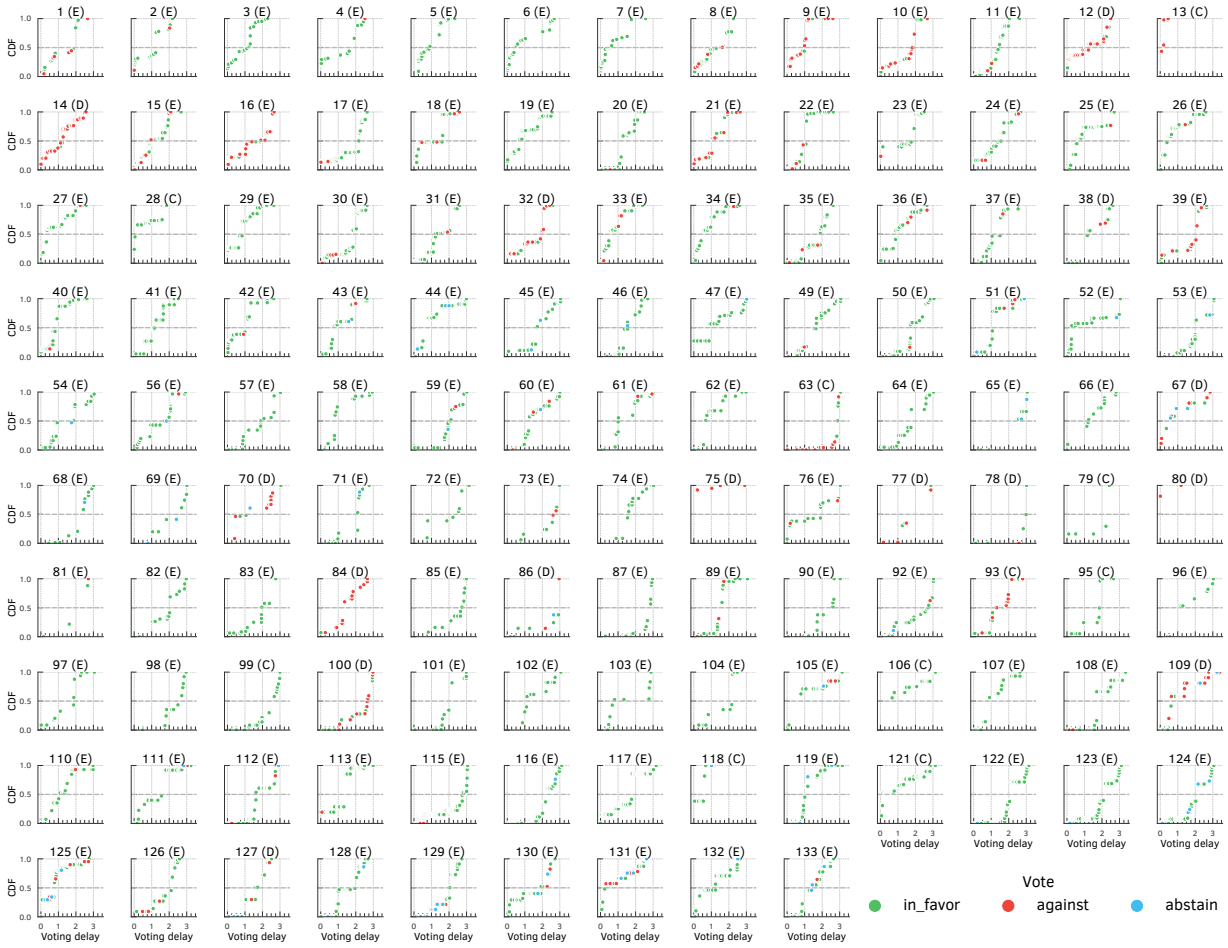

Figure C.4: Voting delays for all votes cast per proposal in chronological order of vote. On average, voters took 1.4 days (with a standard deviation of 0.95 and a median of 1.34 days) to cast their votes after the voting period began. The subtitles indicate the proposal ID and outcome ("E" for executed, "D" for defeated, and "C" for cancelled).

Bibliography

0x Protocol (2022). 0x: Powering the decentralized exchange of tokens on Ethereum. https://www.0x.org.

AAVE (2022). AAVE - Open Source Liquidity Pool. https://aave.com.

AAVE (2023). AAVE Economics – Governance. https://docs.aave.com/aavenomics/governance. Accessed on May 25, 2023.

Adams, H., Zinsmeister, N., Salem, M., Keefer, R., and Robinson, D. (2021). Uniswap v3 core.

Amico, J. (2023). On crypto governance. https://a16z.com/2021/02/05/on-crypto-governance. Accessed on February 8, 2023.

Andrew Hinkes and Joe Ciccolo (2021). OFAC's Bitcoin Blacklist Could Change Crypto. https://www.coindesk.com/goodbye-fungibility-ofacs-bitcoin-blacklist-remake-crypto/.

Antonio Pierro, G., Rocha, H., Tonelli, R., and Ducasse, S. (2020). Are the gas prices oracle reliable? a case study using the ethgasstation. In *2020 IEEE International Workshop on Blockchain Oriented Software Engineering (IWBOSE)*, pages 1–8.

