The Great Fee Migration: An Evaluation of Blockchain Protocols Without Block Rewards

Edvinas Kruminis, BSc (Hons)

School of Computing and Communications

Lancaster University

A thesis submitted for the degree of

Doctor of Philosophy

December, 2025

The Great Fee Migration: An Evaluation of Blockchain Protocols Without Block Rewards

Edvinas Kruminis, BSc (Hons).

School of Computing and Communications, Lancaster University

A thesis submitted for the degree of *Doctor of Philosophy*. December, 2025.

Abstract

Many established blockchain protocols are programmed to transition from a dual-incentive model of block rewards and transaction fees to one sustained exclusively by fees. This paradigm shift, which we term the "Great Fee Migration", introduces high revenue volatility for miners, which undermines the network's security budget and creates economic incentives for destabilising attacks such as fee sniping. This thesis addresses this critical challenge by evaluating blockchain performance under fee-only conditions, with a particular focus on transaction fee mechanism design. We propose and evaluate a novel mechanism, the Reserve Pool TFM, which recycles surplus fees into a network-shared fund to subsidise miner payouts during periods of low network activity, thereby smoothing revenue over time. To empirically validate this design, we develop the Zero-Block Reward Auction simulator, the first framework designed for comparative TFM analysis under fee-only conditions. Using a dataset of over 12 million real-world Bitcoin transactions, our simulations demonstrate that the Reserve Pool TFM is able to reduce miner payout variance by approximately 75% compared to traditional mechanisms, and achieves a 97% block profitability rate, substantially strengthening overall network security. Our thesis is further supported by a game-theoretic study revealing that miner revenue depends on mempool congestion rather than artificial block size

constraints, and a case study of a privacy-preserving digital advertising system, BB-FLoC, which confirms the viability of the Reserve Pool TFM in non-financial blockchain applications. Our findings offer a concrete path towards ensuring the long-term economic security and sustainability of future fee-only blockchain systems.

Declaration

I declare that the work presented in this thesis is, to the best of my knowledge and belief, original and my own work. The material has not been submitted, either in whole or in part, for a degree at this, or any other university. This thesis does not exceed the maximum permitted word length of 80,000 words including appendices and footnotes, but excluding the bibliography. A rough estimate of the word count is: 27828

Edvinas Kruminis

Publications

This thesis is informed by a series of publications produced during its development. The following works, including both completed and in-progress manuscripts, have directly contributed to the research presented:

Edvinas Kruminis and Keivan Navaie (2022). "Game-Theoretic Analysis of an Exclusively Transaction-Fee Reward Blockchain System". In: *IEEE Access* 10, pp. 5002–5011. DOI: 10.1109/ACCESS.2022.3140921

Edvinas Kruminis, Keivan Navaie, and Onur Ascigil (Aug. 2025). "BB-FLoC: A Blockchain-Based Targeted Advertisement Scheme with k-Anonymity". In: *Distributed Ledger Technologies: Research and Practice* 4.3. DOI: 10.1145/3672404. URL: https://doi.org/10.1145/3672404

Edvinas Kruminis, Onur Ascigil, and Keivan Navaie (2025). "A Transaction Fee Mechanism for Fee-Only Blockchains: The Reserve Pool Approach". Manuscript in preparation

Contents

1 Introduction **1**

1.1 Blockchains and the "Fee Phase" 1

1.2 Research Questions and Contributions 5

1.3 Thesis Outline 8

2 Background and Related Literature **11**

2.1 Basics of Blockchain 12

2.2 Consensus Schemes 14

2.2.1 Proof-of-Work: Consensus Through Computation 14

2.2.2 Proof-of-Stake: Consensus Through Collateral 15

2.2.3 Delegated Proof-of-Stake: Consensus Through Delegation . . . 16

2.2.4 Practical Byzantine Fault Tolerance: Consensus Through Voting 16

2.2.5 Consensus Trade-Offs 17

2.3 Types of Blockchains 19

2.4 Mining Incentives 20

2.5 TFM Basics and Literature 22

2.6 Existing Transaction Fee Mechanism Designs 25

2.6.1 1st-Price Auction 30

2.6.2 2nd-Price Auction 31

2.6.3 EIP-1559 33

2.6.4 Burning 2nd-Price Auction 34

2.7 State-Of-The-Art Landscape 36

3 Phasing Out Block Rewards **38**

3.1 Motivation and Background 38

3.2 Evolution of Miner Revenue 40

3.3 Reserve Pool TFM 46

3.4 Comparative Analysis of Incentive Properties 48

4 Game-Theoretic Analysis of an Exclusively Transaction Fee Reward Blockchain System **52**

4.1 Introduction: A Game-Theoretic View of the Fee Market 53

4.2 Modelling Framework and Assumptions 53

4.3 Simulation Design and Implementation 56

4.4 Core Findings 58

4.5 Broader Implications for Fee-Only Blockchains 60

5 Simulation Setup and Methodology **61**

5.1 Input Datasets 61

5.2 Zero-Block Reward Auction Simulator: Architecture and Design . . . 63

5.3 Simulation Parameters 67

6 Analysis of Results and Implications for Protocol Design **72**

6.1 Experimental Scenarios 73

6.2 Results 75

6.3 Model Analysis 78

6.4 Security Analysis 86

6.5 Transitioning to the Reserve Pool TFM 94

6.6 Recommendations 96

6.6.1 Reserve Pool TFM Deployment Guidelines 96

6.6.2 General Fee-Only System Recommendations 98

7 Case Study — BB-FLoC: Blockchain-Based, K-Anonymous Targeted Advertising **100**

7.1 Application Context 100
7.2 Overview of the BB-FLoC System 101
7.3 Economic Incentive Challenges 105
7.3.1 Scenario 1: EIP-1559 Implementation 106
7.3.2 Scenario 2: Reserve Pool TFM Implementation 107
7.4 Summary and Insights 107

8 Conclusions and Future Work **109**
8.1 Contributions 109
8.2 Future Work 112
8.3 Final Remarks 114

References **115**

List of Tables

2.1 Summary of Consensus Mechanisms 18
2.2 List of sample mempool with 8 transactions 27
2.3 Overview of user payout and miner reward rules in different TFMs . . 28
2.4 Overview of transaction inclusion and confirmation rules in different TFMs 28
2.5 List of TFM notations 30
2.6 1st-Price auction TFM rules 31
2.7 2nd-Price auction TFM rules 32
2.8 EIP-1559 auction TFM rules 33
2.9 Burning 2nd-price auction TFM rules 35

3.1 Reserve pool auction TFM rules 47
3.2 Properties of the main TFM designs. 49

4.1 The impact of different block sizes on fees paid when the mempool ratio is the same 58

5.1 Simulator Parameters 69

6.1 Final Simulation Parameters 74
6.2 Simulation Results by TFM Type (Averages) 75
6.3 Block payouts by TFM type (1 BTC = £21,844.63, June 2023 average). 80
6.4 Simulated Miner Profitability (1 BTC = £21,844.63) 85
6.5 Fee Sniping Viability Threshold by TFM Type 93

List of Figures

2.1 An overview of the rules governing a TFM. 26

3.1 Bitcoin mining income composition by year (2009-2020) 42
3.2 Bitcoin mining income composition by year (2021-2024) 43
3.3 Total fees per block (average per week) and the block reward in Bitcoin (2009-2024) 44
3.4 Average fees per block (per week) and the Bitcoin price (USD) (2009-2024) 45

4.1 Maximum fees users are willing to pay, when linearly increasing the mempool levels 59

5.1 High-level system architecture diagram. 64

6.1 Normalized Averages By TFM Type 79

7.1 The BB-FLoC process for displaying targeted advertisements includes the four steps as illustrated in the picture. 102

Chapter 1

Introduction

Blockchain, ever since its emergence as the foundational technology behind Bitcoin in 2009, has fundamentally transformed our understanding of decentralised systems. At its core, a blockchain is a distributed ledger collectively and independently maintained by its participants. Cryptographic mechanisms secure the transactions, while consensus protocols establish their order and shared history. Its security depends as much on economic incentives as on its cryptographic protocols, which motivate participants to consistently commit their resources and act honestly, making attacks on the network prohibitively costly. Although Bitcoin was the first, and still perhaps the most widely known application, the underlying technology of blockchains has since been adapted into a much broader framework that supports diverse uses, including supply chain management, digital identity systems, and privacy-preserving digital advertising. However, as these systems have matured, it has become unclear whether the very economic models supporting their early growth will be able to guarantee long-term viability.

1.1 Blockchains and the "Fee Phase"

Miners in decentralised blockchains commit their resources - whether in terms of computational power or staked collateral - and expect to get rewards for their

contributions. Their revenue generally comes from two distinct sources: fixed block rewards that are issued by the protocol itself, and any transactional fees offered by users of the system. While block rewards are typically guaranteed with each new block produced, transaction fees tend to be voluntary. That is, users are not obligated to include them unless the total number of proposed transactions in a block cycle exceeds the available block space, thereby creating a fee market for inclusion.

Additionally, many cryptocurrencies have a deliberately limited supply; Bitcoin started out awarding 50 BTC for each new block mined, halving it (approximately) every 4 years, and in the year 2025 stands at just 3.125 BTC. Transaction fees, on the other hand, account for only a small percentage of current mining revenue in Bitcoin (refer to Figures 3.1 and 3.2); between 2014 and 2024, transaction fees constituted an average of only 4.55% of annual miner income, peaking at just 13% in 2017. This fixed block reward will eventually disappear entirely, and although it will take many decades for this to happen (sometime in the year 2140), it will reach insignificant levels much sooner than that. At that point, the incentive structure underpinning the system would be severely weakened, as unpredictable income from fees would leave miners and validators unable to cover their ongoing hardware, energy or collateral costs. This may force them to drop out of the network entirely, or consolidate into larger pools, in turn magnifying the risks of centralisation. Moreover, erratic swings in miner payouts encourage strategic behaviours that aim to maximise an individual's revenue streams, such as withholding or forking blocks in an attempt to capture short-term gains, thereby threatening the network's overall security.

Perhaps the biggest concern here will be a lack of stability; the fluctuation and random arrival of transactions versus a fixed reward offered with each block mined. As an example, in Ethereum, during October 2022, as much as 18% of blocks were underfilled ($\leq 50\%$ of target block size), and 22% of blocks were congested ($\geq 150\%$ of target block size). Other than the packing of the blocks, individual transaction fees themselves can vary from just a few cents, to as much as £2.1 million per

transaction (Etherscan, 2020). If the security of the entire network is dependent on fees alone, it is doubtful whether such fluctuations can lead to a stable system. In fact, such volatility not only enhances the feasibility of existing attacks, but can also introduce some new ones:

- **Mining gap** (Tsabary and Eyal, 2018): A situation where miners decide not to create new blocks due to a lack of profitability/reward. This will become particularly prevalent in fee-only blockchains, as without a fixed block reward, miners will have no rational reason to expend their resources until a sufficient total in fees has been accumulated. On top of no new blocks being mined, this, in fact, incentivises miners to fork the blockchain and carry out various types of (profitable) attacks.

- **Undercutting** (Carlsten et al., 2016, Gong et al., 2022, Tang and A. Zhang, 2023): An attack in which a malicious miner forks the blockchain and leaves a subset of transaction fees unclaimed, inciting others to mine upon their fork. In such a way, miners can monetarily encourage others to assist in a double-spending attack. Undercutting can also manifest itself in a variant called *fee sniping*. If fees are the only reward for miners to collect, high fees, whether due to human error or natural urgency of the transaction, could spur miners to occasionally fork the head of the chain to snatch high-paying fees, causing consensus and confirmation issues, potentially stalling the entire system.

- **Whale transactions** (Liao and Katz, 2017): A form of a double-spending attack, where after making a valid transaction with another party, a miner forks the blockchain to a point before the confirmation of the (large) transaction. Here, the miner uses the recouped transaction amount as transaction fees to incentivise other miners to contribute to the fork, again stalling network consensus.

Variability in fee values, unpredictable influx of transactions, network latency, information propagation, and a plethora of other issues associated with not just

blockchain, but distributed systems in general, make intentional malicious behaviour difficult to identify and sanction. Therefore, as block rewards dwindle, transaction fees will become the primary source of compensating miners, making the design of Transaction Fee Mechanisms (TFMs) critical for network sustainability. A TFM can be simply defined as a set of rules governing how blockchain validators are compensated for block production, and how blockchain users pay for transaction inclusion. The majority of public blockchains utilise a 1st-Price Auction as their TFM, where users bid a fee to have their transaction included, and miners rationally prioritise the highest bidders. However, several research studies (e.g. Basu et al., 2019, Yao, 2018) have highlighted the inefficiency of such a design. At the heart of the problem is the incentive structure, which encourages users to act strategically in order to minimise costs: instead of bidding according to the true value or urgency of their transaction, they aim to bid just enough to secure confirmation. In an attempt to address this, Ethereum has recently implemented the EIP-1559 protocol (Ethereum Foundation, 2021), which requires users to pay a (dynamic) base fee for inclusion into a block, and leaves *tips* - optional priority fees paid directly to the miner - as a voluntary payment medium for transaction ordering. This removes the need for users to speculate on what is a competitive fee to offer in a block cycle – they either pay the current base fee, or wait until it adjusts to a more affordable level for them.

An important area that has yet to be addressed is whether, and to what extent, a fee-only-reward blockchain model affects any potential TFM design; in the absence of fixed block rewards, the security of the network is dependent on the consistent arrival of high fee-paying transactions. Current TFM design research has ignored this issue completely, simply assuming that there is always a high enough payout to incentivise miners to continuously engage in block building. As a result, current attempts to design a so-called *dream* TFM Roughgarden, 2021b in blockchains have instead focused on the following three properties being satisfied:

- **User-Incentive Compatibility *(UIC)*:** Users must bid according to their

true value/urgency of the transaction, rather than shading their bids.

- **Miner-Incentive Compatibility *(MIC)*:** Miners must implement the transaction inclusion mechanism honestly and according to the specifications of the protocol.

- **Off-Chain Agreement Proofness *(OCA)*:** No coalition of miners and users conspiring together should be able to increase their joint utility compared to that of the honest implementation of the protocol.

Presently, no current design has successfully managed to achieve all such properties. The closest to attaining the desiderata has been the EIP-1559 protocol; however, when there is network congestion, i.e., more valid proposed transactions than the block size limit allows, EIP-1559 returns to behaving like a 1st-Price Auction by asking users to differentiate their priority demands through offering higher tips, once again failing to satisfy the UIC notion. Our work proposes a TFM design where instead of burning fees (i.e., permanently removing them from circulation), or paying them directly to the miners, the fees are aggregated into a network-shared reserve pool account. Miners are paid a fixed, rolling-average amount drawn from that reserve rather than the per-block fee total, ensuring consistent, smoothed rewards regardless of a particular individual block's fee volatility.

1.2 Research Questions and Contributions

This thesis aims to address three fundamental research questions that are critical to the viability and future sustainability of fee-only blockchain systems:

- *How does the transition from a hybrid reward model to a fee-only model alter the economic incentives and strategic behaviour of participants, and what are the resulting security implications for the blockchain?*

Moving away from diminishing block rewards to a fee-only model represents an inevitable shift for most established blockchain networks. This research question stresses the need to understand how this transition would affect network security and overall stability by analysing the strategic interactions that drive fee market dynamics, such as mempool congestion.

- *Can a novel transaction fee mechanism design provide a stable and secure alternative to block rewards across a diverse set of fee-only blockchain applications?*

Current transaction fee mechanism research assumes constant presence of block rewards, failing to address the unique challenges present in fee-only systems; these include volatility of payouts in between block cycles, mining gaps, and enhanced success rates of new and already-established attack vectors. This question drives the development of novel mechanisms specifically tailored for sustainable fee-only operations.

- *How can the practical viability of different TFM designs be empirically evaluated, and their security vulnerabilities mathematically quantified in a fee-only context?*

Transaction fee mechanism design research has thus far been mostly theoretical. Our aim here will be to use historical transaction datasets to simulate a blockchain mining environment, and evaluate the practicability of different TFM designs, such as the 1st-Price Auction, 2nd-Price Auction, EIP-1559 protocol, Burning 2nd-Price Auction (Chung and Shi, 2023), and our own Reserve Pool TFM design. The performance metrics that we will pay close attention to in our simulation will be the average miner payout, the average transaction fee that users pay, the average block sizes, and the variance between all such metrics. We will also investigate the viability of security attacks that can occur within fee-only blockchains, such as fee-sniping.

In answering these questions, our thesis has made the following original research contributions:

- We conducted **Game-Theoretic Analysis of an Exclusively Transaction Fee Reward Blockchain System** (Kruminis and Navaie, 2022), where we modelled transaction inclusion as a dynamic, time-sensitive game. Our work posited that miner revenue is primarily driven by mempool congestion levels, rather than block size constraints. We also showed how dynamic block size adjustments can benefit both users (with lower fees), and miners (with higher fee totals collected).

- We proposed a novel transaction fee mechanism, the **Reserve Pool TFM**, which redirects part of the user fees into a network-shared reserve to act as a buffer in cases of network volatility. We showed that under stable network conditions, our mechanism is *UIC, MIC,* and *OCA-proof*, while simultaneously mitigating volatility, both in terms of the fees users pay, and the final payouts that miners receive.

- We developed a modular, discrete-event **Zero-Block Reward Auction (ZBRA-TFM) simulator**, which is the first of its kind designed specifically to evaluate TFMs under fee-only conditions. Unlike purely theoretical analysis which may overlook operational constraints, or historical data analysis which lacks controlled comparability across mechanisms, we showed that simulation enables reproducible and parameterised experiments that can realistically isolate the impact of each TFM design, providing researchers with valuable results.

- We conducted **thorough empirical and quantitative evaluation of TFMs in fee-only conditions**, both in terms of their security and economic viability. Our findings clearly demonstrated the superiority of the Reserve Pool TFM; it significantly reduced miner payout variance by $\approx 75\%$ whilst also maintaining competitive revenue levels, achieved a $\approx 97\%$ block profitability

rate for miners, and strengthened the overall security of the network, raising the hashrate threshold required for making fee-sniping a rational miner strategy from 13.5% in 1st-Price Auctions, to over 28% for systems under the Reserve Pool mechanism.

- We provided **practical design guidelines** for future protocol architects, addressing the challenges associated with transitioning to fee-only models, and offering concrete recommendations for best practices related to securely and effectively integrating the Reserve Pool TFM into established blockchain systems.

- Building on these recommendations, we conducted a real-world case study of **BB-FLoC: A Blockchain-Based Targeted Advertisement Scheme with k-Anonymity**. We demonstrated the practical applicability of the Reserve Pool TFM, showing how its integration resolved BB-FLoC's free-rider and incentive challenges, while also ensuring sustainable miner compensation and stable network operation in a non-cryptocurrency application.

Together, these contributions create a comprehensive framework that spans theory, simulation, empirical analysis, and practical recommendations that thoroughly evaluate and advance the next generation of transaction fee mechanisms for secure and sustainable fee-only blockchain ecosystems.

1.3 Thesis Outline

Our thesis progresses from foundational analysis of novel mechanism design, to empirical validation, and practical implementation guidance:

Chapter 2 - Background and Related Literature - first establishes the groundwork by examining blockchain fundamentals such as consensus schemes, mining incentive structures, and showcasing a wide variety of blockchain applications. Next, we provide a thorough introduction of transaction fee mechanisms;

we evaluate and define several distinct fee mechanisms and find that none are able to simultaneously satisfy all ideal properties under every condition. Finally, we conclude with a summary of the current state-of-the-art research in the field. This foundational material shapes our research questions, and directly informs the methodology developed in the subsequent chapters.

Chapter 3 - Phasing Out Block Rewards - delves deeper into the inevitable transition from block rewards to fee-only reward systems. We present a comprehensive analysis of Bitcoin mining revenue evolution from its inception in 2009 to the present day, demonstrating how transaction fees still historically constitute only 4.55% of annual miner income on average. Motivated by this ongoing shortfall, we present our primary novel contribution of the thesis: the Reserve Pool TFM design, the first TFM designed specifically for fee-only blockchain environments. The theoretical evaluation of the TFMs presented earlier is conducted in Section 3.4, and forms the foundation for our empirical methodology, directly informing the design, metrics, and overall structure of the practical simulations presented in future chapters.

Chapter 4 - Game-Theoretic Analysis of an Exclusively Transaction Fee Reward Blockchain System - building on the progressive shrinking of block rewards, we incorporate our published research work, Kruminis and Navaie, 2022, in which we modelled transaction inclusion as a dynamic, time-sensitive game. We examine how mempool congestion and block size limits influence miner revenue and the fees users pay for their transactions, providing insights into optimal block sizing in fee-only environments, and further advancing the understanding of purely fee-based blockchain systems.

Chapter 5 - Simulation Setup and Methodology - showcases another one of our primary novel contributions: the Zero-Block Reward Auction (ZBRA-TFM) simulator, the first comprehensive simulation framework designed specifically for comparative TFM analysis in fee-only conditions. We thoroughly describe its design and internal structure, introduce our dataset of over 12 million real-world

Bitcoin transactions, and justify the network parameters chosen based on historical analysis of Bitcoin network statistics.

Chapter 6 - Analysis of Results and Implications for Protocol Design - presents our empirical findings. We detail our experimental scenarios and simulation results, evaluate miner profitability, and demonstrate the advantages of the Reserve Pool TFM over other designs. We find that the Reserve Pool mechanism is able to achieve significantly higher miner profitability rates whilst also maintaining low variance in between payouts. We also provide the first comprehensive security analysis of TFMs in fee-only environments, with mathematical models quantifying attack viability thresholds; the Reserve Pool TFM shows its strength once again, with significantly higher hashrate threshold rates required for fee-sniping attacks compared to other designs. We finish the chapter by offering concrete guidelines for future protocol architects on integrating the Reserve Pool TFM, and giving general recommendations for fee-only blockchain systems.

Chapter 7 - Case Study — BB-FLoC: Blockchain-Based, K-Anonymous Targeted Advertising - presents our published work on a blockchain-based privacy-preserving advertising system, Kruminis, Navaie, and Ascigil, 2025, demonstrating the adaptability of the Reserve Pool TFM beyond cryptocurrency contexts, and its ability to enhance security and stability in non-financial blockchain contexts.

Finally, **Chapter 8 - Conclusions and Future Work -** summarises our contributions, outlines future research directions for fee-only blockchain ecosystems, and provides concluding remarks.

Chapter 2

Background and Related Literature

Overview

This chapter provides the technical and conceptual foundations necessary to understand blockchain technology and its associated transaction fee mechanisms. We begin by defining blockchain as a distributed ledger technology, explaining block structure and the process of transaction propagation and confirmation. Next, we survey the major consensus mechanisms used in practice, summarising their respective trade-offs. We classify different types of blockchain applications, from simple cryptocurrencies to programmable smart-contract platforms used in decentralised finance (DeFi), non-fungible tokens (NFTs), gaming and beyond. We examine the dual-revenue model that motivates block production, before finally introducing transaction fee mechanism design research. We outline the most commonly used TFM designs in practice, defining their rule set and highlighting their respective limitations.

2.1 Basics of Blockchain

Blockchain represents a specific implementation of the broader category of Distributed Ledger Technologies (DLTs), which, at its core, is a database that is shared, replicated, and kept in sync by its network participants. Unlike traditional databases where control is maintained by a single entity, a DLT operates without a central authority, enabling its participants to transfer digital assets in a (typically) peer-to-peer manner. This type of distributed architecture is its defining characteristic, as by maintaining multiple synchronised copies of the ledger across numerous untrusting nodes, the system eliminates a single point of failure, and additionally takes away the need for trusted intermediaries to mediate disputes.

The term *"blockchain"* is fairly descriptive of its internal nature: a series of data structures, called *blocks*, that are cryptographically linked together in a chronological order, a *chain*. Each block unit consists of two main components:

- a **block header**, which contains the metadata that gives blockchain its unique security properties,
- and a **block body**, which holds the ordered list of validated transactions.

The block header contains several essential fields, such as a version number that specifies the protocol rules, a timestamp that marks the block's creation time, and a special property called a Merkle Root (Nakamoto, 2008) which contains a cryptographic summary of all the transactions inside the block's body. Most importantly, it also stores a hash of the previous block's header, a core design feature that gives the blockchain its immutability. A hash is simply a cryptographic function that takes as input some data, and converts it into a fixed-length string of characters that uniquely represents that data, such that even the slightest of changes to the original data would result in a vastly different output hash. By referencing the previous block, each block is linked to the one before it, thus establishing a chronologically consistent chain that extends all the way back to the very first block, known as the *"genesis block"*, which is programmatically created at the inception of

the blockchain. Consequently, any attempt at modifying a past block would alter its hash, which in turn invalidates the previous hash reference in the subsequent block, cascading through every following block in the chain - this inconsistency can be quickly detected by network participants, who would reject the altered block and the chain it claims to form.

The overall blockchain process can be summarised as follows - a user creates a transaction object and broadcasts it to the network, where it in turn enters the *"mempool"*, a list of insofar unconfirmed transactions from which miners select candidates for inclusion into the next block. Once a miner successfully mines a block and it gets accepted by the majority of the network participants, all the transactions included in it can be deemed as being confirmed. This block is also appended onto the blockchain's longest chain, which represents the single, agreed-upon version of its history. Any attempt to spend the same digital asset as part of a different transaction will be automatically rejected because the public blockchain ledger has already deemed it as having been previously spent. While a traditional database can be deemed immutable through access controls, a privileged administrator can still modify its history. In blockchain, changing a previous block requires recalculating not only that block's hash, but also the hashes of all its subsequent blocks. This process, depending on the consensus scheme used (refer to Chapter 2.2), could be immensely costly for its attacker. This point is crucial in understanding the fundamental focus of our thesis - the security of blockchain is not merely a cryptographic guarantee, but also an economic one. That is, in order to rewrite transaction history or undermine the network in any capacity, an attacker must be willing to outspend the combined economic power of all the network's honest participants, which on a secure blockchain should be both computationally and economically infeasible.

2.2 Consensus Schemes

For a blockchain to function, the majority of its participants must reach agreement on a single, valid version of its history, which should be enforced through a set of rules known as a consensus mechanism. Blockchain systems employ a variety of such approaches, including Proof-of-Work (PoW), Proof-of-Stake (PoS), Delegated Proof-of-Stake (DPoS), and Practical Byzantine Fault Tolerance (PBFT).

2.2.1 Proof-of-Work: Consensus Through Computation

This was the first consensus mechanism to solve the double-spending problem in decentralised systems, and to this day remains a popular choice with major cryptocurrencies such as Bitcoin. Here, active network participants, also known as *miners*, compete to solve a complex, but arbitrary, mathematical puzzle; a process that involves repeatedly hashing a block's header with a variable value, known as a *nonce*, until the resulting hash output satisfies a protocol-set difficulty target. Since the output from hashing is unpredictable, miners essentially have to perform a brute-force search when solving the puzzle, thus consuming substantial computational resources. The first miner that finds a valid solution broadcasts their block to the network, and once it is verified and approved by other miners, the block is appended onto the current longest chain of blocks.

The main advantage of PoW is its proven reliability, demonstrated by its long-term success as the mechanism behind blockchains like Bitcoin. Since its security stands on the collective computational power of all of its participants, it makes attempting an attack on the network prohibitively expensive, and therefore practically unfeasible on a large network, as long as $> 50\%$ of its nodes are honest. At the same time, PoW also carries notable disadvantages. For a start, its high energy consumption has sparked environmental concerns, with Bitcoin's carbon footprint comparable to that of entire countries like Belgium or Poland (Digiconomist, 2025). Additionally, the economies of scale inherent to PoW mining

infrastructure have contributed to a progressive consolidation of hash power within large-scale operations, posing notable risks to network decentralisation. Compared to alternative mechanisms, PoW design also often leads to slower transaction processing and variable block creation times, limiting the scalability of its network.

2.2.2 Proof-of-Stake: Consensus Through Collateral

PoS largely emerged as an alternative to the energy-intensive nature of PoW. Here, instead of relying on *miners* who expend their computational resources, PoS relies on *validators* who lock up some amount of their owned network cryptocurrency as an economic collateral, or *stake*. The mechanism then randomly selects a validator to propose the next block in the chain, with the probability of being chosen being proportional to the size of their stake. Other validators in the network must then attest to the validity of the proposed block. If the block proposer is found to have acted maliciously, they risk losing part (or all) of their staked collateral through a penalty mechanism known as *slashing*.

The primary advantage of PoS lies in its energy efficiency; by eliminating the need for computational-hungry mining, PoS reduces its environmental impact dramatically. As an example, when Ethereum moved away from PoW to a PoS consensus mechanism, its energy consumption reduced by over 99% (Kapengut and Mizrach, 2023). PoS also generally has a lower barrier to entry for participation in the consensus process, as it removes the requirement for specialised mining hardware. That being said, the central drawback of PoS is its concentration of power - participants with larger stakes have a higher probability of being selected as the block proposer, which means they have a higher chance of receiving the rewards associated with the production of the block, reinforcing the *"rich-get-richer"* effect. Furthermore, as a newer mechanism, it lacks the long track record of PoW, and can be seen to introduce new attack vectors, such as the *"long-range attack"*, where an attacker attempts to construct an alternative chain starting far back in the chain's history.

2.2.3 Delegated Proof-of-Stake: Consensus Through Delegation

DPoS builds upon the principles introduced from PoS, additionally incorporating a representative governance model. Here, instead of the validators directly participating in block validation, token holders use their stake to proportionally vote on a limited number of *"delegates"*, also known as *"witnesses"*. These elected members are entrusted with the responsibility of confirming transactions and producing new blocks, typically in a round-robin order. The mechanism maintains accountability by allowing token holders to vote out and replace underperforming or malicious delegates, thereby ensuring that block production remains efficient and secure.

DPoS offers high scalability and faster transaction confirmation times, since the limited number of delegates can reach consensus much more quickly. It maintains its energy-efficient advantage compared to PoW, and additionally encourages stakeholder engagement through its voting system. However, DPoS, just like its more general PoS variant, is particularly vulnerable to centralisation. Since block production is consolidated in a small group of representatives, the network is especially susceptible to collusion and vote manipulation, potentially allowing a few powerful users to gain control over the entire network. As such, under DPoS, the security and integrity of the blockchain is strongly dependent on active and informed participation of its token holders in electing honest delegates.

2.2.4 Practical Byzantine Fault Tolerance: Consensus Through Voting

This is a consensus mechanism designed specifically for permissioned blockchains, that is, where all network participants are verified and identifiable. It addresses the challenge described as the *Byzantine Generals' Problem*, where some network participants may fail or act maliciously within the system. Here, participants

are organised into one designated leader/primary node, and the remaining as secondary/backup nodes, with consensus achieved in a series of rounds with three phases:

- **Pre-Prepare Phase:** The leader node begins the process by proposing a new block to all secondary nodes.
- **Prepare Phase:** Each secondary node checks the block for validity, and in turn broadcasts their confirmation of the block to all other nodes.
- **Commit Phase:** If enough (typically 2/3rds) of the nodes agree on the block's validity, they broadcast a *commit* message, signalling that they intend to add the block to their chain. If enough (again typically 2/3rds) commit confirmations are received, the block is considered as finalised.

PBFT can offer high transaction throughput, low latency, low energy impact, and quick transaction finality as it does not need to wait for multiple block confirmations. However, PBFT faces significant limitations in scalability. Achieving consensus requires extensive communication among all the nodes, with the number of messages growing rapidly as the network expands. As such, it is best suited for smaller, permissioned blockchain networks whose participants are well-known and trusted, such as in enterprise or consortium networks. In a permissionless setting, PBFT is vulnerable to Sybil attacks, where a single malicious entity could control multiple identities in an attempt to disrupt the network. This is particularly concerning since PBFT can tolerate at most 1/3 malicious nodes (33%), whereas PoW or PoS mechanisms only fail if an attacker controls the majority of hashing power or stake $> 50\%$.

2.2.5 Consensus Trade-Offs

The choice between consensus mechanisms involves fundamental trade-offs, with each mechanism — PoW, PoS, DPoS, and PBFT — having its own strengths and

Feature	**PoW**	**PoS**	**DPoS**	**PBFT**
Resource Requirement	Hashrate	Staked Coin	Reputation & Stake	Limited (pre-approved validators)
Security Model	Cost to acquire $> 50\%$ of hashrate	Cost to acquire $> 50\%$ of staked assets	Collusion among $> 50\%$ of delegates	Tolerates $< 1/3$ of nodes being faulty
Energy Consumption	Very High	Low	Low	Low
Scalability	Low	Medium	High	Poor (limited to small networks)
Main Risk	Cost to entry, centralisation	Wealth concentration (rich get richer)	Limited number of delegates, vote buying	High communication overhead, Sybil attacks in permissionless networks
Examples	Bitcoin, Dogecoin	Ethereum, Solana	EOS, Tron	Hyperledger Fabric, Zilliqa

Table 2.1: Summary of Consensus Mechanisms

weaknesses. Choosing the right consensus approach depends on the specific goals, participant structure, and operational requirements of the blockchain application at hand. We summarise the consensus mechanisms in Table 2.1.

As we detailed, PoW can establish robust security guarantees through computa-

tional work but suffers from poor energy efficiency and limited throughput, making it ideal for cryptocurrency-like applications where immutability and censorship resistance are critical. PoS offers improvements in energy efficiency and scalability whilst still maintaining reasonable security guarantees, but it leads to increased centralisation risks where wealth accumulation can lead to disproportionate network control. DPoS, as a variant of the more general PoS, delivers similar benefits in even more superior transaction throughput, however, it also amplifies the risks of centralisation as this efficiency requires trust and reliance on active governance by token holders. PBFT delivers exceptionally high overall performance, but is only scalable in small, permissioned environments where all participants are verified and vetted, sacrificing the central blockchain aspect of decentralisation.

2.3 Types of Blockchains

The first practical implementation of blockchain technology was in the form of a cryptocurrency, Bitcoin, which demonstrated its success in securely transferring BTC, its native currency, between participants without the need of a central intermediary. Since then, the evolution of blockchain technology has expanded far beyond its original view of a peer-to-peer digital cash system. A paradigm shift in its capabilities came with the introduction of smart contracts, pioneered most prominently by the Ethereum platform. A smart contract can be simply defined as a self-executing computer program in which the terms of an agreement are embedded directly into code. Once deployed on a blockchain, they run autonomously, ensuring execution that is transparent and irreversible. This programmability resulted in the creation of a whole new domain of blockchain-based applications called Decentralised Applications (dApps). Operating on top of blockchain networks, dApps inherit the core properties of decentralisation, transparency and cryptographic security, making them a compelling model for a wide range of new systems.

As an example, DeFi applications aim to replicate and extend the traditional

finance system on blockchain, enabling services like lending and borrowing (e.g. Aave, 2020), or decentralised exchanges (e.g. Uniswap Labs, 2018). NFTs are unique cryptographic tokens that represent ownership of a specific asset, creating new markets for digital art, music, or collectibles that regularly trade for millions of pounds (e.g. DappRadar, 2025). DApps have also entered the sphere of gaming and metaverses (e.g. Sky Mavis, 2018), giving rise to *"play-to-earn"* models where users acquire in-game assets that can be exchanged for real-world value. Novel use-cases in areas such as digital advertising, to be further explored in Chapter 7.1, further illustrate the widening scope and versatility of blockchain-based systems.

This evolution from simple value transfers, such as peer-to-peer Bitcoin payments, into complex programmable computations found in smart-contract-based dApps, has fundamentally altered the nature of what a *transaction* can be defined as. Such complex and intensive operations will typically consume significantly more block space, and can often carry greater personal value to their users. Blockchain's expanding use-cases have in turn increased the urgency of solving the economic challenges this will present. As these diverse applications mature, they will all require sustainable economic models that do not rely on inflationary block rewards, making the research questions addressed in this thesis increasingly critical for the broader blockchain ecosystem.

2.4 Mining Incentives

The security and integrity of public blockchains is not guaranteed by technology alone; it relies on well-designed economic incentives that motivate block producers[1] to continuously validate transactions, produce new blocks, and secure the network in general. This compensation has historically been composed of a dual-revenue stream which, thus far, has been able to secure blockchain networks:

[1]More recent protocol designs, such as Ethereum's Proposer-Builder Separation (PBS), further decompose this role into distinct block-proposing and block-building functions, though a detailed treatment of such mechanisms is outside the scope of this thesis.

- **Block Reward:** A fixed amount of newly issued cryptocurrency that is awarded to the producer of a block. This guaranteed payout serves as a predictable, stable baseline income that motivates participation in the consensus even during periods of low transaction activity. It also serves as the primary source for bootstrapping the network's security budget, and in distributing the initial supply of its native currency.

- **Transaction Fees:** The implicit difference between a transaction's total input value and its outputs, retained by the miner as compensation for block inclusion. Users may choose to voluntarily leave some fees in order to incentivise faster inclusion, or for priority ordering within a block. Because block space is scarce, users compete by offering higher fees, and miners rationally prioritise transactions with larger fees to maximise earnings. Unlike the predictable block reward, transaction fees fluctuate with network demand, making them a volatile source of income.

These revenue streams are designed to compensate network participants for tangible, real-world costs and risks associated with actively engaging in block creation and validation. In PoW systems, the *work* performed by miners entails significant financial investment in specialised, high-performance hardware, and substantial ongoing cost of electricity required to power it. In PoS, validators *stake* a substantial amount of their own cryptocurrency as collateral, which represents both a major investment and carries with it an opportunity cost. These locked-up funds cannot be used for other purposes, and additionally, they carry the risk of it potentially being *slashed*, for example in the event of operational failure. Thus, active participation in the blockchain network is not a charitable act, but rather an economic venture that is sustained only if the rewards can adequately offset the associated risks.

Some early research into sustaining mining rewards, such as the work by Lin et al., 2018, mainly focused on procedural improvements within the hybrid-reward systems, therefore still operating under the assumption of some form of block

rewards persisting as part of miner compensation. Crucially, they do not tackle the fundamental economic instability that emerges once these subsidies are removed, a problem that necessitates a full redesign of the overall incentive model.

2.5 TFM Basics and Literature

Transaction fee mechanism design was first formally introduced by Roughgarden, 2021a, who in his work observed that the current blockchain fee mechanisms, such as the most commonly found 1st-Price Auction (used by popular blockchains like Bitcoin and the early version of Ethereum) were afflicted with significant inefficiencies and strategic vulnerabilities that classical auction theory had not anticipated.

The fundamental difference stems from the unique characteristics of blockchains; unlike in classical auctions, where an auctioneer is a trusted entity and implements the auction mechanism honestly, in blockchains the "auctioneer", or miner/validator, is in itself a strategic player who may deviate from the set protocol in order to maximise their own utility/gain or intentionally cause harm. It is this particular departure from traditional mechanism design that necessitated the development of new theoretical frameworks that are more specifically tailored to decentralised blockchain systems.

However, this new transaction fee mechanism design approach draws heavily from classical auction theory and mechanism design, rather than representing an entirely novel paradigm. It takes in particular inspiration from Myerson's work on optimal auctions (Myerson, 1981), and the Vickrey–Clarke–Groves mechanism (Nisan and Ronen, 1999), especially with relation to honest user bidding. The novel issues arising from the decentralised blockchain environment typically require adaptations of these classical results. As such, this new approach consists of three main pillars, that when satisfied, should create a *"perfect"* transaction fee mechanism in blockchain systems:

- **User-Incentive Compatibility *(UIC)*:** This property requires that users must bid according to their true value/urgency of the transaction, rather than shading their bids. *UIC* extends the classical concept of dominant-strategy incentive-compatibility (*DSIC*) to blockchain environments, in that honest bidding is a blockchain user's optimal strategy, regardless of what others in the system do. More specifically, in an attempt to avoid overpayment, a user should not gain any utility from intentionally submitting lower bids than their honest valuation of the transaction. Since the space inside a blockchain block is typically a finite resource, without *UIC*, a user may be forced to engage in complex strategic reasoning about other user bidding behaviour. This may lead to system inefficiency where the most valued transactions are not the ones prioritised, fee market uncertainty, and users paying more than the optimal amount for block inclusion.

- **Miner-Incentive Compatibility *(MIC)*:** This property states that miners must implement the transaction inclusion rule(s) exactly as stated by the specifications of the fee mechanism. *MIC* represents a unique characteristic that arises from blockchain systems where a miner, analogous to the "auctioneer" in classical auction theory, is in itself a strategic actor who has the power to not only determine which transactions get confirmed, but may even determine the fees that users pay to the network. Blockchain's environment creates unique incentive challenges where a miner has the ability to engage in various forms of strategic manipulation. These may include injecting fake transactions in order to inflate fee prices, selectively prioritising, or even totally excluding particular transactions from confirmation (censorship). A transaction fee mechanism that satisfies *MIC* is one in which the miner's dominant strategy in maximising their profit from each block comes from implementing the confirmation rules honestly.

- **Off-Chain Agreement Proofness *(OCA)*:** This property states that no

coalition of miners and users conspiring together should be able to increase their joint utility compared to that of the honest implementation of the protocol. The issue stems from miners' ability to deploy side (smart) contracts, which can potentially be used for exploitation by colluding with users or other miners in the system. Classical mechanism design made such agreements difficult to coordinate or enforce, however, blockchain technology enables both users and miners to create verifiable and self-enforcing off-chain agreements that could be used to circumvent honest protocol or mechanism specifications. A TFM design that is *OCA*-proof is one where even when miners and users can communicate and coordinate their strategies, they cannot collectively benefit by deviating from honest protocol implementation.

We must note that recent TFM literature has at times differed in terms of its nomenclature, particularly in regards to characterising miner-user collusion. This comes in terms of two similar, but separate notions; Off-Chain Agreement Proofness (*OCA*), and Side-Chain Agreement of x-user Proofness (x-*SCP*). The distinction is simple; a TFM design is considered *SCP*-proof if any group of users and a miner cannot increase their total utility, whilst *OCA*-proofness compares any collusions of users and a miner against the honest implementation of the protocol. Whilst *SCP* is considered a stronger principle, we will focus on *OCA*-proofness in our work. This is because although it is theoretically possible, *SCP* appears to be too strict for functional designs (Gafni and Yaish, 2022). It is also unclear how to implement it in practice, considering it would potentially require recurring communication with multiple users, made even more difficult if the coordination of the attack is limited to a short block cycle (e.g. 12sec as in the case of Ethereum).