Antonopoulos, A. M. (2014). *Mastering Bitcoin: unlocking digital cryptocurrencies*. O'Reilly Media, Inc.

Antonopoulos, A. M. and Wood, G. (2018). *Mastering ethereum: building smart contracts and dapps*. O'Reilly Media, Inc.

AntPool (2022). Prioritize Transaction. https://www.antpool.com/user/prioritiseTransaction.htm.

Arruñada, B. and Garicano, L. (2018). Blockchain: The birth of decentralized governance. *Pompeu Fabra University, Economics and Business Working Paper Series*, 1608.

Asayag, A., Cohen, G., Grayevsky, I., Leshkowitz, M., Rottenstreich, O., Tamari, R., and Yakira, D. (2018). A fair consensus protocol for transaction ordering. In *2018 IEEE 26th International Conference on Network Protocols (ICNP)*.

at Berkeley, B. (2021). Compound proposal 89 - lower proposal threshold to 25k comp. ***https://compound.finance/governance/proposals/89***. Accessed on May 22, 2023.

Atzori, M. (2017). Blockchain technology and decentralized governance: Is the state still necessary? *Journal of Governance and Regulation/Volume*, 6(1).

Babel Finance (2021). Economic Daily: Babel Finance Launches Ethereum Mining Pool.

Bahack, L. (2013). Theoretical Bitcoin Attacks with less than Half of the Computational Power (draft). *CoRR*, abs/1312.7013.

Baird, L. (2016). The swirlds hashgraph consensus algorithm: Fair, fast, byzantine fault tolerance. Technical Report SWIRLDS-TR-2016, Swirlds, Inc.

Balancer.fi (2023). Governance – Balancer. ***https://docs.balancer.fi/ecosystem/governance***. Accessed on May 25, 2023.

Balancerl (2022). Balancer AMM DeFi protocol. ***https://balancer.fi***.

Bancor (2022). Bancor: Grow your ETH. ***https://bancor.network***.

Basu, S., Easley, D., O'Hara, M., and Sirer, E. G. (2019). Towards a Functional Fee Market for Cryptocurrencies. *CoRR*, abs/1901.06830.

Beck, R., Müller-Bloch, C., and King, J. L. (2018). Governance in the blockchain economy: A framework and research agenda. *Journal of the Association for Information Systems*, 19(10).

Behrens, J. (2017). The origins of liquid democracy. *The Liquid Democracy Journal*, 5(2).

Binance (2023). ***https://www.binance.com***. Accessed on May 25, 2023.

Bitcoin Abuse (2020). Bitcoin Abuse. ***https://www.bitcoinabuse.com/reports***. Accessed on March 19, 2020.

Bitcoin Wiki (2023a). getblocktemplate. ***https://en.bitcoin.it/wiki/Getblocktemplate***. Accessed on May 26, 2023.

Bitcoin Wiki (2023b). Script. *https://en.bitcoin.it/wiki/Script*. Accessed on August 3, 2023.

BitcoinAbuse (2021). Recently Reported Addresses. *https://www.bitcoinabuse.com/reports*.

bitcoin.org (2023). Bitcoin Core. *https://bitcoin.org/en/bitcoin-core*. Accessed on May 26, 2023.

Blockchain.com (2021). Ethereum Explorer. *https://www.blockchain.com/explorer?view=eth*.

Blockchair (2023). Ethereum Explorer. *https://blockchair.com/ethereum*. Accessed on May 25, 2023.

BloXroute Labs (2022). BloXroute Labs. *https://bloxroute.com*.

Blum, C. and Zuber, C. I. (2016). Liquid democracy: Potentials, problems, and perspectives. *Journal of political philosophy*, 24(2).

Braiins (2021a). Stratum mining protocol V1. *https://braiins.com/stratum-v1*. Accessed on May 26, 2023.

Braiins (2021b). Stratum mining protocol V2. *https://braiins.com/stratum-v2*. Accessed on May 26, 2023.

Breidenbach, L., Cachin, C., Chan, B., Coventry, A., Ellis, S., Juels, A., Koushanfar, F., Miller, A., Magauran, B., Moroz, D., et al. (2021). Chainlink 2.0: Next steps in the evolution of decentralized oracle networks. *Chainlink Labs*.

BTC.com (2022). BTC.com Transaction Accelerator. *https://pushtx.btc.com*.

Buterin, V., Conner, E., Dudley, R., Slipper, M., Norden, I., and Bakhta, A. (2019a). "EIP-1559: Fee market change for ETH 1.0 chain," Ethereum Improvement Proposals.

Buterin, V., Hitzig, Z., and Weyl, E. G. (2019b). A flexible design for funding public goods. *Management Science*, 65(11).

Capital, P. (2021). Compound proposal 60 - address whitelist for submitting proposals. *https://compound.finance/governance/proposals/60*. Accessed on May 22, 2023.

Carlsten, M., Kalodner, H., Weinberg, S. M., and Narayanan, A. (2016). On the instability of bitcoin without the block reward. In *Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security*, CCS '16.