Despite extensive research efforts over the past several years, the literature has insofar presented an impossibility landscape that fundamentally constrains the design space of transaction fee mechanisms. There is currently no existing TFM that has successfully satisfied all three properties simultaneously under all operational conditions. This impossibility result was proved by Chung and Shi, 2023, as they

showed that within a finite block resource space, no transaction fee mechanism can satisfy all three *UIC*, *MIC*, and *OCA* properties. Specifically, the only design that can satisfy both *UIC* and *OCA*-proofness is a trivial mechanism that must always pay the miner nothing, and confirm no transactions, an approach that is clearly unsuitable for real-world use.

2.6 Existing Transaction Fee Mechanism Designs

Currently existing blockchain transaction markets have been heavily influenced by classical auction theory, aiming to efficiently allocate block space while ensuring a fair alignment of incentives between users and miners. The design of these transaction fee mechanisms is typically structured around a set of rules that govern transaction inclusion and confirmation, and the subsequent user payout and miner rewarding processes. These rules also involve determining how transactions are prioritised based on their fee bids, how miners select which transactions to include in the next block, and how (if any) rewards are distributed to miners. A transaction fee mechanism design is conventionally made up of the following rules:

- A *user payout rule* (U_x) which is implemented according to the specifics of the blockchain protocol, and determines how much each user pays for the confirmation of their transaction.

- A *miner reward rule* (M_x) which is also executed by the blockchain protocol, and decides how much revenue a miner receives for creating a valid block.

- An *inclusion rule* (I_x) carried out by the miner, where given a mempool $M = (tx_1, tx_2, .., tx_n)$, a miner decides which of the transactions are to be included in a block.

- A *confirmation rule* (C_x) implemented by the blockchain, which selects a subset of the included transactions to be declared final. Transactions that

were included in a block, but not picked for confirmation typically pay nothing and subsequently return back to the mempool for future consideration.

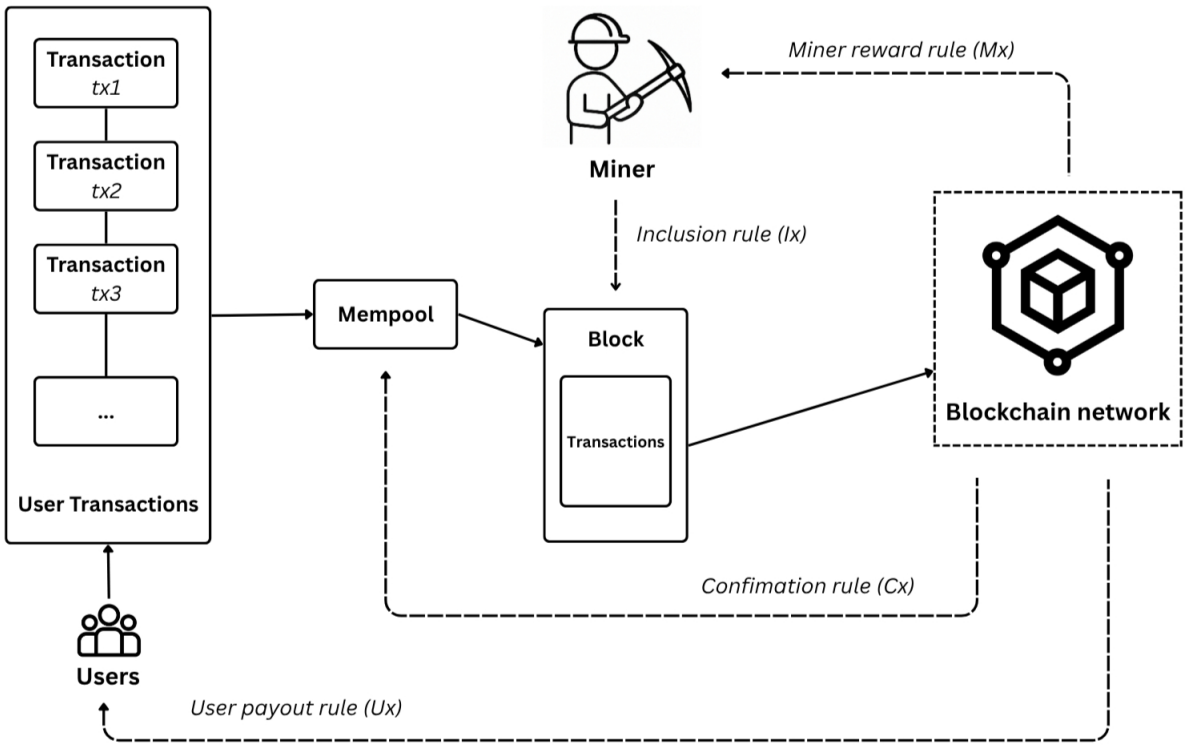

Figure 2.1: An overview of the rules governing a TFM.

A simplified diagram illustrating the sequential operation of a transaction fee mechanism can be seen in Figure 2.1, highlighting the roles of users and miners, and the protocol-specified decision points. Users, governed by the specifics of the *user payout rule* (U_x), independently create and submit their transactions to the mempool, attaching fees or other rewards; this could entail a minimal base fee, a tip to the miner(s) for prioritisation, or any other form of payment. The miner(s) then applies the TFM's *inclusion rule* (I_x), selecting a subset of transactions to incorporate into a block, ordering them, typically prioritising the ones with highest rewards. Once the block is assembled, it is then broadcast to the blockchain network, and the *confirmation rule* (C_x) determines which transactions

are confirmed as finalised, and which are to be returned back into the mempool. Following confirmation, the blockchain network awards the miner(s) for their "work" within the constraints of the *miner reward rule* (M_x); this could be the sum of all user fees, only "tip" amounts, or some type of fixed block rewards. Additionally, some part of the user paid fees may be burned (by sending them to an unreachable account), or allocated into a shared reserve-pool. This entire process encapsulates how a TFM ensures both network security (by incentivising miners), and fairness (by shaping user bidding strategy), key goals sought after by the four rules formally defined above.

Table 2.2: List of sample mempool with 8 transactions

TX Name	**User Fee**
tx1	10
tx2	5
tx3	1
tx4	6
tx5	9
tx6	12
tx7	2
tx8	3

Table 2.3: Overview of user payout and miner reward rules in different TFMs

TFM Style	Miner Revenue	User Fees	Fees Burned
1st-Price Auction	$\underbrace{45}_{=12+10+9+6+5+3}$	$\underbrace{45}_{=12+10+9+6+5+3}$	Ø
2nd-Price Auction	$\underbrace{18}_{=3+3+3+3+3+3}$	$\underbrace{18}_{=3+3+3+3+3+3}$	Ø
EIP-1559[1]	$\underbrace{10}_{=5+3+2}$	$\underbrace{31}_{=12+10+9}$	$\underbrace{21}_{=7+7+7}$
Burning 2nd-Price Auction	$\underbrace{14}_{=6+5+3}$	$\underbrace{18}_{=6+6+6}$	$\underbrace{4}_{=18-14}$ [2]
Reserve Pool[1]	$\underbrace{21}_{=7+7+7}$ [3]	$\underbrace{31}_{=12+10+9}$ [4]	Ø

[1] Base fee = 7, Target block size = 3

[2] User fees (18) − Miner Revenue (14)

[3] Miner payout would also be dependent on past historical payout

[4] Surplus of 10 (=31 − 21) gets sent to the reserve pool

Table 2.4: Overview of transaction inclusion and confirmation rules in different TFMs

TFM Style	Block Size	Confirmed TX	Unconfirmed TX
1st-Price Auction	6	{tx6, tx1, tx5, tx4, tx2, tx8}	Ø
2nd-Price Auction	6	{tx6, tx1, tx5, tx4, tx2, tx8}	Ø
EIP-1559[1]	3	{tx6, tx1, tx5}	Ø
Burning 2nd-Price Auction	6	{tx6, tx1, tx5}	{tx4, tx2, tx8}[3]
Reserve Pool[1]	3	{tx6, tx1, tx5}	Ø

[1] Base fee = 7, Target block size = 3

[3] Unconfirmed transactions return back into the mempool for the next block cycle

For a quick summary of how the designs work in practice, refer to Tables 2.2, 2.3, and 2.4, which include a sample mempool and an overview of TFM rule sets in different TFM designs. In a 1st-Price Auction, miners just fill the block prioritising the transactions with highest fees offered, with users paying exactly what they bid; so, {tx6, tx1, tx5, tx4, tx2, tx8} transactions occupy the block, and the miner collects 45 units. In a 2nd-Price Auction, the same six transactions are confirmed, but each miner only pays the lowest winning bid, which in that case is just 3 units, and as a result the miner collects only 18 units as reward. Alternatively, under the EIP-1559 design, users pay a (dynamically adjusting) base fee, which in our example is 7 units, and any optional tip for transaction ordering. Since the base fees are burned, miners only receive the leftover tip amounts, however, the users still pay the full amounts they offered. The Burning 2nd-Price Auction includes a novel concept of unconfirmed transactions, which are used to set the paying price for users, while the miners get the difference between the actual offered amount minus the set fee. Finally, our reserve pool mechanism involves the miners only being rewarded the base fee amounts of the confirmed bids, and although the users pay the full offered bid, the surplus is redirected into the network-shared reserve rather than being burned, therefore its value is not lost.

Table 2.5: List of TFM notations

Notations	Definitions
U_x	User payout ruleset for TFM
M_x	Miner reward ruleset for TFM
I_x	Transaction inclusion ruleset for TFM
C_x	Transaction confirmation ruleset for TFM
b_i^u	Transaction fee offered by user u for transaction i, expressed as a price per byte or gas unit
bs	Block size limit
bf	Base fee of the block
t_i^u	Tip value i from user u

2.6.1 1st-Price Auction

A 1st-Price Auction, also referred to as a *pay-your-bid* auction, is one of the most common classical auction formats, and is currently in use by Bitcoin. Here, users independently submit their bids, and the dominant strategy of a miner is to append the highest bids into a block for confirmation. Ergo, miners try to attach as many fee-paying transactions up to the block size limit as possible. Transactions in a block are ordered in ascending order based on their fee, with each user paying what they offered, and any ties are settled arbitrarily.

Table 2.6: 1st-Price auction TFM rules

1st-Price Auction	
I_{1st}	Miner includes the highest paid bids in ascending order, up to the block size limit bs, breaking ties arbitrarily.
C_{1st}	All included transactions are confirmed.
U_{1st}	Users of confirmed bids pay their offered gas price b_i times the gas used by their transaction.
M_{1st}	Miner receives the sum total of fees from all the confirmed transactions in a block.

Whilst satisfying the conditions of *MIC (since miner payout is the sum of highest bids)* and *OCA (users pay whatever they offered and miner gets all of the fees, so there is no need for any off-chain contracts)*, it is trivial to see how the 1st-Price Auction is not UIC. Suppose a block size of 1, and transactions $tx_6 = 12$, $tx_8 = 3$, $tx_3 = 1$; bidder of tx_6 could save 8 by bidding 4 instead of 12, and still get their transaction confirmed. This approach of bid shading is the key downfall of the 1st-Price Auction as it encourages overbidding. It fails to provide users with easy fee estimation, and can lead to instability of miner revenue due to the variation of fees offered in between different block cycles.

2.6.2 2nd-Price Auction

The 2nd-Price Auction is among the most popular classical auction formats and is renowned for its ability to foster truthful bidding behaviour. The central idea here is simple; users submit their bids independently and according to their true value, and miners attach the highest offering bids for inclusion (and confirmation). However, users only pay the fee offered by the lowest included transaction, rather than what they offered themselves. In this way, users either pay what they value their transaction to be at, or pay the price that others also offered for successful

Table 2.7: 2nd-Price auction TFM rules

2nd-Price Auction	
I_{2nd}	Miner includes the highest paid bids in ascending order, up to the block size limit bs, breaking ties arbitrarily.
C_{2nd}	All included transactions are confirmed.
U_{2nd}	Users of confirmed bids pay the gas price of the lowest included bid b_{bs-1}, times the gas used by their transaction.
M_{2nd}	Miner receives the total of b_{bs-1} bid times the total gas used by all the confirmed transactions in the block.

inclusion, whichever is lower, satisfying *UIC* principles. [2]

Although very effective in traditional settings, the 2nd-Price Auction is a prime example of classical mechanisms failing in a blockchain environment. The main factor for this is that 2nd-Price Auctions require a trusted third party to correctly implement the protocol in order to determine the winner and the bid to be paid. Alternatively, the key idea behind blockchains is decentralised consensus; with such a TFM design, miners can inject fake transactions, thus failing the *MIC* property. Imagine a scenario where $bs = 3$ and the mempool consists of 4 transactions $tx_6 = 12$, $tx_1 = 10$, $tx_7 = 2$, $tx_3 = 1$. An honest miner implementing the 2nd-Price Auction protocol would be paid 6 *(each user paying the lowest winning bid, which in this case is 2)*. However, a miner can inject a fake transaction $tx_9 = 9$, and get paid 18 *(users of bids tx_6 and tx_1 paying 9, and the third fake bid of 9 going back to the miner with no penalty)*. Likewise, a miner can collude with bidder of tx_3 by telling them to bid 9, increasing the joint miner/user utility from *2(3)+0=6* to *9(3)+0=24*, failing to satisfy *OCA* guarantees.

[2] We must note that additionally, a 2nd-Price Auction can be setup in a way that users pay the fee of the highest excluded transaction. However, due to the specifications of blockchains and in the interest of not wanting to waste block space by including unconfirmed transactions, our model will utilise the value of the lowest winning bid.

Table 2.8: EIP-1559 auction TFM rules

	EIP-1559
I_e	Miner includes any transactions offering $\geq bf$, breaking ties arbitrarily. In the case of the block being filled beyond its capacity, transactions are ordered by the additional t_i value.
C_e	All included transactions are confirmed.
U_e	Users of confirmed bids pay their offered gas price b_i times the gas used by their transaction, and any t_i donated.
M_e	The base fee associated with each transaction is burned, therefore, miner revenue is solely the t_i users attach to their transactions for transaction ordering.

2.6.3 EIP-1559

In a departure from the 1st-Price Auction and its shortcomings, Ethereum has recently implemented the EIP-1559 protocol. This mechanism departs significantly from the previous TFMs. Here, users must pay a base fee (bf) for any inclusion into a block. The base fee is determined by the size of the preceding blocks; a target block size is set by the blockchain, however, the block may expand to double (or some other extent) its size. If the previous block was underfilled ($< targetblocksize$), the fee is lowered, and on the other hand, if the previous block was overfilled ($> targetblocksize$), then the fee is increased. The aim is to dynamically adjust the base fee depending on latest network demands, whilst at the same time allowing for potential uncertainty by varying block size limit as needed. Crucially, the base fee that is paid by the users is burned (that is, sent to an account that cannot be accessed by anyone), and the only revenue a miner may attain is from the additional *"tips"* (t_i) that users may attach within their transaction.

EIP-1559 succeeds in several respects. It enables straightforward fee estimation as users pay a predetermined base fee for confirmation, whilst tips allow for priority ordering within a block, supporting *UIC*. Since the base fee is burned, off-chain agreements between users and miners become costly to execute, thereby satisfying

OCA. Miner revenue comprises the sum of included tips, and fake transactions incur a real cost that is not returned to the miner, making honest inclusion the optimal strategy *(MIC)*. However, these properties only hold in an uncongested environment. In the extreme case where the mempool contains more valid base-fee paying transactions than the block size limit permits *(≥200% of bs)*, inclusion priority is determined by the tip t_i amount offered, effectively returning to the dynamics of a 1st-Price Auction and once again failing to attain *UIC* guarantees.

2.6.4 Burning 2nd-Price Auction

One of the most promising TFMs, the Burning 2nd-Price Auction, was recently introduced by Chung and Shi, 2023. Their design makes use of the novelty of the blockchain auction environment by introducing a relatively new concept of unconfirmed transactions. Such a transaction subset is used to set the price for the confirmed transactions, and determines the miner payout. It works in the following manner; imagine a scenario of $bs = 4$, a subset of 2 transactions being confirmed, and a mempool of $tx_6 = 12$, $tx_5 = 9$, $tx_8 = 3$, and $tx_7 = 2$. All 4 transactions are included into the block, with tx_6 and tx_5 being confirmed, tx_8 and tx_7 unconfirmed. The confirmed bids all pay the value of $tx_8 = 3$ *(user payment is 3+3=6)*, miner payout being 5 *($tx_8 = 3$ + $tx_7 = 2$)*, and the remainder of 1 is burned.

Burning 2nd-Price Auction	
I_b	Miner includes the highest paid bids in ascending order, up to the block size limit bs, breaking ties arbitrarily.
C_b	Only a subset N of the included transactions are confirmed. All other $bs - N$ transactions are included into the block but deemed unconfirmed, and thus are later returned back into the mempool for potential future confirmation.
U_b	All confirmed bids pay the value of the highest offering $bs - N$ bid, unconfirmed bids pay nothing.
M_b	Miner gets paid the sum of the $bs - N$ bids. Any extra amount remaining from the user payout is burned.

Table 2.9: Burning 2nd-price auction TFM rules

Similarly to the classical 2nd-Price Auction format, the Burning 2nd-Price Auction is *UIC*, with each user's optimal behaviour to bid according to their true value. In addition, the idea of an unconfirmed transaction subset and the burning mechanism brings forth additional perks. With respect to fake bids injected by the miner, these now incur a *future* cost, called the γ-discount parameter. Whilst it would incur an extra cost to the users' individual payout in the moment, the unconfirmed fake bid would still be eligible for future confirmation, and if successful, would suffer some relative cost. Chung and Shi, 2023 claims that under weak incentive compatibility *(i.e. γ=1, where the fake transaction suffers the full cost of its bid)*, the Burning 2nd-Price Auction is UIC, MIC and OCA-proof. Nevertheless, this mechanism fails to address the fundamental challenges of revenue stability and miner incentivisation in fee-only blockchain environments, as it still relies on block rewards for sustained network security.

2.7 State-Of-The-Art Landscape

The current state-of-the-art research landscape for TFM design has largely focused on innovative mechanism designs under the assumption of continued block rewards. The impossibility results demonstrated by Chung and Shi, 2023, showing that no single mechanism can simultaneously achieve perfect user-incentive compatibility, miner-incentive compatibility, and collusion resistance without forgoing miner revenue, has effectively closed the search for a "dream" TFM. As such, the field has pivoted towards two main avenues; designing mechanisms that make explicit but pragmatic trade-offs (e.g. Roughgarden, 2021b), and exploring models that relax the traditional assumptions (e.g. Wu, Shi, and Chung, 2024, Bahrani, Garimidi, and Roughgarden, 2023, Ganesh, Thomas, and Weinberg, 2024). This post-impossibility perspective reframes the challenge from finding a perfect solution, to engineering the least imperfect one tailored to specific priorities, whether it be revenue stability, user experience, or network efficiency.

In parallel, recent blockchain security research has begun to refine the threat landscape in fee-only environments. Bao, 2024, extending directly upon Carlsten et al., 2016, showed that undercutting is not necessarily an ever-present threat, but rather one that is highly contingent on specific mempool dynamics. Blockchains remain largely safe from undercutting when the block size limit is small relative to existing transaction volume, but become increasingly vulnerable when infrequent, high-value *"whale"* transactions, that is transactions with unusually high fees attached to it, end up creating outlier blocks that incentivise miners to fork the chain. This suggests that the main risk lies in the *fee gradient*, and the statistical properties of the mempool, indicating that effective long-term security may therefore depend as much on managing the mempool's composition, as on any transaction fee mechanism itself.

A notable recent shift in transaction fee mechanism research has been the movement beyond pure economic efficiency, and more into goals such as fairness and inclusivity. Kiayias et al., 2024 introduced the notion of *"traffic diversity"*,

proposing tiered pricing mechanisms, whereby a predefined portion of the block space would be reserved for low-urgency transactions that are paying less. This presented a paradigm shift from the single, unified fee market into one that is multi-tiered, aiming to preserve accessibility to a wide variety of users, not just those who can afford to consistently outbid others.

Beyond economic properties, attention has also turned to how TFM design must evolve alongside fundamental shifts in consensus architecture. The emergence of leaderless, DAG-based consensus protocols in blockchains like *Aptos* completely breaks apart the single-leader, single-auctioneer model that currently underpins nearly all existing TFM literature. In such multi-proposer environments where validators build and broadcast blocks at the same time, a new incentive challenge arises in ensuring that validators honestly report transactions, and agree on a final, canonical block. This has given rise to new TFMs like the *"first-price auction with equal sharing"*, proposed in Garimidi, Heimbach, and Roughgarden, 2025, which was designed specifically for systems with multiple simultaneous block producers. Their work suggested that future TFMs will likely need to co-evolve with consensus protocols, moving beyond *one-size-fits-all* solutions, and towards mechanisms that are specifically tailored to complex decentralised network topologies.

Ultimately, whilst the literature has made significant progress in refining TFM design under traditional assumptions, a critical gap remains largely unaddressed: the viability of TFMs in the absence of block rewards entirely. The studies surveyed above either assume a continued presence of block rewards, or focus on isolated aspects of fee market behaviour without considering the broader systemic implications of a fee-only environment.

Chapter 3

Phasing Out Block Rewards

Overview

In this chapter, we establish the foundational premise of the thesis: that the transition to a fee-only model represents an inherent protocol-level challenge that threatens the long-term sustainability of blockchain systems. We examine this problem in detail and support our analysis with empirical evidence drawn from historical Bitcoin network data. Based on this analysis, we develop a data-driven argument demonstrating that the current fee market is inadequately prepared to operate in a fee-only environment. In response, we introduce our novel Reserve Pool TFM design, discussing its advantages and shortcomings relative to existing mechanisms.

3.1 Motivation and Background

The security of PoW blockchains, such as Bitcoin, can be thought of as being fundamentally rooted as an economic construct, rather than a purely cryptographic one. At its core, the design idea is to make honest participation more profitable than malicious behaviour. Within this kind of framework, miners are modelled as rational, yet strategic economic actors who commit their computational resources

under the expectation of financial rewards for it. Therefore, blockchain security can also be viewed through a game-theoretic lens, where the protocol's resilience against attack vectors is determined by its security budget, that is, the total economic value distributed to its miners or validators. This *"budget"* must remain sufficiently high in order to render attacks (such as double-spending) prohibitively costly for rational miners to attempt. Accordingly, the stability and predictability of miners' revenue streams, which sustain the blockchain's security budget, are paramount to the long-term survival of the system.

Miner compensation has historically come from two distinct sources; a predictable block reward, designed to establish a baseline level of network security, and variable transaction fees, which operate as a market-driven payment mechanism for the validation and ordering of transactions. This dual-revenue model has thus far ensured a stable security floor, shielding the network's security budget from fluctuations of user demands. However, as can be seen, they serve fundamentally different roles. This distinction between them is essential in understanding the profound impact that removing the block reward would have on the whole ecosystem of blockchains; it dismantles the primary pillar of stability that blockchain networks have historically depended on.

It must be noted that the programmed elimination of block rewards is not some unforeseen flaw, but an intentional design feature that aims to enforce digital scarcity and control monetary inflation. For instance, in Bitcoin, the block reward has been halved approximately every 4 years, going from an initial value of 50 BTC to just 3.125 BTC by 2025. Thus, over time, the network will have to self-fund its security budget entirely through user-paid fees. Whether the volatile fee market is actually capable of reliably sustaining the network remains an unresolved question, and forms the central focus of our thesis. Importantly, this problem is not some distant theoretical concern, as while Bitcoin's block rewards will not completely disappear until sometime around the year 2140, their value will drop to negligible levels long before then, making the transition to dependency on fees an urgent and

pressing issue for blockchain systems today.

The primary risk is that without a stable and sufficient security budget, the economic incentives that currently drive honest miner participation will erode, forcing rational miners to instead adopt behavioural strategies that destabilise overall network consensus. For example, fluctuations in miner payouts may lead to scenarios where it becomes more profitable for a miner to attempt to fork a chain and capture rewards from a previous block, rather than continuing to build on the current chain and collect rewards from a subsequent low-value block. This inversion of rational miner strategies introduces a whole new class of vulnerabilities, whilst also amplifying the viability of existing ones - a dynamic that we will examine in greater detail in Chapter 6.4.

3.2 Evolution of Miner Revenue

As blockchain technology is still a relatively recent technology, there are no real-world examples to draw upon to illustrate the risks of this fee-only shift. Nonetheless, that does not mean that we cannot provide an empirical foundation for the concerns outlined in Section 3.1. By examining historical blockchain data, we can build a data-driven argument stating that the currently existing fee market in major blockchains is fundamentally unprepared to take over the role currently fulfilled by block rewards. For this, we believe that Bitcoin, as the first and still the most prominent blockchain application, stands as a particularly suitable subject for our historical analysis. With a market capitalisation of approximately £1.7 trillion, comparable to the GDP (Worldometer, 2025) of Russia (£1.58 trillion) or Canada (£1.67 trillion), and a 58% share of the overall cryptocurrency market cap (CoinMarketCap, 2025), its long operational history and scale provide a rich and reliable dataset for empirical study. Additionally, Bitcoin is also at risk from the challenges posed by the transition to a fee-only phase, making it directly relevant to the issues under our investigation. Of course, Bitcoin is not alone in

this. An evaluation of the top 50 blockchain systems by market capitalisation (CoinMarketCap, 2025) shows that such problems are also inherent in other major blockchains - 74% of the leading blockchains still rely on fixed block rewards in rewarding miners, and 58% enforce a hard supply cap. These figures highlight that the long-term sustainability concerns associated with the fee-phase transition are not confined to just Bitcoin, but are indeed relevant to the majority of the top blockchain networks.

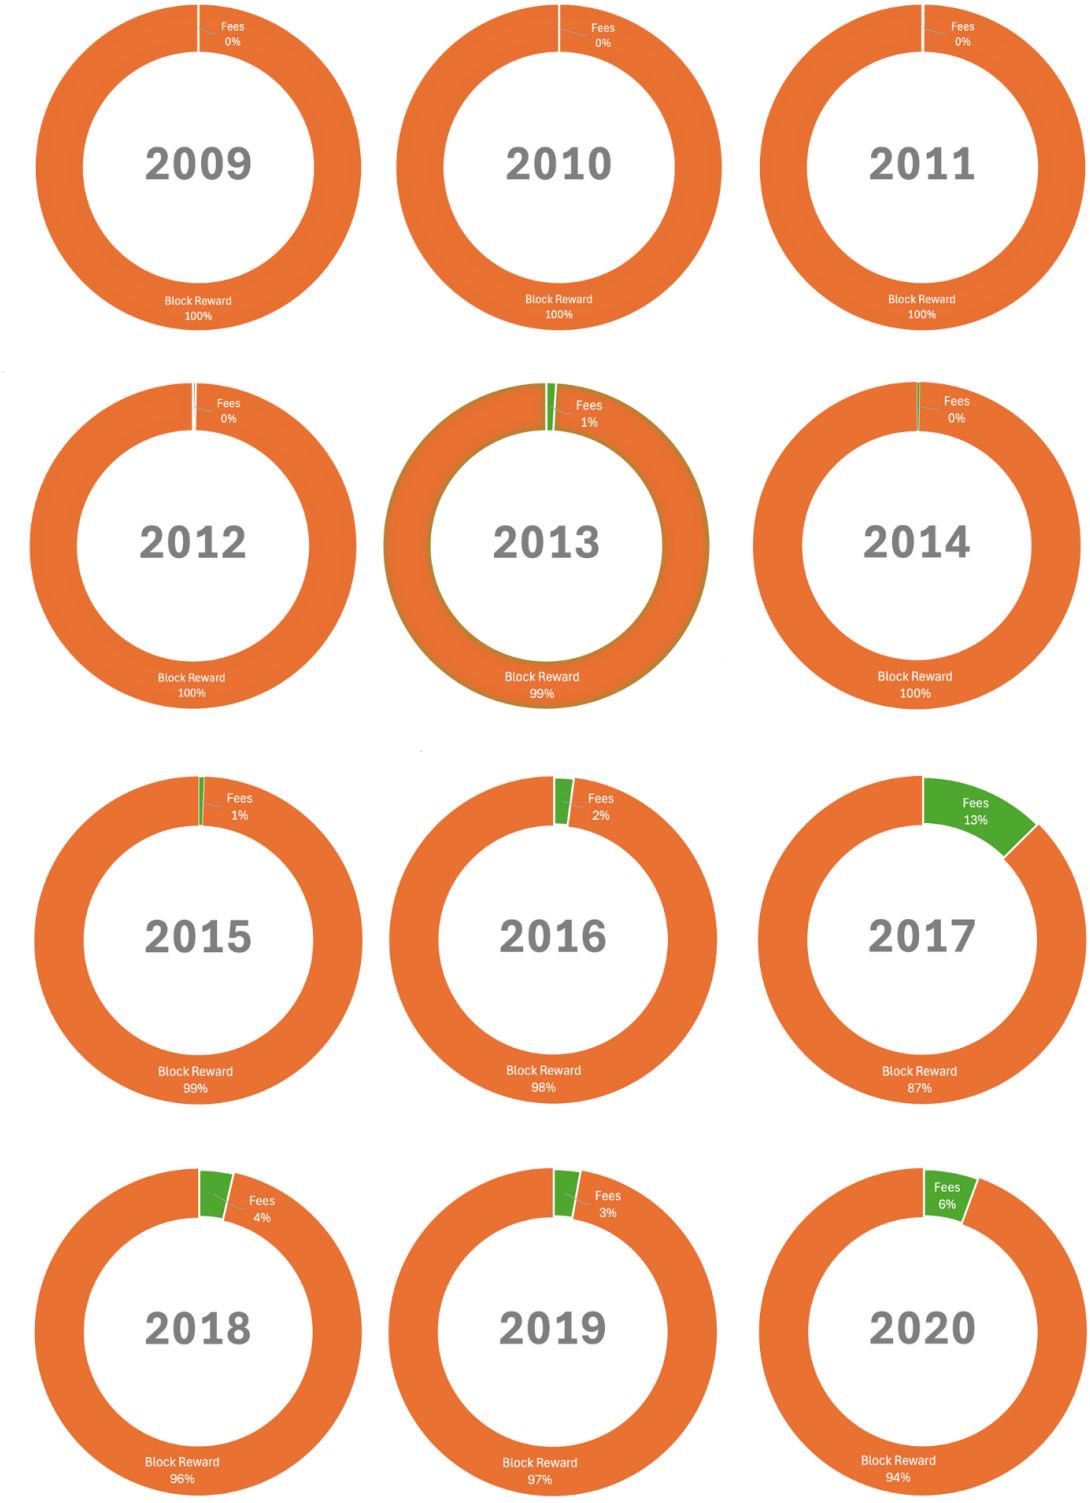

Figure 3.1: Bitcoin mining income composition by year (2009-2020)

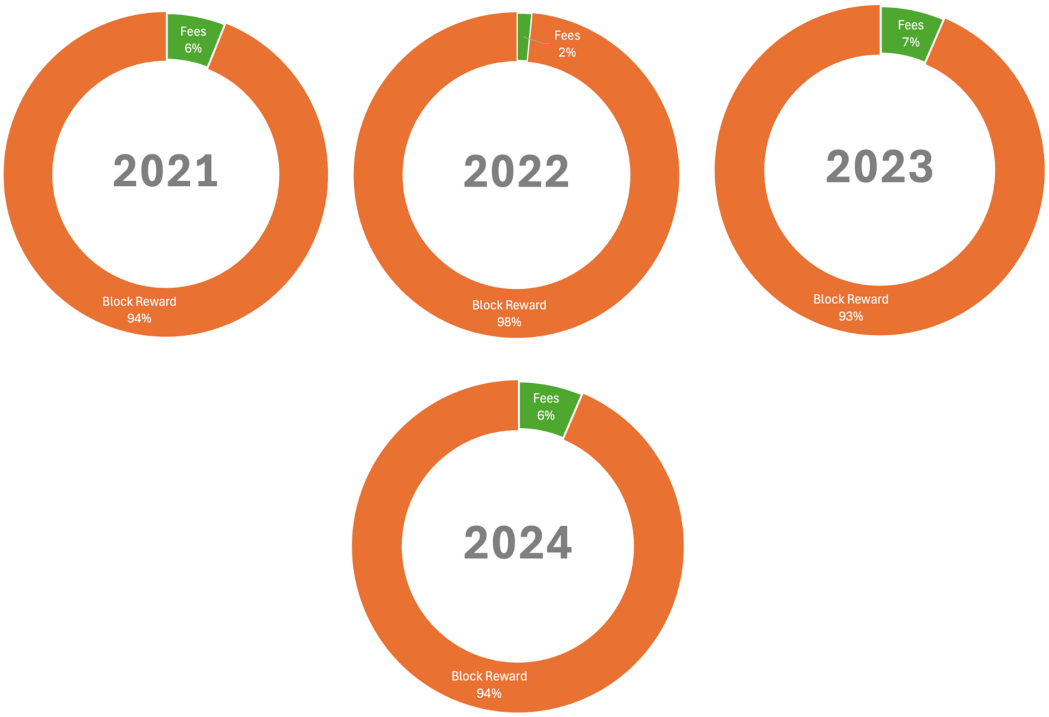

Figure 3.2: Bitcoin mining income composition by year (2021-2024)

Historical data underscores the overwhelming and persistent dominance of block rewards in funding Bitcoin's security budget. Miner income composition from Bitcoin's inception in 2009 through to 2024, illustrated in Figures 3.1 and 3.2, reveals how transaction fees have consistently constituted only a negligible fraction of total miner revenue across their 16-plus years of operation. For instance, in 2019, a whole decade after its launch, fees still represented just 3% of miner revenue, compared to 97% obtained from block rewards. Even as block rewards diminished to 3.125 BTC per block, transaction fees were a mere 6% of miner income. This stark disparity clearly illustrates that the network has never been meaningfully funded by its users, instead being wholly reliant on the subsidy of block rewards that are programmatically disappearing.

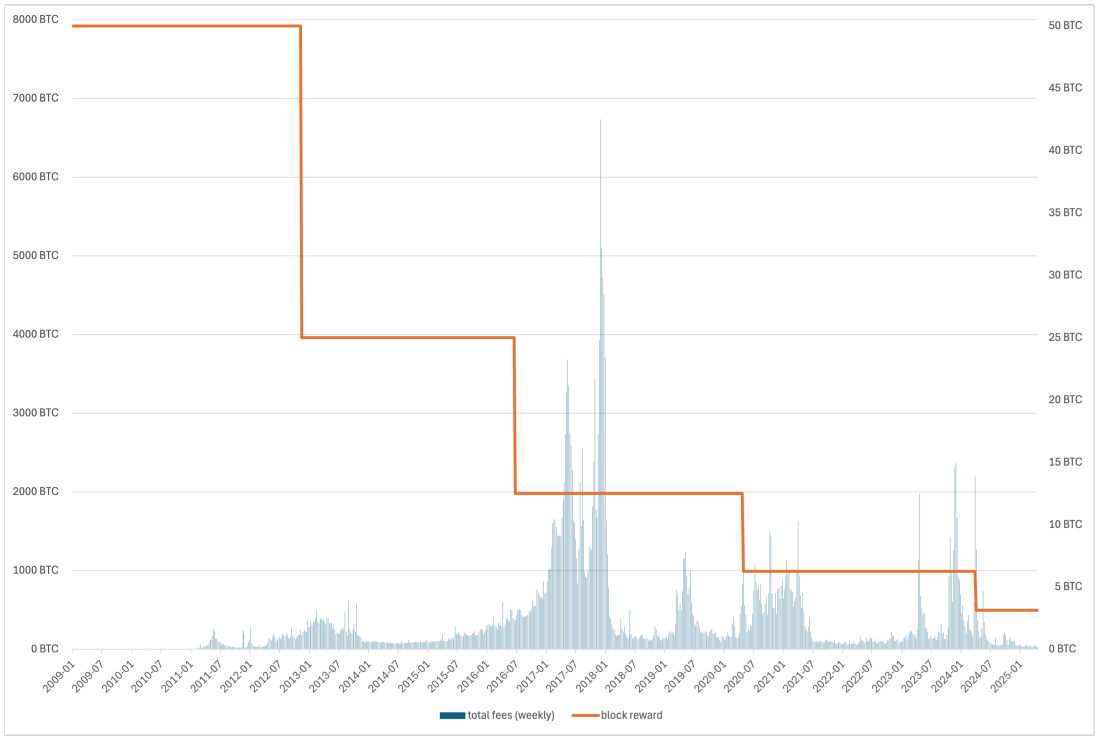

Figure 3.3: Total fees per block (average per week) and the block reward in Bitcoin (2009-2024)

Figure 3.3 plots the lifetime of block rewards against the total fees collected per block, once again highlighting the fundamental contrast between the predictable, stable block rewards, and the volatile nature of transaction fees. In the absence of these block rewards, miner revenue becomes an erratic, market-driven and utterly unpredictable revenue stream, with profound implications on the mining industry's entire business model. As will be explored in greater depth in Section 6.3, mining is a capital-intensive operation that requires substantial upfront investment in specialised hardware, and significant ongoing operational costs; such investments are only feasible when revenue streams are predictable, as is provided by block rewards. A fee-only model would, in turn, transform mining into a high-risk, speculative venture. Such an economic reality threatens to drive out all but the most risk-averse mining participants, potentially increasing centralisation and reducing overall network security.

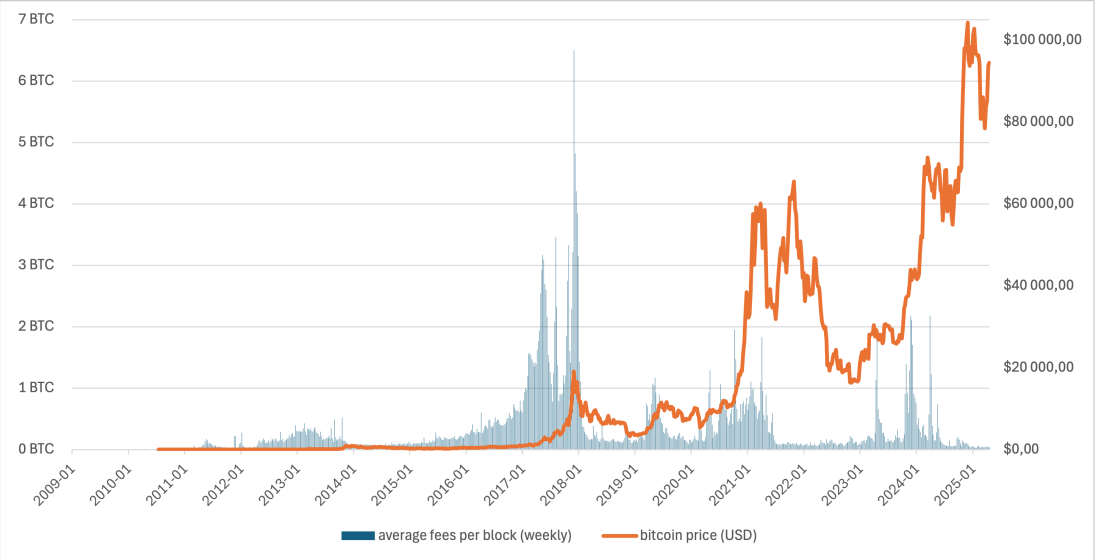

Figure 3.4: Average fees per block (per week) and the Bitcoin price (USD) (2009-2024)

The vulnerability of fee-only networks is further illustrated in Figure 3.4, which plots the average fees per block against the USD value of Bitcoin. This data indicates that in a fee-only model, the network's security budget would peak when the price of Bitcoin is high, and would conversely fall sharply when the price is low, coincidentally leaving it most exposed when it can least afford to be. Because the cost of attacking a PoW network is directly proportional to its hashrate (refer to Section 6.4), which in itself is financed by the security budget, a lower security budget results in a lower cost to attack. Hence, a well-funded miner could potentially undermine the security of the network at a fraction of the cost required when Bitcoin is more valuable.

All such historical data and its analysis reinforces the hypothesis that transaction fees could never even come close to replacing the inherent value that comes from block rewards, foreshadowing a predictable shortfall in the economic and security budget of future fee-only blockchains. From 2009 to 2024, fees accounted for only 3.19% of yearly miner revenue on average, and Figures 3.1 – 3.4 offer no indication that a sustained, and monumental increase in fee revenue is even plausible. All of

this taken together points to a concerning, yet unavoidable conclusion: The Bitcoin network faces a significant structural deficit, underscoring the urgent need for novel protocol redesign that is capable of sustaining the stability and revenue that is necessary to maintain network integrity in the absence of block rewards.

3.3 Reserve Pool TFM

Finally, we introduce an innovative TFM design where rather than paying the fees directly to the miners or burning them, we transfer them to a publicly-accessible reserve pool account for potential future utilisation. Similarly to the EIP-1559 design, our protocol includes a dynamic base fee that is adjusted according to historical block space usage. However, rather than being paid solely by the tips, miners here get paid a sum of the base fees calculated at the target block size; the eventual miner payout for each block is then set as a rolling window average of these target-size base fee amounts over the past n-blocks (e.g. 144-blocks as per our simulations, or a full day's worth of blocks at the 10-minute block cycle as in Bitcoin), which further helps in smoothing out historical fluctuations. As an example, if the base fee rises to a high enough point that the next block ends up being severely underfilled, the miner of such a block may use up the resources of the reserve pool to yield the same revenue that it would under normal network conditions. At the same time, if the current block ends up being overfilled, any extra user payments beyond the miner payout are directed to the reserve pool. To allow for transaction ordering, users may still attach an additional tip value, however, the sum of such values is not fully directed to the miner of the block; it gets sent to a miners' *private* pool, from which they can withdraw a portion ω of it over time if they successfully mine blocks in the future.