Carroll, L. (1884). Principles of parliamentary representation.

Chainlink (2022). Decentralized Data Feeds. ***https://data.chain.link***.

Chainlink Foundation (2023). What Is Layer 2? ***https://chain.link/education-hub/what-is-layer-2***. Accessed on August 22, 2023.

Chapman, J. W. (1971). Rawls's theory of justice.

Chen, X., Papadimitriou, C., and Roughgarden, T. (2019). An axiomatic approach to block rewards. In *Proceedings of the 1st ACM Conference on Advances in Financial Technologies*, AFT '19.

Chen, Y., Richter, J. I., and Patel, P. C. (2021). Decentralized governance of digital platforms. *Journal of Management*, 47(5).

City, S. (2023). Swarm city. ***https://swarm.city***. Accessed on February 2, 2023.

Coin Dance (2021). Bitcoin Nodes Summary. ***https://coin.dance/nodes***.

Coinbase (2021). What are miner fees and does Coinbase pay them? ***https://help.coinbase.com/en/coinbase/trading-and-funding/pricing-and-fees/what-are-miner-fees-and-does-coinbase-pay-them.html***.

CoinGecko (2023). Compound Tokenomics. ***https://www.coingecko.com/en/coins/compound/tokenomics***. Accessed on February 2, 2023.

CoinMarketCap (2023). Cryptocurrency Prices by Market Cap. ***https://coinmarketcap.com***. Accessed on May 23, 2023.

CoinStaker (2018). Bitcoin CPFP Experience—Bitcoin Child Pays for Parent. ***https://www.coinstaker.com/bitcoin-cpfp/***.

Compound (2022). Compound. ***https://compound.finance***.

Compound Labs, Inc. (2022a). Compound Governance. ***https://docs.compound.finance/governance***. Accessed on Dec 10, 2022.

Compound Labs, Inc. (2022b). Compound Leaderboard. ***https://compound.finance/governance/leaderboard***. Accessed on Dec 2, 2022.

Compound Labs, Inc. (2022c). User Distribution. *https://compound.finance/governance/comp*. Accessed on Dec 10, 2022.

Convex (2023a). Convex Finance Proposals. *https://vote.convexfinance.com/*. Accessed on May 25, 2023.

Convex (2023b). Voting and Gauge Weights. *https://docs.convexfinance.com/convexfinance/general-information/why-convex/voting-and-gauge-weights*. Accessed on May 25, 2023.

Curve (2022). Curve.fi. *https://curve.fi*.

Curve (2023). Curve.fi Governance. *https://gov.curve.fi*. Accessed on April 2, 2023.

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

Daian, P., Kell, T., Miers, I., and Juels, A. (2018). On-Chain Vote Buying and the Rise of Dark DAOs. *https://hackingdistributed.com/2018/07/02/on-chain-vote-buying*. Accessed on December 15, 2022.

David Siegel (2013). Understanding The DAO Attack. *https://www.coindesk.com/learn/understanding-the-dao-attack*. Accessed on February 16, 2023.

Easley, D., O'Hara, M., and Basu, S. (2017). From Mining to Markets: The Evolution of Bitcoin Transaction Fees. *SSRN*.

Easley, D., O'Hara, M., and Basu, S. (2019). From mining to markets: The evolution of bitcoin transaction fees. *Journal of Financial Economics*.

Eden Network (2022). Eden Network. *https://www.edennetwork.io/*.

EigenPhi (2022). EigenPhi Crypto & DeFi Analytics. *https://eigenphi.io*.

Ekblaw, A. and Azaria, A. (2017). MedRec: Medical Data Management on the Blockchain. *Viral Communications*.

Eskandari, S., Moosavi, S., and Clark, J. (2020). Sok: Transparent dishonesty: Front-running attacks on blockchain. In Bracciali, A., Clark, J., Pintore, F., Rønne, P. B., and Sala, M., editors, *Financial Cryptography and Data Security*. Springer International Publishing.

Ethereum Foundation (2022a). Proof-of-Stake (PoS). ***https://ethereum.org/en/developers/docs/consensus-mechanisms/pos***.

Ethereum Foundation (2022b). The Merge. ***https://ethereum.org/en/upgrades/merge***.

Ethereum Foundation (2023a). ERC-20 token standard. ***https://ethereum.org/en/developers/docs/standards/tokens/erc-20***. Accessed on April 10, 2023.

Ethereum Foundation (2023b). ERC-721 Non-Fungible token standard. ***https://ethereum.org/en/developers/docs/standards/tokens/erc-721***.

Ethereum Foundation (2023c). Non-fungible tokens (NFT). ***https://ethereum.org/en/nft***. Accessed on May 25, 2023.

Ethereum Foundation (2023d). Optimistic Rollups. ***https://ethereum.org/en/developers/docs/scaling/optimistic-rollups***. Accessed on August 23, 2023.

Ethermine (2022). Ethermine MEV-Relay. ***https://ethermine.org/mev-relay***.

Etherscan (2023a). Etherscan apis documentation – contracts. ***https://docs.etherscan.io/api-endpoints/contracts***. Accessed on May 25, 2023.

Etherscan (2023b). Etherscan (ETH) Blockchain Explorer. ***https://etherscan.io***. Accessed on May 25, 2023.

Eyal, I. and Sirer, E. G. (2018). Majority is not enough: Bitcoin mining is vulnerable. *Commun. ACM*.

F2Pool (2022). Pushtx. ***https://www.f2pool.com/pushtx***.

Feichtinger, R., Fritsch, R., Vonlanthen, Y., and Wattenhofer, R. (2023). The hidden shortcomings of (d)aos – an empirical study of on-chain governance.

Fiat, A., Karlin, A., Koutsoupias, E., and Papadimitriou, C. (2019). Energy equilibria in proof-of-work mining. In *Proceedings of the 2019 ACM Conference on Economics and Computation*, EC '19.

Fisher, R. A. (1992). *Statistical Methods for Research Workers*. Springer New York.

Flashbots (2022a). Flashbots Blocks API. ***https://blocks.flashbots.net***.

Flashbots (2022b). Flashbots Docs. ***https://docs.flashbots.net***.

Flashbots (2022c). How much hashrate is currently on Flashbots Auction? ***https://docs.flashbots.net/flashbots-auction/searchers/faq#how-much-hashrate-is-currently-on-flashbots-auction***.

Francisco Rodrigues (2022). Maintaining decentralization: Are custody services a threat to DeFi protocols? ***https://cointelegraph.com/news/maintaining-decentralization-are-custody-services-a-threat-to-defi-protocols***. Accessed on April 2, 2023.

Fritsch, R., Müller, M., and Wattenhofer, R. (2022). Analyzing voting power in decentralized governance: Who controls daos?

Gauntlet (2021). Compound proposal 40 - compound grants program. ***https://compound.finance/governance/proposals/40***. Accessed on May 25, 2023.

Gencer, A. E., Basu, S., Eyal, I., van Renesse, R., and Gün Sirer, E. (2018). Decentralization in Bitcoin and Ethereum Networks. In *Financial Cryptography and Data Security 2018*.

Gervais, A., Karame, G. O., Wüst, K., Glykantzis, V., Ritzdorf, H., and Capkun, S. (2016). On the security and performance of proof of work blockchains. In *Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security*, CCS '16.

Goldreich, O. and Oren, Y. (1994). Definitions and properties of zero-knowledge proof systems. *Journal of Cryptology*, 7(1):1–32.

Goldwasser, S., Micali, S., and Rackoff, C. (2019). The knowledge complexity of interactive proof-systems. In *Providing sound foundations for cryptography: On the work of shafi goldwasser and silvio micali*, pages 203–225.

Goren, G. and Spiegelman, A. (2019). Mind the mining. In *Proceedings of the 2019 ACM Conference on Economics and Computation*, EC '19.

Guidi, B., Michienzi, A., and Ricci, L. (2021). A graph-based socioeconomic analysis of steemit. *IEEE Transactions on Computational Social Systems*, 8(2).

Haig, S. (2022). Balancer ends long governance battle with whale. ***https://thedefiant.io/balancer-ve-tokenomics-whale***. Accessed on April 28, 2023.

Hassan, S. and De Filippi, P. (2021). Decentralized autonomous organization. *Internet Policy Review*, 10(2).

Helen Partz (2021). Second-largest Ethereum mining pool to suspend all operations. Cointelegraph.

Hobbes, T. (1651). Leviathan. *Project Gutenberg*.

Huberman, G., Leshno, J. D., and Moallemi, C. (2021). Monopoly without a Monopolist: An Economic Analysis of the Bitcoin Payment System. *The Review of Economic Studies*.

Huillet, M. (2020). Report: Hamas, Iran-Tied Militants Intensify Bitcoin Fundraising Action. ***https://cointelegraph.com/news/report-hamas-iran-tied-militants-intensify-bitcoin-fundraising-action***. Accessed on August 2, 2023.

Jeff Kauflin and Emily Mason (2022). How Did Sam Bankman-Fried's Alameda Research Lose So Much Money? ***https://www.forbes.com/sites/jeffkauflin/2022/11/19/how-did-sam-bankman-frieds-alameda-research-lose-so-much-money***. Accessed on December 2, 2022.

Jota Missias (2021). Jota Missias. ***https://pt.wikipedia.org/wiki/Jota_Missias***.

Judmayer, A., Zamyatin, A., Stifter, N., Voyiatzis, A. G., and Weippl, E. (2017). Merged mining: Curse or cure? In *Data Privacy Management, Cryptocurrencies and Blockchain Technology*.

Karame, G. (2016). On the security and scalability of bitcoin's blockchain. In *Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security*, CCS '16.

Kelkar, M., Zhang, F., Goldfeder, S., and Juels, A. (2020). Order-fairness for byzantine consensus. In Micciancio, D. and Ristenpart, T., editors, *Advances in Cryptology – CRYPTO 2020*, pages 451–480, Cham. Springer International Publishing.