	Reserve-Pool TFM
I_r	Miner includes any transactions offering $\geq bf$, breaking ties arbitrarily. In the case of the block being filled beyond its capacity, transactions are ordered by the additional t_i value.
C_r	All included transactions are confirmed.
U_r	Users of confirmed bids pay their offered gas price b_i times the gas used by their transaction, and any t_i donated.
M_r	Miners get paid the sum of the bf times the optimal block size (bs) gas amount. If the block is below the optimal capacity ($< targetblocksize$), withdraw the deficit from the reserve pool. Conversely, if ($> targetblocksize$), any surplus beyond the optimal block size gets sent to the reserve pool.

Table 3.1: Reserve pool auction TFM rules

A clear concern with such a mechanism would be miners creating empty blocks and draining the reserve pool; if the reserve pool is abundant enough to fully cover the miner at a particular block cycle, then mining an empty block versus a full block returns the same utility for the miner, yet a completely different one for the regular users of the network. This type of behaviour would not be exclusively malicious, but also rational. Transactions take up extra computation and storage, and at the same time, a smaller block will propagate faster throughout the network, giving the greedy miner an advantage in having their block being the first one accepted. It is critical that the reserve pool assets are never depleted, since an empty pool would cause the entire system to halt.

To address this issue, we will set a limit on the amount that can be extracted from the reserve pool within a single block cycle, e.g. only 50% of the current target block payout. This way the malicious miner would be irrationally missing out on extra income by not including some valid transactions. Additionally, we will stipulate that miners are only allowed to take out of the reserve pool what they positively contributed to it in the past. This way, adding to the reserve pool acts as

a form of a future security blanket. For example, in the extreme case of an empty mempool, an honest miner could take the conventional maximum of 50% from the reserve pool, and then also be allowed to get the rest of the 50% because they added to the pool in the past. [1].

There are several other considerations we could add to our protocol. For example, if the reserve pool reaches a set target amount, future block miners, for a short period, could withdraw extra revenue from the pool even if network conditions are stable e.g. via the reinstatement of block rewards. This way the resources of a rich reserve pool are not just wasted away with little purpose. Alternatively, the "surplus" beyond the target amount could be re-awarded to previous pool contributors via rebates. In any case, future blockchain architects would need to decide on the best-fit scenario based on their specific blockchain application.

3.4 Comparative Analysis of Incentive Properties

As shown, each transaction fee mechanism has its own strengths and weaknesses, and certain designs may be more applicable to specific blockchain applications than others. For a summary of key properties for each TFM design, refer to Table 3.4. 1st-Price Auctions are the simplest to implement and incentivise honest protocol following from the miners, however, they encourage bid shading and off-chain collusion between users and miners. In contrast, 2nd-Price Auctions promote honest bidding from users but crucially rely on miners being trustworthy, a condition that may not hold in adversarial or strategic environments in decentralised systems. The EIP-1559 protocol appears to satisfy all three properties - *UIC*, *MIC*, and *OCA* proofness, however, such properties only hold under uncongested network conditions. When there are more base fee paying transactions that can fit within a block size limit, the EIP-1559 mechanism reverts back to 1st-Price Auction behaviour; some base fee-paying transactions may need to be excluded from block inclusion, so users

[1] Assuming that the past contribution of such a miner surpasses or equals the value of that 50%.

need to submit strategic bids, making it no longer *UIC*-compliant. Similarly, the Burning 2nd-Price Auction can also satisfy all three properties, though once again, only under specific operational conditions. Therefore, there still remains a need for a transaction fee mechanism that holds up not just in theory, but in practice as well.

TFM	**UIC**	**MIC**	**OCA-proof**
1st-Price Auction	×	✓	×
2nd-Price Auction	✓	×	×
EIP-1559	✓[1]	✓	✓
Burning 2nd-Price Auction	✓	✓	✓[2]
Reserve-Pool Auction	✓[1]	✓	✓[3]

[1] It fails to satisfy *UIC* during congestion periods when there are more base-fee paying transactions than available space inside a block, as it returns to behaving like a first-price auction and ordering transactions based on "tips" value.

[2] Only satisfies *OCA* under weak incentive compatibility assumptions

[3] Provided the reserve pool is adequately built-up to cover for periods of low activity.

Table 3.2: Properties of the main TFM designs.

The newly proposed Reserve Pool TFM, similarly to the EIP-1559 protocol, is *UIC*-compliant only under stable network conditions, as a congested block cycle can push away some base fee paying transactions from inclusion, requiring users to supplement their fees with extra strategic bidding. That said, the Reserve Pool TFM is clearly *MIC*-compliant, as miners gain no value from submitting fake transactions. The protocol enforces base fees for valid inclusion within a block, forcing fake transactions to incur a loss by paying the base fee as a minimum -

blocks that include any invalid transactions will be rejected by the network. Thus, the dominant strategy for miners is to honestly follow the protocol. Finally, our design can also be deemed as *OCA*-proof, though again, under certain conditions. For a transaction to be valid, it must at least pay the base fee; any additional fees beyond it are transparently transferred into the network-shared reserve pool. As an additional incentive, miners' positive (and negative) contributions to this public reserve pool are documented in a miner-specific private pool. As a reminder, in the case of an empty block being mined, only a certain percentage of revenue can be extracted from the public reserve pool. In order to get the full nominal payout reward, a miner needs to have positively contributed to the pool in the past, and such a value would need to be taken from said private pool.

Example: Suppose a user values their transaction inclusion at $v = 10$, with the base fee set at $bf = 5$. The protocol rules state that only transactions paying at least the base fee are to be deemed valid, so the user must offer at least 5 on-chain for inclusion.Off-chain deals, such as directly paying the fees to the miner, offer no meaningful utility advantage, since the base fee must still be paid on-chain and the extra *tips* can be transparently recouped by the miner in the event of low network activity, also giving the miner extra security as a buffer.

Therefore, the joint utility obtainable via off-chain collusion is never greater than what the coalition could achieve by simply following the protocol honestly over time. The only case where this property may fail is if the public reserve pool is empty, or has not yet been sufficiently replenished. In such an event, the miner who has positively contributed to the pool in the past may be unable to withdraw their earned rewards, though again, future block cycles may make it possible again. Overall, the enforced base fee and private pool documentation ensures that the Reserve Pool TFM is *OCA*-proof under stable, mature blockchain network conditions.

In evaluating transaction fee mechanisms, it is also crucial to go beyond theoretical incentive properties. We also need to consider their practical use-cases when deployed in real-world blockchain environments. For example, while

the Burning 2nd-Price Auction may appear particularly desirable because of its theoretical properties, in practice, it suffers from several disadvantages. Specifically, its real-world implementation effectively reduces network throughput in half as a significant portion of the block space must be reserved for unconfirmed transactions. These transactions occupy the valuable and already limited space within a block, yet end up going back to the mempool for consideration in the next block cycle. For cryptocurrency systems, such an inefficiency would further bottleneck the already comparably low transaction throughput (G. Zhang, Zhao, and Si, 2023) typical of major blockchains like Bitcoin and Ethereum, and would also result in higher latency and congestion under heavy load.

Additionally, both the 2nd-Price Auction and EIP-1559 protocols specify fee-burning approaches, where some portion of the submitted transaction fees are permanently removed from circulation. Although this can have the benefit of deflationary pressure and potentially increase the value of the remaining coins, it also restricts the monetary base available for essential system functions such as transaction liquidity and user participation. If fee burning is continued over a long period of time, it may even completely deplete the entire supply of coins, threatening the sustainability and economic viability of the system. As an example, after the implementation of EIP-1559 in Ethereum in August 2021, more than 5.32 million ETH has already been burned (Beaconcha.in, 2025) as of August 2025, representing approximately 19 billion USD in value. As an additional side effect, with fee burning reducing miner rewards, overall network security is also weakened, making the whole system increasingly more vulnerable to attacks in fee-only blockchain systems.

Chapter 4

Game-Theoretic Analysis of an Exclusively Transaction Fee Reward Blockchain System

Overview

Having established the empirical reality of the fee market's volatility and inadequacy in Chapter 3.2, it is also imperative to examine the theoretical mechanisms underlying such a dynamic. While historical data can show what has happened, a theoretical framework is better suited to explaining *why* it happens, allowing us to explore the strategic interactions that produce them in the first place. We present our earlier publication work, *Game-Theoretic Analysis of an Exclusively Transaction-Fee Reward Blockchain System* (Kruminis and Navaie, 2022), which showed how mempool congestion and block size limits can influence miner revenue and user fees, highlighting strategies that can enhance outcomes for both miners and users under fee-only network conditions.

4.1 Introduction: A Game-Theoretic View of the Fee Market

The initial idea of the paper was simple; increasing the block size allows more transactions to be included into a block, thereby potentially increasing the total fees collected. As there is more block space available, the competition between users would decrease, which in turn would also lower the fees that individual users would pay. All of this could lead to a symbiotic dynamic where miners benefit from a higher overall revenue resulting from more fee-paying transactions, while the users enjoy reduced fees and quicker confirmation times.

To further investigate the feasibility of this dynamic, we modelled transaction inclusion as a dynamic, time-sensitive game in which users compete for the limited space inside a block. This approach was novel in the sense that it moved beyond the static, single-shot auction perspective that was adopted by previous research attempts, for example Dimitri, 2019. Instead, we modelled user utility as a direct function of the time to confirmation, simply treating the fees as a means to an end. This framework more accurately reflects the operational reality of blockchains, as the transactions that fail to be included into the current block still remain in the mempool, continuing their competition for space in subsequent blocks. If the fee to be offered remains the same throughout subsequent block cycles, the property that matters is the time to inclusion; hence, user satisfaction is inherently linked to confirmation time, not the fee itself.

4.2 Modelling Framework and Assumptions

To maintain our focus on user-to-user interaction, the model assumed a single collective of miners in charge of block production, who behave honestly and strictly adhere to protocol guidelines. Additionally, our analysis in the paper abstracted away from mining costs, mining difficulties, and was not concerned with the

possibility of a *mining gap*. By and large, we presumed an idealised environment in which no attacks on the network are attempted, and all submitted transactions are valid. All transactions were also assumed to have a uniform fixed size of 500 bytes, which is representative of the average Bitcoin transaction size, giving a capacity of $\beta = 4{,}200$ transactions per a 2 MB block.

To formally analyse this competition among users, we defined three core components:

- **Transaction worth (ω_{tx})**: This represents the intrinsic, personal urgency a user assigned to their transaction. It was modelled as a randomly generated value between 0 (lowest priority), and 1 (highest priority). The ω_{tx} value ultimately dictates the user's bidding behaviour, and their tolerance for confirmation delay.

- **User Utility**: This function depends on how long it takes for a transaction to be included in a block, rather than the fee paid. It is formally defined as:

$$U_t = (1 - \omega_t)^{\frac{d-d_t}{\iota \times p}}, \tag{4.1}$$

  where $\omega_t \in [0, 1]$ is the worth of the transaction, d is the current time, d_t is the time the transaction was submitted, ι is the block interval (600 seconds in Bitcoin), and p is the patience of the user, fixed at $p = 2e$ for all users in the model. The function is exponentially decreasing with respect to both the elapsed time and the transaction worth; high-worth transactions lose utility rapidly if not confirmed promptly, whilst low-worth transactions can tolerate confirmation delays and still retain high utility levels.

- **Network State**: This variable captures the level of competition in the system, and is another component influencing a user's bidding behaviour. It is formally defined as:

$$c = \frac{\rho}{\frac{\beta}{\phi}} \tag{4.2}$$

where $\rho \in \mathbb{W}$ is the number of transactions currently in the mempool, $\beta \in \mathbb{W}$ is the block size in terms of how many transactions it can accommodate, and $\phi \in \mathbb{R}^+$ is the exchange rate of one Satoshi to pounds sterling. A higher network state value indicates greater mempool congestion, so consequently, users must increase their fee bids in accordance with this heightened competition for block space. Conversely, as the BTC price rises and ϕ increases, users are less willing to pay the same nominal fee in BTC, since its real-world cost has increased. The function therefore adjusts fee-paying behaviour accordingly, keeping the effective fiat cost stable.

The above components ultimately interact through the fee function, which determines whether a user is willing to pay the current minimum fee (m_i) required for block inclusion:

$$\text{if} \quad \left(\frac{c}{m_i}\right)^{\omega_{tx}} \geq \gamma, \quad \text{then} \quad f_{tx} = m_i, \tag{4.3}$$

where $\gamma \in \mathbb{R}^+$ is a calibrated threshold representing a user's inclination to pay the fee. If the fee function value meets or exceeds γ, the user submits a transaction with the required fee amount and thus joins the top of the queue for block inclusion; otherwise, they hold off, leaving their transaction for potential inclusion for another block cycle. In summary, this formulation captures the competitive, iterative nature of the block inclusion game - users incrementally raise their bids until the fee is just sufficient for inclusion within a block. But, if the conditions have deteriorated, and the prevailing fee amount has exceeded the worth they assigned to their transaction, they exit the game and re-evaluate their position under the next block cycle's conditions.

To anchor our parameters within a realistic operating regime, we draw upon a full year of empirical Bitcoin statistics, spanning from May 2020 to May 2021. Over this period, Bitcoin received approximately 3.525 transactions per second, or, around 2,115 transactions submitted per a 10-minute block cycle, of which roughly 2,141 were confirmed per block (3.57 tx/s). The mempool contained on

average 31,596 pending transactions, while a standard 2 MB block accommodates $\beta = 4{,}200$ transactions. The average BTC to GBP exchange rate over this period was £24,490 per BTC, corresponding to a per-Satoshi rate of $\phi = £0.00024490$. The average transaction fee was 0.00033984 BTC, whilst the median was 0.00016348 BTC. From these statistics, the core model parameters were derived as follows. Since only 4,200 of the average 31,596 mempool transactions were confirmed per block, the proportion of winners was approximately 13.3%, implying a representative transaction worth of $\omega_{\text{tx}} \approx 0.867$. Substituting the mempool size, block capacity, and per-Satoshi exchange rate into Equation 4.2 yielded a calibrated network state of $c = 30{,}718.077$. Finally, using the average winner fee $m_i = 0.00033984$ BTC alongside $\omega_{\text{tx}} = 0.867$ and $c = 30{,}718.077$ in Equation 4.3, the fee-paying probability threshold was set to $\gamma = 7{,}905{,}978.559$. These calibrated values collectively ensure that the model's simulated user behaviour is consistent with what was empirically observed in the Bitcoin network during this period.

4.3 Simulation Design and Implementation

In order to evaluate our model dynamics, we developed a publicly available game-theoretic simulation model (Kruminis, 2022), which captures the interaction between competing users striving to have their transactions confirmed. To aid our simulations, some network parameter values were informed by a historical analysis of Bitcoin statistical data, as per the above calculations. The model made assumptions by fixing transaction sizes, and also assumed that each user had perfect knowledge of all bids across the network. As such, the variance between the offered fees was uniform, with users incrementally increasing their bids until the fee was good enough for confirmation, or the current fee rate exceeded the transaction's assigned worth. The simulation ran for 100 rounds, and its procedure is summarised as pseudocode in Algorithm 1:

Input: Block size β; new transactions per round ρ_{in}; calibrated γ, ϕ; rounds T

Output: Per-round average fee, user utility, network utility, mempool size

Populate the mempool $\mathcal{M}$ with an initial batch of ρ_{in} transactions, each assigned a random worth $\omega_{tx} \in [0, 1]$;

for *each block cycle round* $t = 1, \ldots, T$ **do**

Add ρ_{in} new transactions to $\mathcal{M}$;

Compute network state c based on current mempool size, block size, and exchange rate;

Seed the winner list $\mathcal{W}$ with any transactions that lost last round but still have a competitive fee;

for *each remaining transaction* x *in the mempool* **do**

Determine whether x is willing to pay the current minimum fee m_i, based on its worth ω_{tx}, the network state c, and threshold γ;

if *block is not yet full* **then**

add x to $\mathcal{W}$;

else if x *outbids the lowest-paying winner* **then**

replace the lowest-paying winner with x; loser re-enters competition;

else x cannot outbid anyone — carry it over to the next block cycle round;

end

Confirm the block: remove all winners from $\mathcal{M}$;

Record average fee paid, average winner utility, average mempool utility, and mempool size;

end

Algorithm 1: Game-theoretic fee-market simulation

Block Size Limit	TX	Avg. Fee in BTC	Avg. Fee Total per Block
4200 (2MB)	15000	0.00000386	0.01619648
8400 (4MB)	30000	0.00000386	0.03240829
16800 (8MB)	60000	0.00000382	0.06410137
33600 (16MB)	120000	0.00000380	0.12761163

Table 4.1: The impact of different block sizes on fees paid when the mempool ratio is the same

4.4 Core Findings

The primary, and most significant finding from our game-theoretic model was that miners have no rational economic incentive to artificially constrain the block size in a fee-only environment. Our analysis demonstrated that the revenue derived from fees is determined not by the block size limit itself, but by the level of mempool congestion. This result stands in direct contrast to prior research attempts (e.g. Dimitri, 2019), which suggested that miners should leave some block space empty to drive up fees. Notably, their conclusions fall apart when block rewards are absent, further underscoring how existing research needs to be reconsidered when evaluating purely fee-driven blockchain systems.

The simulation runs produced several important results under varying parameter configurations:

- As shown in Table 4.1, when the mempool level is fixed, increasing the block size lowers the average fees paid by users, and in turn, the total revenue collected by miners.

- If users were to set their bids in line with their maximum valuation, the total fee revenue to be collected by miners is independent of the block size limit. Our simulations showed that as long as the block is full, smaller blocks just drive up the per-transaction fee, but do not necessarily increase overall revenue.

- Figure 4.1 illustrates that miner revenue is a function of demand, not scarcity; when block size is fixed, a bigger mempool intensifies the bidding competition between users and thus increases the total fees to be collected by miners.

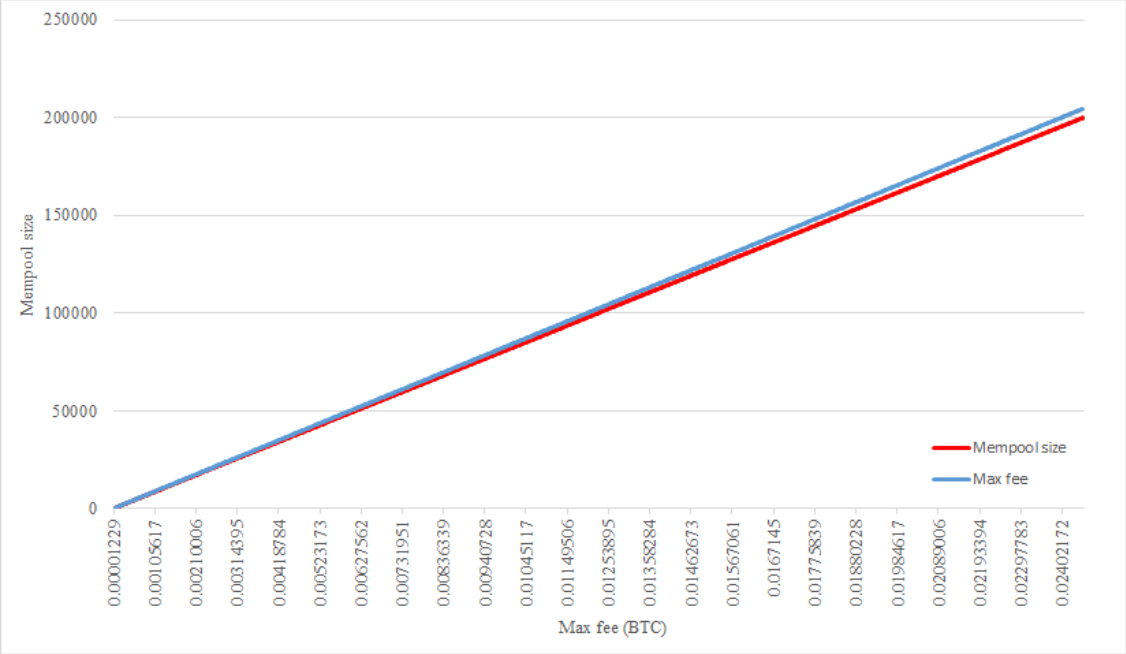

Figure 4.1: Maximum fees users are willing to pay, when linearly increasing the mempool levels

Therefore, the key finding is that a dynamic block size works to the benefit of both the users and the miners. As shown in Table 4.1, total revenue from fees can be maximised by capturing as much of the existing user demand as possible. By adjusting the block size limit in accordance with mempool congestion, we can lower the individual fees that users pay while simultaneously increasing the total fees collected by miners, producing a win-win dynamic.

4.5 Broader Implications for Fee-Only Blockchains

The game-theoretic analysis presented in this chapter dismantled the previously held consensus view that artificially limiting block space will increase miner revenue streams, a notion that simply does not seem to hold up under fee-only conditions. The collection of our overall results points to a clear, and actionable strategy for optimising mining revenue in fee-only blockchains: to dynamically adjust the block size, expanding or contracting in response to mempool congestion.