Kharif, O. (2017). Cryptokitties mania overwhelms ethereum network's processing. ***https://www.bloomberg.com/news/articles/2017-12-04/cryptokitties-quickly-becomes-most-widely-used-ethereum-app***.

Kiayias, A., Koutsoupias, E., Kyropoulou, M., and Tselekounis, Y. (2016). Blockchain mining games. In *Proceedings of the 2016 ACM Conference on Economics and Computation*, EC '16.

Kiayias, A. and Lazos, P. (2023). Sok: Blockchain governance.

Kiffer, L., Levin, D., and Mislove, A. (2017). Stick a fork in it: Analyzing the ethereum network partition. In *Proceedings of the 16th ACM Workshop on Hot Topics in Networks*, pages 94–100.

Kursawe, K. (2020). Wendy, the good little fairness widget: Achieving order fairness for blockchains. In *Proceedings of the 2nd ACM Conference on Advances in Financial Technologies*, AFT '20.

kybx86 (2020). Compensation proposal: Distribute comp to affected users in the dai liquidations. ***https://www.comp.xyz/t/compensation-proposal-distribute-comp-to-affected-users-in-the-dai-liquidations***. Accessed on February 9, 2023.

Labs, C. (2020). Compound proposal 7 - distribute comp to users. ***https://compound.finance/governance/proposals/7***. Accessed on May 22, 2023.

Labs, C. (2021). Compound proposal 42 - migration to governor bravo. ***https://compound.finance/governance/proposals/42***. Accessed on May 22, 2023.

Labs, P. (2023a). Ipfs powers the distributed web. ***https://ipfs.tech***. Accessed on February 2, 2023.

Labs, S. (2023b). Snapshot. ***https://snapshot.org***. Accessed on February 2, 2023.

Lalley, S. P. and Weyl, E. G. (2018). Quadratic voting: How mechanism design can radicalize democracy. *AEA Papers and Proceedings*, 108.

Lavi, R., Sattath, O., and Zohar, A. (2019). Redesigning bitcoin's fee market. In *The World Wide Web Conference*, WWW '19.

Lee Mathews (2017). How WannaCry Went From A Windows Bug To An International Incident. ***https://www.forbes.com/sites/leemathews/2017/05/16/wannacry-ransomware-ms17-010***.

Leshner, R. and Hayes, G. (2019). Compound: The money market protocol.

Lev-Ari, K., Spiegelman, A., Keidar, I., and Malkhi, D. (2020). Fairledger: A fair blockchain protocol for financial institutions. In *23rd International Conference on Principles of Distributed Systems (OPODIS 2019)*. Schloss Dagstuhl-Leibniz-Zentrum für Informatik.

Li, J., Yuan, Y., Wang, S., and Wang, F.-Y. (2018). Transaction queuing game in bitcoin blockchain. In *2018 IEEE Intelligent Vehicles Symposium (IV)*.

Liu, C., Liu, H., Cao, Z., Chen, Z., Chen, B., and Roscoe, B. (2018). Reguard: Finding reentrancy bugs in smart contracts. In *Proceedings of the 40th International Conference on Software Engineering: Companion Proceeedings*, ICSE '18.

Liu, F., Wang, X., Li, Z., Xu, J., and Gao, Y. (2020). Effective gasprice prediction for carrying out economical ethereum transaction. In *2019 6th International Conference on Dependable Systems and Their Applications (DSA)*, pages 329–334.

Liu, Z., Li, Y., Min, Q., and Chang, M. (2022). User incentive mechanism in blockchain-based online community: An empirical study of steemit. *Information & Management*, 59(7).

Loewen, T. (2021a). Compound proposal 62 - split comp rewards distribution and bug fixes. **https://compound.finance/governance/proposals/62**. Accessed on May 16, 2023.

Loewen, T. (2021b). Compound proposal 64 - fix comp accrual bug. **https://compound.finance/governance/proposals/64**. Accessed on May 16, 2023.

Lombrozo, E., Lau, J., and Wuille, P. (2015). BIP-141: Segregated Witness (Consensus layer). **https://github.com/bitcoin/bips/blob/master/bip-0141.mediawiki**. Accessed on May 26, 2023.

Luu, L., Velner, Y., Teutsch, J., and Saxena, P. (2017). Smartpool: Practical decentralized pooled mining. In *26th USENIX Security Symposium (USENIX Security 17)*.

MakerDAO (2023). Governance Module – Maker Protocol Technical Docs. **https://docs.makerdao.com/smart-contract-modules/governance-module**. Accessed on April 2, 2023.

Mars, R., Abid, A., Cheikhrouhou, S., and Kallel, S. (2021). A machine learning approach for gas price prediction in ethereum blockchain. In *2021 IEEE 45th Annual Computers, Software, and Applications Conference (COMPSAC)*, pages 156–165.

Martin Ruubel (2018). World's First Blockchain Platform for Marine Insurance Now in Commercial Use. **https://guardtime.com/blog/world-s-first-blockchain-platform-for-marine-insurance-now-in-commercial-use**.

Matter Labs (2023). zkSync Era basics. **https://era.zksync.io/docs/reference/concepts/zkSync.html**. Accessed on August 22, 2023.