However, our work also heavily relied on several simplifying assumptions, ones that we will address extensively throughout the remainder of our thesis. Specifically, we disregarded any security challenges that become prevalent in a fee-only regime, a topic that we explore in detail in Chapter 6.4. Additionally, the costs of mining, and the overall profitability of mining were not evaluated in the paper, the sole focus being to attempt to maximise the revenue as much as possible; miner profitability in fee-only environments will be analysed comprehensively in Chapter 6.3.

In summary, the game-theoretic analysis conducted here provides a strong theoretical backdrop for the subsequent themes to be discussed in our thesis. It reconfirms that fee-based revenue is fundamentally dictated by unpredictable, stochastic arrival of user transactions and their resulting congestion. Although our model recommended a dynamic block size to help manage this congestion and thus improve fee collection when activity levels are high, it fails to resolve the core problem of revenue volatility when transaction demand is low. Given this inherent instability, we require a mechanism that extends beyond just block space management, and is able to actively moderate miners' income fluctuations. This challenge motivates the development of our novel Reserve Pool TFM, which we introduce and examine in detail in the following chapter.

Chapter 5

Simulation Setup and Methodology

Overview

This chapter presents the Zero-Block Reward Auction simulator, henceforth referred to as *ZBRA-TFM*, which will serve as the foundational experimental model for our evaluation of transaction fee mechanisms. We describe its architecture and operational logic in detail, justify our usage of real-world Bitcoin transaction data as input for the simulations, and explain our selection of core simulation parameters, emphasising their fairness and relevance for robust comparative analysis.

5.1 Input Datasets

The credibility of any simulation model is fundamentally dependent on the relevancy of its input data. Therefore, utilising a real-life dataset comprising of transactions executed by actual blockchain users is preferable to relying on artificially created data. By basing our simulations on empirical transaction records, our results will more accurately reflect how transaction fee mechanisms operate in practice.

Given the vast abundance of publicly accessible blockchain ledgers, each with its

own distinct architectures and data structures, selecting an appropriate blockchain for data extraction is no trivial task. However, Bitcoin, as the most popular blockchain network worldwide and almost synonymous with blockchain technology in general, emerges as an especially suitable candidate not only because of its widespread adoption but also due to its sheer volume of available data.

Since its launch in early 2009, Bitcoin has processed over 1.2 billion transactions, comprising of approximately 675 GB of on-chain data as of August 2025 (Blockchain.com, 2025). Whilst of course analysing the entire transactional history of Bitcoin would certainly be the most comprehensive approach, such an undertaking is also impaired by the blockchain's various evolutionary phases; protocol upgrades, hard fork shifts, fluctuations in its fiat valuation, block reward amounts, etc. Hence, shifting our focus towards a more bounded sample size from a period characterised by protocol maturity and operational stability is both more analytically robust and practical for our purposes.

Therefore, for our simulations, the elected dataset to be used contains all the confirmed Bitcoin transactions from the first block on June 1, 2023, through to the last block mined on June 30, 2023, amounting to a total of 12,150,245 transactions. This particular interval was selected due to its appropriateness in representing normal operational dynamics of an actively used blockchain network. Importantly, June 2023 falls midway between the May 2020 and April 2024 block reward halving events, thus ensuring that the dataset is unaffected by the fee and price volatility that typically surrounds halving periods. The month was therefore characterised by relatively typical and sustained network usage; neither unusually quiet nor particularly congested, allowing us to analyse the essential characteristics of a mature blockchain network. The transaction fees were also sustainable, which is of special importance to our examination of transaction fee mechanisms.

It should be noted that confirmed transactions represent only a subset of the full network activity. Some transactions may enter the mempool but remain unconfirmed for an extended period of time, whether it be due to its offered fee uncompetitiveness,

censorship, or other factors. Moreover, the mempool in Bitcoin is inherently decentralised, with each node maintaining its own version of the mempool. The set of transactions visible to one miner may vary depending on their experienced network latency, their connectivity to other peers, dropped packets, etc., making it difficult to obtain a universal count of transactions entering the mempool at any given time. Consequently, there is no definitive or globally consistent mempool state. Therefore, we are only able to concern ourselves with confirmed transactions.

Our dataset was constructed by first retrieving all the block hashes corresponding to blocks mined between June 1st and June 30th, 2023. For each such block, we extracted the transaction data and stored only relevant attributes, more specifically: the transaction hash, transaction size (in bytes), SegWit weight units, and the total fee paid by the transaction. To make use of storage efficiency and facilitate further data processing, the collected information was stored as JSON objects, enabling both rapid access and analytical flexibility. Nonetheless, we reaffirm again that the ZBRA-TFM model is generalised and can be made compatible with transaction data from any classical blockchain protocol.

5.2 Zero-Block Reward Auction Simulator: Architecture and Design

In order to evaluate the applicability and viability of our proposed Reserve Pool TFM (see Section 3.3), we developed a modular, discrete-event simulator called Zero-Block Reward Auction and Transaction Fee Mechanism simulator (ZBRA-TFM), which was made freely available at Kruminis, 2025. The simulator is written in Java and is specifically tailored for the analysis of transaction fee mechanisms within fee-only blockchain environments.

A high-level overview of the model's architecture is presented in Figure 5.1. The model uses a command-line interface where a user can select the parameters of the simulation process, taking as input a JSON-formatted transaction dataset

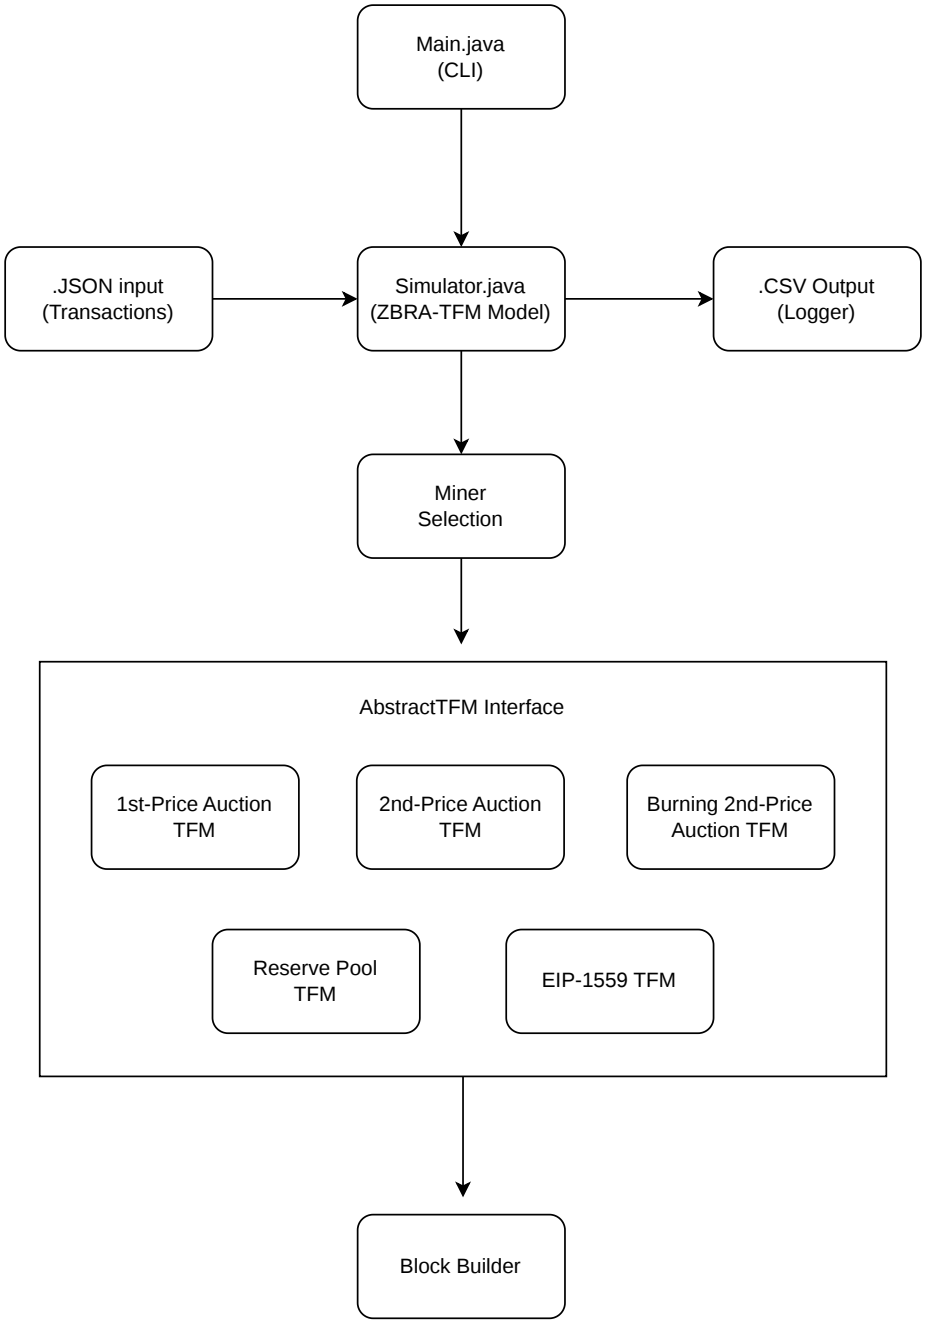

Figure 5.1: High-level system architecture diagram.

and outputting CSV log files: a detailed log file containing all the transactional information for a particular TFM, and a smaller summary log file which contains key information such as miner payouts and network utilisation. A pseudorandom seed is also used to determine factors such as how many transactions to add each block cycle, or which miner(s) get selected as winning the block payout. Core blockchain elements, such as transaction or block structure, were designed with flexibility in mind, making it easy to tailor them to specific use cases.

To ensure future compatibility, our model is easily extendable. Other transaction fee mechanisms can be incorporated into the simulator by extending the *Abstract-TFM* class and overriding the *fetchValidTX()* method, which specifies the logic by which the mechanism behaves. To adjust mining and/or consensus policies, such as transitioning to a permissioned state, or adopting alternative consensus schemes, the *getWinningMiner()* method can be modified in the *Simulation.java* class. Such strict separation and loose coupling of mechanism logic from the simulation core will ensure experimental fairness, so that only the transaction fee mechanism logic impacts the results of the simulation process. The operational flow of the simulator is summarised in the following pseudocode:

We begin by initialising the network state and mempool, then iteratively processing blocks according to the selected TFM ruleset. New transactions join the mempool sampled in a Gamma-Poisson distribution, which models more realistic, bursty network arrivals. Rewards and fees are calculated as per the TFM's logic, and the results are logged for further analysis. This design approach enables our simulator to conduct robust, reproducible, and extensible experiments for different transaction fee mechanisms, thereby enabling controlled comparative analysis under varied blockchain network parameter settings.

Data: Transaction dataset

Result: Logged blockchain blocks

Set initial state of the network according to parameters chosen;

Load transactions from JSON file;

Randomly select transactions to add to initial mempool state;

Create set of miners with assigned stakes;

for *each block cycle* **do**

Add more transactions into mempool, sampled from Gamma-Poisson distribution;

Select block-proposing miner(s) by stake-weighted random process;

Process mempool transactions according to selected TFM logic;

Calculate miner(s) reward according to selected TFM;

Calculate user fees according to selected TFM;

Assemble block and add it to the blockchain;

Update mempool state by removing confirmed transactions;

Payout reward to winning miner(s);

Update reserve pool balance if applicable;

Log block data and statistics to CSV files;

end

Log summary of results

Algorithm 2: ZBRA-TFM Simulator Data Flow

5.3 Simulation Parameters

Setting well-defined and justified parameter values is a critical step in ensuring accuracy and reliability in simulation-based research. The parameters for our model aim to reflect both the empirical characteristics observed in real-world blockchain systems and theoretically motivated values. Table 5.1 summarises the parameter values we used across our main simulation experiments, which include block size constraints, transaction throughput and structure, and initial system state.

Considering our use of real-world Bitcoin transactions as the dataset in our simulations (as defined in Section 5.1), we also closely base our parameter values on the statistical data observed from the Bitcoin network. Conducting analysis between different transaction fee mechanisms poses several issues, perhaps chief among them being the challenge of establishing uniform evaluation conditions. Each TFM operates under its own distinct set of rules (see Section 2.6), and while these differences are central to the comparison itself, meaningful analysis requires controlling for other contextual factors. By fixing some set of simulation parameters, such as network load, mempool state, and fee volatility, we ensure that any differences observed in performance can be attributed to the fee mechanisms themselves, and not to some variation in network and/or system conditions.

Concerning the block size, we define two distinct parameters: `SIZE_LIMIT` and `SIZE_TARGET`. After summing up the individual transaction sizes in our simulation, the `SIZE_LIMIT` refers to the hard upper bound on block size that must not be exceeded under any circumstances. Alternatively, `SIZE_TARGET` serves as a soft target value, to be used for dynamic fee adjustment in EIP-1559 and our Reserve Pool TFMs. When the actual block size exceeds the target size value, it signals that there is excess demand on the network, prompting the base fee to be increased for future blocks. Conversely, if the block size falls below the target, the base fee is reduced to encourage higher inclusion.

Bitcoin originally had a hard limit of 1 MB as its block size, however, after the introduction of SegWit (Bitcoin Developers, 2015), which separated signature data

from the transactions in order to save space, Bitcoin moved towards a more flexible model based on weight units (`WU`). This not only changed how the transaction data inside the block is counted, but also increased the theoretical block size limit to around 4 MB. However, on average, blocks at full `WU` capacity tend to be around 2 MB in practice. Accordingly, for our simulations, we have also set the `SIZE_LIMIT` = 4×10^6 `WU`, matching the current Bitcoin protocol limit. On the other hand, Ethereum defines their block size in terms of gas units, which measures the total computational costs used up by transactions inside the block. With the introduction of the EIP-1559 mechanism, the focus shifted from a hard limit towards a target size value which still remained at the old limit of 15,000,000 gas units, and changed the block size limit to be double its target. While this adjustment theoretically doubled the throughput, the protocol is designed to dynamically regulate usage back toward the target value.

To ensure fair and controlled comparison between the TFMs, we have decided to normalise the effective throughput across all models. More specifically, we will increase the limit (`SIZE_LIMIT` = 8×10^6 `WU`) for the EIP-1559 and Reserve Pool TFMs only, paired with a `SIZE_TARGET` = 4×10^6 `WU`. This configuration mirrors Ethereum's target/limit dynamic and addresses a confounding variable: one TFM outperforming another simply because it is allowed higher throughput due to a lower target size. This setup is also justified by historical network behaviour: while Bitcoin miners tend to include as many transactions as possible (up to the block size limit) to maximise fee revenue, in Ethereum, the average block size (Etherscan, 2025) has consistently remained close to its target value (i.e., 50% of the maximum size limit). Consequently, the effective average block size remains roughly uniform across all TFMs in our simulations, at 4×10^6 weight units (approximately 2 MB).

Next, in order to accurately simulate network congestion, we must define the network state parameters: `MEAN_TX_ARRIVAL_RATE`, which will reflect the average number of new transactions entering the mempool each block cycle, `ALPHA`, a dispersion parameter that will capture the variability in transaction arrival rates,

and MEMPOOL_INITIAL_SIZE, which will refer to the number of transactions already in the mempool at the start of our simulation.

Table 5.1: Simulator Parameters

Parameter	**Value**	**Description**
SIZE_LIMIT	8×10^6 WU *(EIP-1559, RP)*, 4×10^6 WU *(1PA, 2PA, B2PA)*	Maximum block size, expressed in megabytes (MB) or weight units (WU)
SIZE_TARGET	4×10^6 WU *(EIP-1559, RP)*	Target block size for dynamic fee adjustment
MEAN_TX_ARRIVAL_RATE	2,471.0	Avg. new transactions per block cycle
BASE_FEE	0.0000002333	Initial base fee per byte
ALPHA	0.01 – 0.03	Dispersion parameter for transaction arrival variance
RESERVE_POOL_BASE	0 – 134.38	Initial balance of the reserve pool
MEMPOOL_INITIAL_SIZE	93,824	Initial number of transactions in mempool

The dataset to be used in our simulations contains a total of 12,150,245 transactions; given Bitcoin's average block interval of approximately 10 minutes, this translates to an average of around 2,813 transactions per block (30 days × 144 blocks/day). As per statistical data from Blockchain.com, 2025, during the month of June 2023, the daily average number of transactions per second (TPS) added to the mempool ranged from 3.21 (25/06/2023) to 5.61 (07/06/2023), with an average of 4.12 TPS. Extrapolated, this would amount to between 1,925 and 3,366 transactions per block cycle. Therefore, we will set the MEAN_TX_ARRIVAL_RATE at 2,471 (4.12 TPS) transactions arriving per block cycle on average. Whilst this is below our

confirmation rate, it's important to remember the fact that a mature blockchain would already contain a backlog of previously submitted transactions. In fact, the Bitcoin blockchain went under severe congestion following the introduction of the "Ordinals" protocol in early 2023, which allowed for digital inscriptions on Bitcoin transactions.

These transactions weighed a lot more than simple payment transactions, with a single transaction taking up as much as 3.94 MB (mempool.space, 2023). As such, fewer traditional transactions could be fitted inside blocks, contributing to a significant backlog, reaching as far as 430,000+ transactions (Hoenicke, 2024) in the mempool in early May 2023. However, most nodes prune older transactions from the mempool; the default Bitcoin node settings limit the mempool size to 300 MB and discard transactions that are more than 14 days old. We will initialise the `MEMPOOL_INITIAL_SIZE` parameter to a value of 93,824 - the number of transactions observed in the mempool on June 1, 2023, as found in Blockchain.com, 2025. This will represent the initial transaction backlog present at the beginning of our simulations, and since we already deal with a period of just 30 days, no transactions will be discarded.

Transaction arrival rates can also be influenced by exogenous factors, such as time-of-day usage patterns (e.g., reduced activity during nighttime hours in primary user regions) and market events (e.g., dropping of new NFTs). As these factors are inherently variable and unpredictable, simulating these transient effects directly would overcomplicate comparisons between the different transaction fee mechanisms. Instead, we will model transaction arrival in the mempool using a Gamma-Poisson (negative binomial) distribution, which will capture the natural volatility with a single dispersion parameter, `ALPHA`. This approach will preserve realistic fee market dynamics without overfitting to some temporal idiosyncrasies.

For our main simulations, we will set `ALPHA = 0.02`, which should ensure that in 95% of block cycles, there will be between 1,800 and 3,400 new transactions entering the mempool. This value is closely aligned with our previous analysis of the state

of the Bitcoin network during the month of June 2023, where we found 1,925–3,366 transactions arriving per block cycle. A specific `ALPHA` parameter allows us to also conduct tests for high volatility, with `ALPHA = 0.03` and 1,500–3,900 transactions per block cycle, and for low volatility, with `ALPHA = 0.01` and between 2,100–2,900 new transactions per block cycle. In this way, by running simulations under varying stress conditions, we can arrive at a more robust assessment of the different fee mechanisms, supporting findings that will hold across a diverse set of operational states.

For the rest of our simulation parameters, we will once again rely on the historical values drawn from our transaction dataset. The average fee per weight unit observed was 23.33 satoshis, so we set `BASE_FEE` to 0.0000002333 BTC per weight unit. To capture the dynamics of transitioning to a reserve pool transaction fee mechanism, we can simulate scenarios where the reserve pool is initially empty, i.e., `RESERVE_POOL_BASE = 0`. Alternatively, we can also simulate with the reserve pool at a steady state, which we will assume to be the total capacity of a full day's worth of blocks; at the limit of 4 million weight units (and 144 blocks), and considering full blocks, this would amount to `RESERVE_POOL_BASE = 134.38`. This will allow us to observe system behaviour at both a steady state, and during a transitional phase of pool accumulation.

Chapter 6

Analysis of Results and Implications for Protocol Design

Overview

This chapter presents our experimental results and the evaluation of transaction fee mechanisms in fee-only blockchain environments. We begin first with a description of our experimental scenarios, outlining the simulation parameters and the methodological approach used to evaluate each TFM under identical network conditions. We then present comprehensive simulation results, which are the averaged-out performances of each TFM across 50 simulation runs of month-long blockchain activity. Next, we provide detailed analysis comparing our proposed Reserve Pool TFM against the other traditional mechanisms. The analysis will be based on several key performance metrics that evaluate miner payouts, user experience, and network efficiency to assess the Reserve Pool's potential as a viable TFM, and overall system-wide trade-offs. We conclude this chapter by examining the long-term viability of fee-only blockchain systems, exploring the broader societal implications of different TFM designs, identifying the practical challenges associated with transitioning away from block rewards, and finally providing actionable recommendations for future protocol designers aiming to

develop robust and sustainable fee-based incentive models.

6.1 Experimental Scenarios

Our experimental simulation entails the use of our novel ZBRA-TFM simulation model, where we simulate and critically assess the performance of five transaction fee mechanisms: 1st-Price Auction TFM, 2nd-Price Auction TFM, EIP-1559 TFM, Burning 2nd-Price Auction TFM, and our newly proposed Reserve Pool TFM. The core experimental design follows a controlled comparative approach, where each TFM is subjected to identical network conditions, allowing us to isolate the impact of mechanism design from random network quirks:

1. Identical initial network state conditions across all TFM comparisons,

2. Synchronized pseudorandom seeds for reproducibility,

3. Usage of real-world historical transaction dataset to ground the results in external validity.

All the simulations to be done are under fee-only blockchain conditions, meaning that there are no fixed block rewards for miners to collect; the only revenue to be gained is from transaction fees. This focus on fee-only conditions underscores the core objective of the thesis, as we aim to evaluate how blockchain systems will function in the future when block rewards have become negligible, and transaction fees remain the sole incentive for miner participation. Overall, we conducted 50 independent simulation runs, each initialised with a different pseudorandom seed to capture the natural variation in network dynamics, with all five TFMs evaluated under the same seed. The 50-run methodology ensures statistical robustness while maintaining computational feasibility, with each run consisting of 4,320 blocks, or 30 days of blockchain activity at 10-minute block cycle intervals.

Table 6.1: Final Simulation Parameters

TFM	**SIZE_LIMIT**	**SIZE_TARGET**	**RESERVE_POOL_BASE**
`1st-Price Auction`	4×10^6 WU	–	–
`2nd-Price Auction`	4×10^6 WU	–	–
`EIP-1559`	8×10^6 WU	4×10^6 WU	–
`Burning 2nd-Price Auction`	4×10^6 WU	–	–
`Reserve Pool Auction`	8×10^6 WU	4×10^6 WU	134.38

Our evaluation will focus on three main groups of performance metrics, that together should provide an accurate representation of the overall blockchain state:

- **Miner incentives**: average block payout, and the variability in payouts,
- **User experience**: average transaction fees, and fee variability,
- **Network efficiency and stability**: average block size, block size variance, and the block fill ratio.

Each simulation interacts with the same 30-day Bitcoin transaction dataset from June 2023 (refer to Section 5.1), and encompasses 12,150,245 real-world transactions. The experimental network parameters that we used can be found in Table 6.1, with `MEAN_TX_ARRIVAL_RATE = 2,471`, `BASE_FEE` $= 2.333 \times 10^{-}7$ `WU`, `ALPHA = 0.02`, and `MEMPOOL_INITIAL_SIZE = 93,824` remaining constant across all the TFMs respectively, justification for which we already established in Section 5.3. To ensure fair comparison, we normalised the effective throughput across the TFMs. The EIP-1559 and Reserve Pool mechanisms operate with expanded block limits *(*8×10^6 *WU)*, paired with lower target sizes *(*4×10^6 *WU)*, reflecting their dynamic fee adjustment capabilities, whilst the traditional auction mechanisms operate with fixed limits *(*4×10^6 *WU)*. Although that means that they are theoretically capable

of double the throughput, in practice, the mechanism should quickly adjust to 50% of the defined threshold, effectively holding the same throughput on average as the other traditional mechanisms.

6.2 Results

Table 6.2: Simulation Results by TFM Type (Averages)

TFM Type	**Avg. Block Payout**	**Variance Between Block Payout**	**Avg. TX Fee**	**Variance Between TX Fees**	**Avg. Block Size (MB)**	**Block Fill Ratio (%)**	**Block Size Variance**
1st Price Auction	0.2886	0.05497	0.0001162	0.00001688	1.5235	87.99	214764
2nd Price Auction	0.0529	0.05093	0.0001162	0.00001694	1.5243	88.03	212801
Burning 2nd Price Auction	0.1314	0.02022	0.0000784	0.00001322	1.6723	99.79	159836
EIP-1559	0.2761	0.06078	0.0001168	0.00001755	1.5219	43.94	387803
Reserve Pool	0.2371	0.01388	0.0001172	0.00001828	1.5241	44.00	410023

The results from our multiple averaged-out independent runs of month-long blockchain activity reveal notable performance differences between the simulated transaction fee mechanisms under fee-only conditions. We present our final results in Table 6.2. Each transaction fee mechanism exhibits its own distinct strengths and weaknesses when simulated in a practical blockchain environment, with implications that extend beyond just theoretical evaluations, bearing directly onto real-world

economic viability and long-term network sustainability. The performance metrics found in Table 6.2 can be roughly grouped into 3 main operational dimensions.

Miner Incentives: In the crucial domain of miner compensation, our newly proposed Reserve Pool mechanism emerged as the clear leader, achieving the most effective balance between block payout and payout variance, two values essential for sustaining mining operations and network security. The Reserve Pool TFM showcased the overall lowest payout variance (0.01388), while also maintaining competitive average block payouts (0.2371 BTC). It yielded a 77.2% reduction in variance relative to the EIP-1559 mechanism which had the highest variance overall (0.06078), and a 74.8% reduction relative to the 1st-Price Auction mechanism. The 1st-Price Auction did achieve the highest average payout (0.2886 BTC), however, it also experienced one of the highest levels of variance (0.05497) in our simulations. Most critically, our simulations showed economically undesirable results from the 2nd-Price Auction TFM, generating block payouts of only 0.0529 BTC, which is only 22.3% of the average reserve pool payout. The Burning 2nd-Price Auction also showed unfavourable results, and although its variance control was reasonable (0.2022), its miner compensation was substandard at only 0.1314 BTC on average.

User Experience: On the other hand, from the user perspective, the Burning 2nd-Price Auction TFM provided the most favourable fee structure, achieving the lowest average transaction costs (0.0000784 BTC), combined with minimal fee variance (0.00001322). Both the traditional auction mechanisms delivered identical average transaction fees of 0.0001162 BTC, and similar variance levels; 0.00001688 for 1st-Price Auction TFM, and 0.00001694 for the 2nd-Price Auction TFM. The two dynamic fee mechanisms, EIP-1559 and the Reserve Pool TFM, imposed only marginally higher user costs, 0.0001168 BTC and 0.0001172 BTC respectively, with slightly elevated variance of 0.00001755 and 0.00001828 respectively. However, in practical terms, these discrepancies

amount to only minimal difference in real-world cost for users, with an average transaction fee in 1st-Price Auctions amounting to £2.54, and £2.56 for the Reserve Pool TFM, a margin of only 2 cents (PoundSterlingLive, 2023).

Network Efficiency: Perhaps the most clear differences between the transaction fee mechanisms can be seen through their network state metrics. As justified in Section 5.3, the block size limit and the target limit values were adjusted to normalise the effective throughput across the models; such an approach can be seen as validated as the average block sizes in MB remained effectively equivalent across all the TFMs; averaging 1.5239 MB for the 1st-Price and 2nd-Price Auctions, and 1.5230 MB for the EIP-1559 and Reserve Pool TFMs. The Burning 2nd-Price Auction was the sole outlier, attaining a near-optimal block utilisation rate of 99.79%, coupled with the smallest recorded variance of 159836, and an average block size of 1.6723 MB. However, the block fill ratio does not tell the full picture, and indeed it masked the mechanism's fundamental flaw of incorporating unconfirmed transactions, in practice cutting the effective network throughput in half. In the same vein, although the EIP-1559 and Reserve Pool TFMs only reached 43.94% and 44% block utilisation respectively, it is important to note that their block limits were doubled because of the dynamic fee adjustment approach. As evidenced by their average block sizes, their effective throughput is comparable to that of the traditional auction mechanisms. The significantly higher block size variance of the EIP-1559 and Reserve Pool mechanisms, 387803 and 410023 respectively, versus 214764 in 1st-Price and 212801 in 2nd-Price Auctions, also represented a design strength rather than a weakness, showcasing their ability to dynamically respond to network demand fluctuations, competently accommodating transaction volume spikes.

Additionally, in order to further validate our simulation model's viability, we compared our ZBRA-TFM simulated 1st-Price Auction results against historical Bitcoin network data from June 2023, the same time period as our dataset. Using

statistical data from BitInfoCharts, 2023 and Blockchain.com, 2025, the average block size for the month of June 2023 was found to be 1.7185 MB, compared to 1.5235 MB in simulation, a modest -11.3% gap. The average transaction fee was around 0.00122 BTC, closely matching the 0.001162 BTC value in our simulations, a difference of only -4.8% that confirms that our simulated fee market realistically mimics bidding behaviour and transaction inclusion. Finally, miners received on average 0.3231 BTC per block just in fees, versus 0.2886 BTC observed in our simulation runs, a -10.7% shortfall. The close alignment between the historical and simulated values shows that ZBRA-TFM provides conservative yet reliable projections of Bitcoin's core fee market dynamics, thereby proving to be capable of effectively simulating the blockchain fee ecosystem.

6.3 Model Analysis

The simulation results above highlight clear trade-offs between miner revenue stability, user fees, and network utilisation across our five simulated TFMs. Figure 6.1 provides a high-level comparison of the results, showing that the Reserve Pool mechanism achieves the lowest payout variance by a considerable margin, with only a marginal decrease in average block revenue for miners. In contrast, the Burning 2nd-Price Auction provides the best user experience with reduced costs and easily predictable fee amounts, however these come at the expense of network throughput and miner sustainability. Unsurprisingly, the two dynamically adjustable fee protocols, EIP-1559 and the Reserve Pool TFM, most effectively maintained network efficiency and stability by aptly adapting to the fluctuations in transaction volume.

Figure 6.1: Normalized Averages By TFM Type

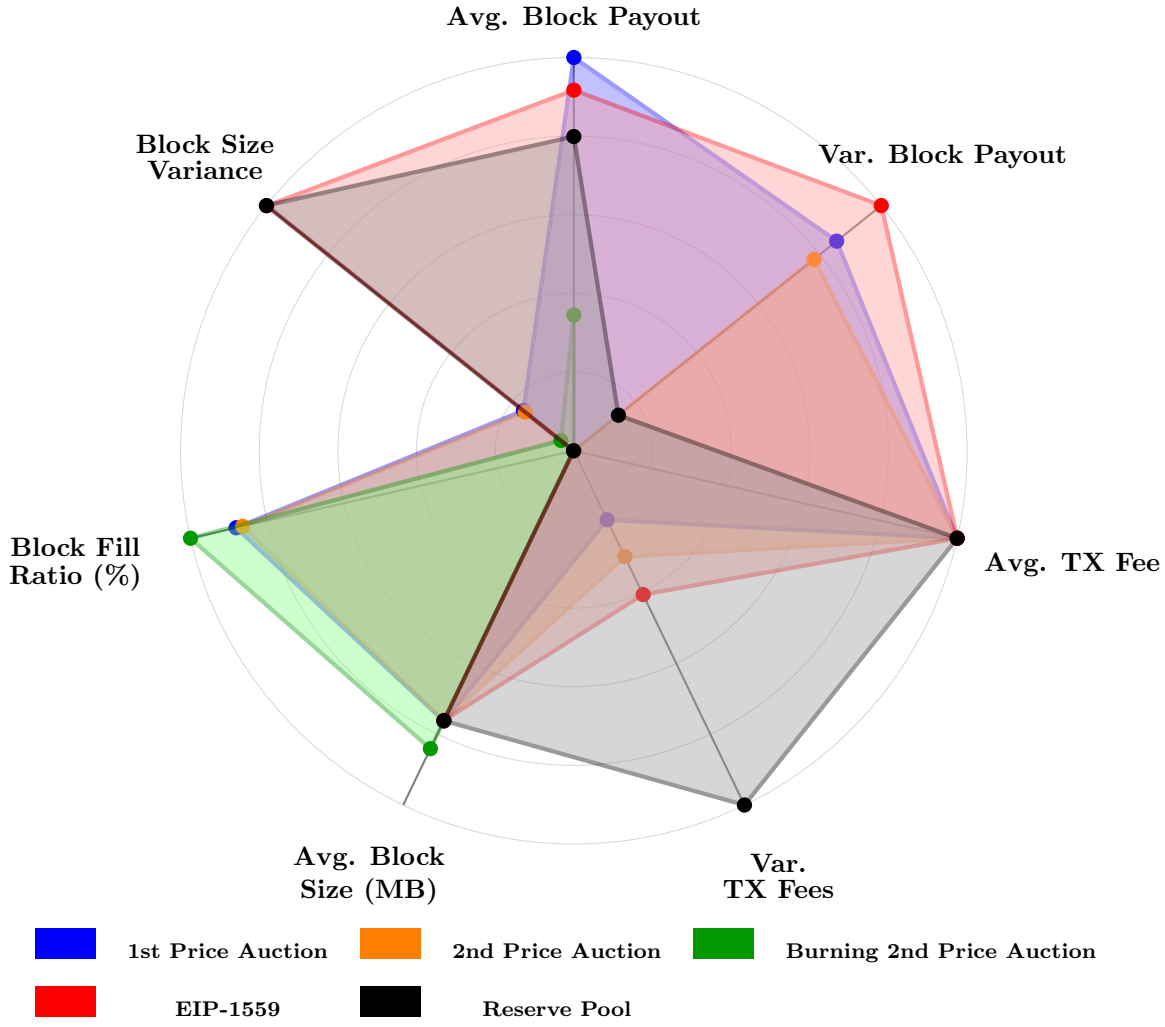

In providing useful analysis of our results, it is also important to view them in a more practical context; while the Reserve Pool TFM yielded the highest average transaction fees at £2.56, it was mere cents more expensive compared to the EIP-1559 protocol or the 1st and 2nd price auctions. Such minor differentials in some performance metrics underscore the importance of evaluating the entire dynamics of the blockchain experience; for example, the Reserve Pool's slightly higher user fees can be seen as effectively functioning as an insurance premium for increased network stability and security.

Table 6.3: Block payouts by TFM type (1 BTC = £21,844.63, June 2023 average).

TFM Type	**Min. Block Payout**	**Avg. Block Payout**	**Max. Block Payout**
1st-Price Auction	0.0174 BTC (£380)	0.2886 BTC (£6,304)	1.599 BTC (£34,930)
2nd-Price Auction	0 BTC (£0)	0.0529 BTC (£1,156)	0.9402 BTC (£20,538)
Burning 2nd-Price Auction	0.0803 BTC (£1,754)	0.1314 BTC (£2,870)	0.559 BTC (£12,211)
EIP-1559	0.016 BTC (£350)	0.276 BTC (£6,029)	1.004 BTC (£21,932)
Reserve Pool	0.139 BTC (£3,036)	0.2371 BTC (£5,179)	0.437 BTC (£9,546)

Perhaps the most clearly evident advantage of the Reserve Pool TFM over the other mechanisms is its significantly higher stability of block payouts to miners. As per Table 6.3, payouts to miners under the 1st-Price Auction TFM ranged from as little as £380 to nearly £35,000, £350 to almost £22,000 under the EIP-1559, and from £0 to £20,538 with the 2nd-Price Auction; meanwhile, under the Reserve Pool TFM, payouts remained relatively stable throughout, £3,036 to £9,546, with an average payout of £5,179.

The Reserve Pool payouts are particularly stable because they are computed over a 144-block average. As a result, a sudden single zero-transaction block can only reduce the reported average by no more than 1/144 of its prior value (approximately 0.69%), with k consecutive zero-payout blocks reducing it by at most $k/144$ percent. This averaging mechanism mitigates short-term volatility, something that the other mechanisms are especially vulnerable with, and enhances the predictability of miners' revenue streams. Such stability is necessary for maintaining network security; substantial fluctuations in payouts can nudge miners towards more opportunistic and strategic behaviours like *undercutting* or *fee-sniping*, which, while would be in the best interests of the individual miner, would have

detrimental effects on the blockchain network as a whole. In extreme cases, it may even lead to a situation called a *mining gap*, in which miners temporarily withdraw from the network entirely due to unprofitable conditions. With the absence of fixed block rewards to fallback on, payouts from just transaction fees must remain sufficiently stable - something that the Reserve Pool design particularly excels at.

The revenue derived from fees represents just a partial perspective of blockchain economics, as it must also be contextualised within the broader framework of overall miner profitability. Although higher payouts would, of course, typically correlate with increased miner satisfaction, what truly matters is the balance between revenue and mining costs, as miners have no economic incentive to participate when operations would result in net financial loss. Mining costs fundamentally depend on the type of consensus scheme used by the blockchain; for example, in *PoS*, the "mining" expenses for validators are mostly financial (e.g. opportunity cost of staked capital) and operational (e.g. bandwidth/storage requirements), which are marginal compared to the capital-intensive requirements of PoW systems that require substantial investments in both specialised hardware and electricity consumption. Given that PoW remains the predominant consensus mechanism for most major blockchain networks, and is the basis of Bitcoin, whose dataset we used in our simulations, we will analyse miner profitability specifically within the PoW paradigm. In order to conduct this analysis, we cannot simply rely on existing models or apply historical Bitcoin network data; such figures cannot be directly mapped onto fee-only models without accounting for additional revenue sources such as block rewards. Bitcoin miners continue to receive substantial block rewards which still constitute for the majority of total miner compensation. As demonstrated in Figure 3.2, fees accounted for just 7% of total miner revenue during 2023, with largest share of revenue coming in the form of block rewards, which during the month of June 2023 were 6.25 BTC per block.

According to Blockchain.com, 2025 data, the Bitcoin network maintained a total hashrate of 369,000,000 TH/s during June 2023, with network difficulty at

approximately 51.8 trillion. To model a large-scale mining operation, we assume an individual miner controlling 1% of the global network hashrate, which would correspond to around 3,690,000 TH/s. Maintaining profitability at this scale requires the use of the most energy-efficient hardware available; the Antminer S19 Pro, with an efficiency of 29.55 J/TH (Hashrate Index, 2025), represents state-of-the-art equipment that would have been available to miners during June 2023.

While a comprehensive economic analysis would incorporate both operational and capital expenditures, our evaluation will focus solely on operational costs, assuming that our hypothetical miner has already absorbed the rig acquisition costs. The total power consumption costs of this scale can be calculated using the following formula:

$$P_{\text{total}} = \text{Miner's Hash Rate} \times \text{Hardware Efficiency} \tag{6.1}$$

Inputting our data, the result is the following;

$$P_{\text{total}} = 3{,}690{,}000 \text{ TH/s} \times 29.55 \text{ W/TH} = 109{,}039{,}500 \text{ W} \tag{6.2}$$

The resulting power requirement of 109 MW represents a substantial energy demand, equivalent to that of a small city. Such a large scale of power consumption moves the operation into the scope of industrial-scale energy infrastructure, far exceeding residential grid capacity. As such, large-scale miners typically establish their facilities in locations with competitive electricity costs, and additionally negotiate Power Purchase Agreements (PPAs) directly with energy producers to secure prices even further below retail market levels. Based on these practices, we can reasonably estimate electricity costs to be approximately around £0.05/kWh for such operations. First, we extend our power consumption costs at a per-day rate: $109{,}039.5\,\text{kW} \times 24\,\text{h} = 2{,}616{,}948\,\text{kW h}$. With this, we can get the daily electricity costs with the following formula:

$$C_{\text{daily}} = \text{Daily Energy Consumption (kWh)} \times \text{Electricity Rate (£/kWh)} \tag{6.3}$$

Therefore, for a miner with 1% of the global network hashrate, we can estimate their daily mining operational costs to be:

$$C_{\text{daily}} = 2{,}616{,}948 \text{ kWh} \times 0.05 \text{ £/kWh} = \text{£}130{,}847.40 \tag{6.4}$$

Supposing a 10-minute block interval, just like in Bitcoin, there are 144 blocks mined per day ((6 blocks/h $\times$ 24 h = 144 blocks/day)). Thus, the estimated mining costs for a miner controlling 1% of the Bitcoin network during June 2023 were approximately £908.66 per block.

On the other hand, the expected revenue of a miner for a given block can be modelled as:

$$\frac{\text{Hashrate of miner}}{\text{Total network hashrate}} \times \text{Total fees in block} \tag{6.5}$$

Basing it on the average block payout values, the expected fee-only revenue is only £63.04 per block under a 1st-Price Auction TFM, and £51.79 per block under the Reserve Pool TFM. Clearly, directly using historical data whilst also ignoring the prevalent effects a block reward has had on the whole network does not provide for valuable insights; mining would seem to be inherently unprofitable, incurring losses of £845.62 per block under the 1st-Price Auction TFM. Ultimately, incorporating the block reward into the payouts brings up the expected reward for the 1% miner to £1,428.33 per block, or £519.67 expected profit per block in 1st-Price Auctions.

A neutral framework for assessing mining expenses in a fee-only context would treat the network hashrate as a dynamic variable that reaches an equilibrium balancing revenue against costs. Conceptually, the global coalition of miners can be thought of as a single rational entity that will expend resources only when the expected revenue at least covers the anticipated costs. If revenue from transaction fees declines, mining expenditures must also scale down; conversely, increased fee revenue intensifies competition between miners, and as such would result in higher operational spending. Under this equilibrium assumption, the projected network hashrate can be calculated with the following formula:

$$\text{Equilibrium Hashrate} = \frac{\text{Total Revenue per Second}}{\text{Cost per TH/s}} \tag{6.6}$$

Using average payout figures, we can compute the equilibrium hashrate for our simulated TFMs; the summary of these results appears in Table 6.4. For our calculations we used the same state-of-the-art mining rig, Antminer S19 PRO, which had an efficiency of 29.55 J/TH (Hashrate Index, 2025), and based our electricity costs on a rate of £0.05/kWh. As expected, TFMs with higher per block average revenue were able to command a more secure network hashrate. However, as mining hardware continues to improve in terms of computational power and energy efficiency, coupled with economies of scale that will reduce operational spendings, it is reasonable to assume that the overall network hash rate would increase over time too. An increased hash rate enhances the security and robustness of a blockchain network, therefore, bolstering user confidence and participation in it. Consequently, this influx of new users leads to higher transaction volumes, which in turn would drive the fees up, increasing miner revenue. This positive feedback loop demonstrates how technological progress, enhanced network security, increased user adoption, and economic incentives reinforce one another in maintaining thriving and sustaining blockchain ecosystems.