McCorry, P., Shahandashti, S. F., and Hao, F. (2017). A smart contract for boardroom voting with maximum voter privacy. In *Financial Cryptography and Data Security*, FC '17.

Messari (2023). Messati: Crypto research, data, and tools. ***https://messari.io***. Accessed on May 16, 2023.

Messias, J. (2023a). Data sets and scripts used to analyze governance protocols in the Ethereum blockchains. ***https://github.com/johnnatan-messias/blockchain-governance***.

Messias, J. (2023b). Data sets and scripts used to analyze the contention and prioritization transparency in both Bitcoin and Ethereum blockchains. ***https://github.com/johnnatan-messias/blockchain-transaction-ordering***.

Messias, J., Alzayat, M., Chandrasekaran, B., and Gummadi, K. P. (2020). On blockchain commit times: An analysis of how miners choose bitcoin transactions. In *KDD Workshop on Smart Data for Blockchain and Distributed Ledger*, SDBD '20.

Messias, J., Alzayat, M., Chandrasekaran, B., Gummadi, K. P., Loiseau, P., and Mislove, A. (2021). Selfish & opaque transaction ordering in the bitcoin blockchain: The case for chain neutrality. In *Proceedings of the 21st ACM Internet Measurement Conference*, IMC '21, page 320–335, New York, NY, USA. Association for Computing Machinery.

Messias, J., Pahari, V., Chandrasekaran, B., Gummadi, K. P., and Loiseau, P. (2023a). Dissecting Bitcoin and Ethereum Transactions: On the Lack of Transaction Contention and Prioritization Transparency in Blockchains. In *Proceedings of the Financial Cryptography and Data Security (FC'23)*.

Messias, J., Pahari, V., Chandrasekaran, B., Gummadi, K. P., and Loiseau, P. (2023b). Understanding blockchain governance: Analyzing decentralized voting to amend defi smart contracts.

Mike Dalton (2022). Build Finance DAO Suffers Governance Takeover Attack. ***https://cryptobriefing.com/build-finance-dao-suffers-governance-takeover-attack***. Accessed on February 2, 2023.

Mosteller, F. and Fisher, R. A. (1948). Questions and answers. *The American Statistician*.

Nakamoto, S. (2008). Bitcoin: A Peer-to-Peer Electronic Cash System.

Nick Martitsch (2021). Compound Treasury Updates, COMP Bug Fix, Dynamic Risk Parameters. ***https://compound.substack.com/p/compound-treasury-updates-comp-bug***. Accessed on April 10, 2023.

Nikhilesh De (2021). US Treasury Blacklists Bitcoin, Litecoin Addresses of Chinese 'Drug Kingpins'. ***https://www.coindesk.com/markets/2019/08/21/us-treasury-blacklists-bitcoin-litecoin-addresses-of-chinese-drug-kingpins/***.

Noda, S., Okumura, K., and Hashimoto, Y. (2020). An economic analysis of difficulty adjustment algorithms in proof-of-work blockchain systems. In *Proceedings of the 21st ACM Conference on Economics and Computation*, EC '20.

Offchain Labs (2023). A gentle introduction to Arbitrum. ***https://docs.arbitrum.io/intro***. Accessed on August 22, 2023.

Omni Layer (2020). Omni Protocol Specification. ***https://github.com/OmniLayer/spec/blob/master/OmniSpecification.adoc***. Accessed on August 3, 2023.

Omni Layer (2023). Omni Layer: An open-source, fully-decentralized asset platform on the Bitcoin Blockchain. ***https://www.omnilayer.org***. Accessed on August 1, 2023.

OpenZeppelin (2023). Security audits - compound. ***https://blog.openzeppelin.com/?s=compound***. Accessed on May 22, 2023.

Optimism Foundation (2023). Optimism. ***https://www.optimism.io***. Accessed on August 22, 2023.

Orda, A. and Rottenstreich, O. (2019). Enforcing fairness in blockchain transaction ordering. In *2019 IEEE International Conference on Blockchain and Cryptocurrency (ICBC)*.

Pass, R., Seeman, L., and Shelat, A. (2017). Analysis of the blockchain protocol in asynchronous networks. In *Advances in Cryptology – EUROCRYPT 2017*.

Pass, R. and Shi, E. (2017). Fruitchains: A fair blockchain. In *Proceedings of the ACM Symposium on Principles of Distributed Computing*, PODC '17.

Paul R. La Monica (2019). Bitcoin's march to $10,000 propelled by Facebook and the Fed. ***https://edition.cnn.com/2019/06/21/investing/bitcoin-price-increase/***.

Perez, D., Werner, S. M., Xu, J., and Livshits, B. (2021). Liquidations: Defi on a knife-edge. In *Financial Cryptography and Data Security*, FC '21.

Philipp Schmidt (2015). Certificates, Reputation, and the Blockchain. ***https://medium.com/mit-media-lab/certificates-reputation-and-the-blockchain-aee03622426f***.

Pilkington, M. (2016). Blockchain Technology: Principles and Applications. In *Research handbook on digital transformations*. Available at SSRN: https://ssrn.com/abstract=2662660.

Poolin (2022). Transaction Accelerator. ***https://pushtx.com***.

Provenance (2015). Blockchain: the solution for transparency in product supply chains. ***https://www.provenance.org/whitepaper***.

Qin, K., Zhou, L., and Gervais, A. (2022). Quantifying blockchain extractable value: How dark is the forest? In *2022 IEEE Symposium on Security and Privacy (SP)*.

Qin, K., Zhou, L., Livshits, B., and Gervais, A. (2021). Attacking the DeFi Ecosystem with Flash Loans for Fun and Profit. In *Financial Cryptography and Data Security*, FC '21.

Reijers, W., O'Brolcháin, F., and Haynes, P. (2016). Governance in blockchain technologies & social contract theories. *Ledger*, 1.

Research, A. (2020a). Compound proposal 13 - set wbtc collateral factor to 65%. ***https://compound.finance/governance/proposals/13***. Accessed on May 22, 2023.

Research, A. (2020b). Compound proposal 14 - set wbtc collateral factor to 65%. ***https://compound.finance/governance/proposals/14***. Accessed on May 22, 2023.

Research, A. (2020c). Compound proposal 16 - set wbtc collateral factor to 40%. ***https://compound.finance/governance/proposals/16***. Accessed on May 22, 2023.

Rikken, O., Janssen, M., and Kwee, Z. (2019). Governance challenges of blockchain and decentralized autonomous organizations. *Information Polity*, 24(4).

Robert Hackett (2017). Walmart and 9 Food Giants Team Up on IBM Blockchain Plans. ***http://fortune.com/2017/08/22/walmart-blockchain-ibm-food-nestle-unilever-tyson-dole***.

Romiti, M., Judmayer, A., Zamyatin, A., and Haslhofer, B. (2019). A deep dive into bitcoin mining pools: An empirical analysis of mining shares.

Roughgarden, T. (2021). Transaction fee mechanism design for the ethereum blockchain: An economic analysis of eip-1559. In *Proceedings of the 2021 ACM Conference on Economics and Computation*, EC '21.

Rousseau, J.-J. (1920). *The social contract: & discourses*. Number 660. JM Dent & Sons.

Sam Kessler (2022). Binance Becomes Second-Largest Voting Entity on Uniswap DAO. ***https://www.coindesk.com/tech/2022/10/19/binance-becomes-second-largest-voting-entity-on-uniswap-dao***. Accessed on April 2, 2023.

Sasson, E. B., Chiesa, A., Garman, C., Green, M., Miers, I., Tromer, E., and Virza, M. (2014). Zerocash: Decentralized anonymous payments from bitcoin. In *2014 IEEE symposium on security and privacy*. IEEE.

Scam Alert (2020). Scam Alert - Cryptocurrency Crime Fighters. ***https://scam-alert.io***. Accessed on March 19, 2020.

Scikit Learn (2023). Cosine Similarity. ***https://scikit-learn.org/stable/modules/generated/sklearn.metrics.pairwise.cosine_similarity.html***. Accessed on April 10, 2023.

Sharma, T., Kwon, Y., Pongmala, K., Wang, H., Miller, A., Song, D., and Wang, Y. (2023). Unpacking how decentralized autonomous organizations (daos) work in practice.

Shaurya Malwa (2022). Binance Denies Allegations It Intends to Use Users' Uniswap Tokens for Voting. ***https://www.coindesk.com/tech/2022/10/20/binance-denies-allegations-that-it-intends-to-use-users-uniswap-tokens-for-voting***.

Sheera Frenkel and Nathaniel Popper and Kate Conger and David E. Sanger (2020). A Brazen Online Attack Targets V.I.P. Twitter Users in a Bitcoin Scam. ***https://www.nytimes.com/2020/07/15/technology/twitter-hack-bill-gates-elon-musk.html***.

Sheera Frenkel, Mark Scott and Paul Mozur (2017). Mystery of Motive for a Ransomware Attack: Money, Mayhem or a Message? ***https://www.nytimes.com/2017/06/28/business/ramsonware-hackers-cybersecurity-petya-impact.html***.

Siddiqui, S., Vanahalli, G., and Gujar, S. (2020). Bitcoinf: Achieving fairness for bitcoin in transaction fee only model. In *Proceedings of the 19th International Conference on Autonomous Agents and MultiAgent Systems*, AAMAS '20.

Solidity Team (2023). Solidity programming language. ***https://soliditylang.org***. Accessed on January 18, 2023.

Sompolinsky, Y. and Zohar, A. (2015). Secure high-rate transaction processing in bitcoin. In *Financial Cryptography and Data Security*, FC '15.

SparkPool (2021). Taichi Network. ***https://taichi.network***.

Strehle, E. and Ante, L. (2020). Exclusive mining of blockchain transactions. In *In Scientific Reports 2020-Conference proceedings of the Scientific Track of the Blockchain Autumn School 2020*.

Sukernik (2021). Auditing Compound Protocol. ***https://www.comp.xyz/t/auditing-compound-protocol/2543***. Accessed on April 10, 2023.

SushiSwap (2022). Decentralized Exchange Made For Everybody. ***https://www.sushi.com***.

Sybil (2023a). Introducing Sybil. ***https://blog.uniswap.org/sybil***. Accessed on May 19, 2023.

Sybil (2023b). Sybil – Top delegated addresses. ***https://sybil.org***. Accessed on February 2, 2023.

team, T. T. (2022). Compound proposal 84 - trueusd market upgrades. ***https://compound.finance/governance/proposals/84***. Accessed on February 2, 2023.

Tether (2023). Tether USDT token. ***https://tether.to/en***. Accessed on August 2, 2023.

Thurman, A. (2022). Tron's justin sun accused of 'governance attack' on defi lender compound. ***https://www.coindesk.com/tech/2022/02/04/trons-justin-sun-accused-of-governance-attack-on-defi-lender-compound***. Accessed on February 2, 2023.

Torres, C. F., Camino, R., and State, R. (2021). Frontrunner jones and the raiders of the dark forest: An empirical study of frontrunning on the ethereum blockchain. In *30th USENIX Security Symposium*.

TrueUSD (2023). Trueusd. ***https://www.trueusd.com***. Accessed on February 2, 2023.

Tsabary, I. and Eyal, I. (2018). The gap game. In *Proceedings of the 2018 ACM SIGSAC Conference on Computer and Communications Security*, CCS '18.

Turksonmez, K., Furtak, M., Wittie, M. P., and Millman, D. L. (2021). Two ways gas price oracles miss the mark. In *2021 IEEE International Conference on Omni-Layer Intelligent Systems (COINS)*, pages 1–7.

Uniswap (2022). Uniswap Protocol. ***https://uniswap.org***.

Uniswap (2023). Uniswap Decentralized Trading Protocol. ***https://uniswap.org***. Accessed on May 23, 2023.

Uniswap Labs (2023). Governance – Uniswap Protocol. ***https://uniswap.org/governance***. Accessed on April 2, 2023.

Van Saberhagen, N. (2013). Cryptonote v2.0.

Vasek, M., Thornton, M., and Moore, T. (2014). Empirical analysis of denial-of-service attacks in the bitcoin ecosystem. In *Financial Cryptography and Data Security*, FC '14.

ViaBTC (2022). Transaction Accelerator. ***https://www.viabtc.com/tools/txaccelerator***.

web3.py team, T. (2022). Web3.py documentation. ***https://web3py.readthedocs.io/en/v5***. Accessed on December 12, 2022.

Weintraub, B., Torres, C. F., Nita-Rotaru, C., and State, R. (2022). A Flash(bot) in the Pan: Measuring Maximal Extractable Value in Private Pools. In *Proceedings of the ACM Internet Measurement Conference (IMC'22)*.

Whale Alert (2021). Scam Alert - Cryptocurrency Crime Fighters. ***https://scam-alert.io***.

Wood, G. (2016). Polkadot: Vision for a heterogeneous multi-chain framework.

Wood, G. et al. (2014). Ethereum: A secure decentralised generalised transaction ledger. *Ethereum project yellow paper*.

Xia, P., Zhang, L., and Li, F. (2015). Learning similarity with cosine similarity ensemble. *Information sciences*, 307.

Yahoo Finance (2023a). Ethereum USD (BTC-USD), Price, Value, News, and History. ***https://finance.yahoo.com/quote/BTC-USD/history?p=BTC-USD***. Accessed on August 3, 2023.

Yahoo Finance (2023b). Ethereum USD (ETH-USD), Price, Value, News, and History. ***https://finance.yahoo.com/quote/ETH-USD/history?p=ETH-USD***. Accessed on January 17, 2023.

Zack Voell and William Foxley (2020a). ***https://www.coindesk.com/markets/2020/09/04/alameda-research-claimed-nearly-70-of-wrapped-bitcoin-minted-in-august***. Accessed on April 10, 2023.

Zack Voell and William Foxley (2020b). Alameda Research Claimed Nearly 70% of Wrapped Bitcoin Minted in August. ***https://www.coindesk.com/markets/2020/09/04/alameda-research-claimed-nearly-70-of-wrapped-bitcoin-minted-in-august***.

Zhang, R. and Preneel, B. (2019). Lay down the common metrics: Evaluating proof-of-work consensus protocols' security. In *2019 IEEE Symposium on Security and Privacy (SP)*.

Zhou, L., Qin, K., Torres, C. F., Le, D. V., and Gervais, A. (2021). High-frequency trading on decentralized on-chain exchanges. In *2021 IEEE Symposium on Security and Privacy (SP)*, pages 428–445.

Zhou, L., Xiong, X., Ernstberger, J., Chaliasos, S., Wang, Z., Wang, Y., Qin, K., Wattenhofer, R., Song, D., and Gervais, A. (2023). Sok: Decentralized finance (defi) attacks. In *2023 IEEE Symposium on Security and Privacy (SP)*.

Zwitter, A. and Hazenberg, J. (2020). Decentralized network governance: blockchain technology and the future of regulation. *Frontiers in Blockchain*, 3.