Once again, our results prove that the proposed Reserve Pool TFM is exceptionally successful in making fee-only blockchains stable; it can hold relatively competitive levels of network hashrate, whilst also maintaining an exceptionally high level of probability for miners. The probability rate in Table 6.4 refers to the percentage of blocks from our simulations whose payouts were at a break-even or profitable level for miners. Whilst other TFMs hovered at around a ≈50% mark, with the Burning 2nd-Price Auction being second highest at just 52.79%, the Reserve Pool TFM held an astonishing 97.08% profitability rate. This means that under the Reserve Pool mechanism, mining continuously remains profitable, even if the network experiences fluctuations in fee payouts or transaction volumes; an ability that further solidifies the security of the network and makes mining attractive for

miners, inviting them to join the system and in turn expending their resources.

Table 6.4: Simulated Miner Profitability (1 BTC = £21,844.63)

TFM Type	**Avg. per Block Revenue**	**Equilibrium Hashrate**	**Avg. per Block Mining Costs**	**Profitability Rate**
1st-Price Auction	£6,304	25,600,000 TH/s	£6,304	47.42%
2nd-Price Auction	£1,156	4,694,416 TH/s	£1,156	37.35%
Burning 2nd-Price Auction	£2,870	11,654,822 TH/s	£2,870	52.79%
EIP-1559	£6,029	24,483,249 TH/s	£6,029	50.89%
Reserve Pool	£5,179	21,031,472 TH/s	£5,179	97.08%

Alternatively, mining was only profitable or just break-even in costs and projected revenue 47.42% of the time for 1st-Price Auctions. When profitability or break-even conditions occur infrequently, miners may adopt strategies that deviate from ideal miner behaviour, such as undercutting previously mined blocks or holding off from mining entirely if there is no profitable revenue to collect. This slowdown in mining activity can have the effect of prolonging the transaction confirmation times (until block difficulty adjustments restore it), further degrading the user experience. Unreliable payouts may also push away smaller miners out of the network, as the inconsistent rewards make it more difficult for them to cover their operational costs. Under such circumstances, they may be compelled into joining larger mining pools as the only viable means of securing a steady and proportionate share of rewards; however, this shift risks weakening the decentralisation property that underpins the resilience of blockchain systems. All such deviations enhance the success probability

rates of attacks on the network, potentially deterring users from participating in the system if, for example, double-spending attacks can be achieved with a relatively high success rate, eroding overall trust in the blockchains fundamental security guarantees, and ultimately jeopardising its long-term sustainability.

6.4 Security Analysis

A transition to a fee-only model represents a significant yet little understood challenge to the long-term viability of many established blockchains. While we addressed the broader economic viability of this paradigm in Section 6.3, the direct implications for network security still require dedicated, rigorous analysis. In this section, we will evaluate the security posture of our proposed Reserve Pool TFM by comparing its inherent revenue stability with the volatility observed in other TFMs.

At its foundational level, the security of a blockchain system relies on the critical assumption that the economic incentives governing its network participants, the users and miners/validators, continuously remain aligned with the preservation of network health and integrity. Historically, substantial and predictable block rewards have served as the cornerstone for this, consistently ensuring that honest block production yields guaranteed, and economically attractive compensation. However, as these subsidies progressively decline, and overtime will be eliminated altogether, the network's security budget will become entirely dependent on transactions fees, which are inherently stochastic and volatile. Such a paradigm shift fundamentally restructures miner incentives, establishing conditions where rational, profit-maximising strategies may diverge from honest protocol compliance. This construct was first comprehensively investigated by Carlsten et al., 2016, where they posited that high variance in block payouts is a fundamental catalyst for consensus instability. Their work pioneered several critical research domains that have become particularly prevalent in fee-only blockchain systems:

1. They introduced the concept of a **mining gap**, where, without the guarantees

of a fixed and consistent block reward, *"immediately after a block is found there is zero expected reward for mining but nonzero electricity cost, making it unprofitable for any miner to mine"*. Consequently, they would rationally wait until enough fee-paying transactions arrive in the mempool to make mining profitable again.

2. **Strategic undercutting**, where a miner may rationally decide to fork the longest chain while deliberately leaving some transaction fees behind for other miners. In turn, this makes their fork the more economically attractive option for other miners to mine upon, fundamentally breaking the blockchain's status-quo of extending upon the longest valid chain.

Their research work found that multiple attack vectors have enhanced success probability rates in a fee-only regime, primarily due to the increased volatility of such an environment. They identified the block reward as an integral part of maintaining stability in mining, proposing to make it a permanent feature. Instead, we contend that the Reserve Pool TFM is a viable alternative, capable of sustaining stability within fee-only blockchain ecosystems. Our simulation results also provide strong empirical support for this hypothesis. As shown in Table 6.2, the block payout variance for the Reserve Pool TFM was remarkably low, at just 0.01388, whereas comparatively, the more traditional mechanisms had significantly higher variance; 1st-Price Auctions displaying a variance of 0.05497, and EIP-1559 showing variance of 0.06078, approximately 4.4 times higher than the Reserve Pool TFM. These findings are further corroborated by the data in Table 6.3, which showed drastic disparities in block payout ranges; from as little as £350 to peaks approaching nearly £35,000 in 1st-Price Auctions, whereas the Reserve Pool TFM exhibited much moderate swings, from the lowest of £3,036 to £9,546. Whilst high variance in miner payouts is not in itself a direct attack vector, it serves as the underlying condition from which multiple, distinct behavioural patterns can emerge, each of which has the potential to undermine network integrity:

1. **Survival-Driven:** This encompasses miner behaviour that is motivated by fundamental economic survival, wherein blocks with unprofitable or even near-zero rewards force the miner to abandon honest following of the protocol. As the expected revenue falls below their operational costs, miners are rationally incentivised to explore alternative strategies in an attempt to restore profitability, manifesting in phenomena such as *'mining gaps'*, and strategic attacks like *'undercutting'*.

2. **Greed-Driven:** This category characterises miner behaviour in which they deliberately target *'wealthy'* blocks containing exceptionally high rewards, a direct consequence of the high variance of miner payouts. The existence of such high-reward outliers creates powerful incentives for miners to act selfishly and intentionally fork the blockchain in an attempt to steal lucrative blocks from competing miners, constituting attacks such as *'fee sniping'*.

The results presented in Table 6.4 provide a good basis for assessing survival-driven mining behaviour. Under the 1st-Price Auction TFM, 52.58% of block cycles would have exhibited conditions in which a strategic miner could have rationally considered undercutting previously mined blocks, or ceasing mining operations entirely. Comparable rates could be observed in other TFMs, with EIP-1559, for example, reporting 49.11% of blocks as unprofitable at the equilibrium hashrate. By contrast, our proposed Reserve Pool TFM demonstrated exceptional profitability or break-even rates, with only 2.92% of blocks deemed unprofitable over a 30-day simulation period. This performance can be largely attributed to the presence of the smoothed payout mechanism, which ensures that the utility gained from honest behaviour is easily predictable and almost always positive. Thus, any deviation from fairly following the protocol guidelines becomes an economically irrational decision. This combination of high payout stability and low variance between block payouts makes attacks such as *undercutting*, or the emergence of a *mining gap* exceedingly unlikely under the Reserve Pool TFM. This resilience is even further reinforced by the substantial mining power that would be required to make survival-driven

behaviour economically viability for a strategic miner.

Whilst survival-driven attacks are driven by desperation, the existence of exceptionally high-reward blocks due to payout volatility drives attacks of greed. *Fee sniping* represents a specialised variant of undercutting, wherein the primary objective is to capture revenue from an unusually profitable block that has just been mined by a competing miner. Here, an attacker may find that the more strategic behaviour is to attempt to fork the chain, and in turn collect the fees from an already mined high-payout block for themselves, rather than honestly mining the next, likely average-value block. The most basic countermeasure against fee sniping would be to constrain the extreme volatility of block payout ranges. Further analysis of the data from Table 6.3 reiterates this sentiment:

- In 1st-Price Auctions, block payouts ranged from as little as £380 to £34,930, and with such an extreme range, *fee sniping* could become a particularly compelling move. Compared to the average payout of £6,304, the maximum payout was over 5.5 times the average, and over 91.2 times the minimum of £380.

- The EIP-1559 protocol showed swings of 3.63 relative to the average (£6,029), and over 62.7 times from the minimum (£350).

- Comparatively, payouts from the Reserve Pool TFM were significantly tighter constrained. With an average of £5,179, the maximum (£9,546) was only 1.8 times from the mean, and 3.1 times from the minimum (£3,036). There were simply no extreme outliers that could have made *fee sniping* a strategically attractive move.

What makes the Reserve Pool TFM particularly secure from this is the rolling window of payouts, meaning that any single block with an outlier payout will only minimally impact the final block payout. With this, we can reasonably assert that miners operating under the Reserve Pool TFM can be assured that no single block

will contain a lottery-like prize, which effectively neutralises the attack, because its very own prerequisite, a high-value target, is specifically designed out of the system.

Nonetheless, a comprehensive security analysis still needs to evaluate the threshold at which such attacks become rational for a miner, regardless of how attractive the potential reward might appear. In the most simple terms, a rational miner would only ever attempt an attack if the potential reward to be gained from it, adjusted for its probability of success, is greater than the reward from just mining honestly. This gives us two separate utility functions:

- **Expected utility of Honest Mining** ($\mathrm{U}_{\mathrm{honest}}$): This is the payout a miner can expect from following the longest chain, and mining upon it, typically represented as their share of the networks hash power multiplied by the historically average payout.

- **Expected utility of Strategic Behaviour** ($\mathrm{U}_{\mathrm{strategic}}$): This represents the anticipated utility a miner would have by deviating from honest mining behaviour, and more specifically when targeting a high-value block. It is normally calculated as their share of the network hashrate, justified against the probability of succeeding in such an attack, multiplied by the targeted payout amount.

In simple terms, a strategic deviation for a miner is rational when:

$$\mathrm{U}_{\mathrm{strategic}} \; > \; \mathrm{U}_{\mathrm{honest}} \tag{6.7}$$

As previously discussed, the expected utility of honest behaviour can be given by a miner's proportional share x of the total network hashrate, multiplied by R_{average}, the expected reward for the next block, which we will infer from our simulation records from Table 6.3. Therefore:

$$\mathrm{U}_{\mathrm{honest}} \; = \; x \; \times \; R_{\mathrm{average}} \tag{6.8}$$

Modelling the expected utility of strategic behaviour is a particularly complex task. To narrow down our approach, we will focus specifically on fee sniping, which

to recap, involves a strategic miner forking the blockchain and attempting to mine an already published block in an attempt to steal away the high-valued collection of transactions within it. Following the framework established in Bitcoin Optech, 2024, we define the expected utility of *fee sniping* as:

$$\mathrm{U}_{\text{strategic}} \;=\; \left(\frac{x}{1-x}\right)^{2} \times R_{\text{target}} \tag{6.9}$$

This formula links the probability of successfully forking the chain with the potential revenue that could be received from the targeted block. Ultimately, a rational strategic miner would only ever attempt to fee snipe a block whenever:

$$\left(\frac{x}{1-x}\right)^{2} \times R_{\text{target}} > x \;\times\; R_{\text{average}} \tag{6.10}$$

Using formula 6.10, we can determine the minimum share of overall network hashrate a miner would require at which fee sniping would become a viable strategy. We draw once again from the simulation data in Table 6.3 for our calculations. For a 1st-Price Auction, taking in the average payout value, we have $R_{\text{average}} = £6,304$, and using the maximum observed value of £34,930 as the target payout R_{target}, our formula becomes as follows;

$$\left(\frac{x}{1-x}\right)^{2} \times £34,930 > x \;\times\; £6,304 \tag{6.11}$$

In turn, we must solve for the smallest $x \in (0,1)$ that satisfies:

$$\left(\frac{x}{1-x}\right)^{2} \cdot 34930 \;>\; 6304\,x.$$

Multiplying both sides by $(1-x)^2$ gives:

$$34930\,x^2 \;>\; 6304\,x(1-x)^2.$$

Expanding the terms on the right side yields:

$$34930\,x^2 \;>\; 6304x - 12608x^2 + 6304x^3.$$

Now, we rearrange the terms to one side of the inequality:

$$0 \ > \ 6304x^3 - (12608 + 34930)x^2 + 6304x.$$

$$0 \ > \ 6304x^3 - 47538x^2 + 6304x.$$

Factoring out x from the expression:

$$0 \ > \ x(6304x^2 - 47538x + 6304).$$

Since we are solving for $x \in (0, 1)$, x is positive. We can divide by x and flip the inequality to make the leading term positive:

$$6304x^2 - 47538x + 6304 < 0.$$

Solving for the roots of the corresponding quadratic equation using the quadratic formula:

$$x = \frac{47538 \pm \sqrt{(-47538)^2 - 4(6304)(6304)}}{2(6304)}.$$

Calculating the smaller root gives the minimum value:

$$x_{\min} \approx \frac{47538 - \sqrt{2100903220}}{12608}.$$

$$x_{\min} \approx \frac{47538 - 45835.61}{12608}.$$

$$x_{\min} \approx 0.1350.$$

Therefore, the minimum value x satisfying our inequality is 13.5%; for a fee-sniping attack to be a rationally viable move (over mining the next block), a miner would need at least 13.5% of the total network hashrate. Comparatively, using the formula for the Reserve Pool TFM, and inputting the values of £5,179 as the average, and £9,546 as the target payout:

$$\left(\frac{x}{1-x}\right)^2 \times £9,546 > x \ \times \ £5,179 \tag{6.12}$$

We get a minimum value of $x \approx 0.2807$, meaning that for fee sniping to be a rational decision in a Reserve Pool TFM, a miner would need 28.07% of total

network hashrate. We present overall fee sniping security analysis results in Table 6.5.

These results reveal significant differences in vulnerability threshold levels of fee sniping across different TFMs. As can be seen, the 2nd-Price Auction is the most vulnerable, as a miner with just 5.07% of network hashrate would find it strategically rational to attempt to fork the blockchain, compromising overall network integrity. By contrast, the Reserve Pool TFM is far more resilient than all the other evaluated TFMs (28.07%), requiring over 5.5 times the network share rate of the 2nd-Price Auction, and roughly twice that of the 1st-Price Auction TFM. Such thresholds are considerably more difficult to pull off in practice, and would necessitate immense capital investment and computational resources. For context, as of August 2025, no single mining pool in Bitcoin would posses sufficient hashrate (Blockchain.com, 2025) to mount a feasible attack under the Reserve Pool mechanism. The largest mining pool, AntPool, controls only 15.95% of the network hashrate, which indicates that if a sufficiently high-paying block were to enter the Bitcoin network, it could become a rational decision for them to attempt to fee snipe that block.

Table 6.5: Fee Sniping Viability Threshold by TFM Type

TFM Type	**Avg. Block Payout**	**Target. Block Payout**	**Hashrate Threshold for Strategic Behaviour**
1st-Price Auction	£6,304	£34,930	13.5%
2nd-Price Auction	£1,156	£20,538	5.07%
Burning 2nd-Price Auction	£2,870	£12,211	16.42%
EIP-1559	£6,029	£21,932	18.33%
Reserve Pool	£5,179	£9,546	28.07%

In summary, the combination of low variance in block payouts, stable and predictable revenue streams for miners, the absence of disproportionally high-value outlier blocks, and the substantially elevated hashrate threshold share required to render attacks on the network economically rational, positions our novel Reserve Pool TFM as not only significantly more secure than all the currently existing transaction fee mechanisms, but also as a broadly adaptable defensive architecture in blockchains on its own. Furthermore, our findings highlight the critical role that transaction fee mechanisms can play in blockchain systems, not only in shaping the incentive structures that govern their economic dynamics, but also in fostering robust and resilient network security.

6.5 Transitioning to the Reserve Pool TFM

Adopting the Reserve Pool TFM, whether incorporating it into an established blockchain or deploying it in a new system, introduces a critical vulnerability period during its inception that demands careful planning and risk mitigation. Unlike traditional TFMs, where the effects can be experienced in full immediately, the Reserve Pool TFM design relies on the gradual accumulation of its reserves to deliver the intended stability benefits. This creates a *bootstrapping problem*, where the risk stems from the possibility of volatile transaction volumes, which could produce insufficient fee revenue to meet the target payout. This would force a reliance on the reserve pool before it has had time to reach adequate capacity. If the reserves are insufficient, or, worse, completely depleted, miners would likely encounter block payouts that fall below their operational costs, increasing the susceptibility of the network to attacks such as undercutting, fee sniping, or even a mining gap in the system. Such a scenario exposes the network to the very risks that the Reserve Pool mechanism is designed to mitigate, presenting blockchain stakeholders with a paradox: the mechanism requires operational stability to build up its reserves, yet it relies on those reserves to provide that stability in the first place.

The data from our comprehensive simulation results of a 30-day period of blockchain activity can provide us with some key insights into the dynamics of the Reserve Pool operations. As per the justifications in Chapter 5.3, we initialised our simulations with the reserve pool at an already stable state, which we took the assumption of basing at the total capacity of a full day's worth of blocks, all paying the average base fee; this value came out at 134.38 BTC, or its fiat equivalent of £2,935,481. Our simulation results presented us with several key findings:

- Over a 30-day period, or 4320 blocks, the average total value contributed to the reserve pool was 219.21 BTC, and the total taken from the pool was 23.90 BTC - this gives us a net pool change of +195.32 BTC.
- The reserve pool was positively contributed to in 79.7% of the blocks, with an average addition of 0.064 BTC per block.
- Comparatively, miners took out from the pool in 20.3% of blocks, with an average withdrawal of 0.027 BTC.

Although our simulation experiments were executed with the reserve pool in an already relatively stable state, the collected results still allow us to evaluate the potential impact of a transitional phase. In this context, the pre-existing reserve total simply served as a buffer against the potentially adverse effects that an insufficient pool could possible cause on the system; nevertheless, our simulations were still susceptible to fluctuations during the initial adaptation of the rolling 144-block payout window, up until the payouts converged towards an equilibrium value. Our results showed that over a 4320-block simulation period, on average, 32 block cycles, or 0.7% of blocks, experienced the effects of a negative pool balance. Most critically, all such instances occurred within the first 144 blocks of the simulations; we attribute this phenomenon to the transaction fee mechanism adjusting its initial network parameters during its early operational period, most particularly with regards to optimal base fee rates. Specifically, such deficits generally happened between blocks indexed 4 and 38, with the highest deficit amount peaking at –1.2576

BTC (£27,471.81). Beyond this initial deficit period, the pool balance remained consistently positive, underscoring the importance of adequately preparing for the integration phase of the Reserve Pool mechanism prior to deployment.

Ultimately, the transition to the Reserve Pool TFM presents blockchain stakeholders with a trade-off between short-term implementation vulnerability, and long-term operational stability. Whilst the initial deployment phase introduced bootstrapping risks where miners faced some periods of possible unprofitability, if the risks can be carefully mitigated, our simulation results demonstrate conclusively that the mechanism provides superior stability and security guarantees in fee-only blockchain ecosystems compared to other existing TFM alternatives.

6.6 Recommendations

As blockchain systems transition towards becoming fee-only ecosystems, protocol architects must carefully craft policies that preserve network stability, sustain economically viable mining incentives, and maintain a predictable user experience. It is important to note that protocol modifications in decentralised systems require majority approval, which can be difficult to coordinate among diverse stakeholders. In this section, we offer recommendations for future protocol designers that address both the implementation of the Reserve Pool TFM, as well as general considerations applicable to fee-only blockchain systems.

6.6.1 Reserve Pool TFM Deployment Guidelines

As discussed in Section 6.5, the Reserve Pool TFM is particularly vulnerable during its initial deployment phase. To minimise disruption to ongoing operations, we recommend a phased rollout of the mechanism, rather than an instant "flip-the-switch" deployment. The specific details of such a rollout would, of course, depend on the blockchain application in question, but a possible approach could perhaps involve a hybrid stage, in which a portion of the existing block rewards are instead

redirected to the reserve pool, rather than the miner. This would allow the pool to build-up its balance before the transition to a fee-only model is finalised. To further ensure a smooth transition, it is imperative that miners are presented with a clear and transparent transition schedule, allowing them to adapt to the changing economic environment. Protocol update guidelines should specify clear activation thresholds, such as the reserve pool reaching a set balance amount, before the transition to a fee-only model is completed, and should alternatively include rollback contingencies to revert to a previous/safe state if key metrics are not met.

Regarding specific parameter tuning, we offer the following guidelines:

- Just as we did in our simulations, the initial transaction fee base cost value should be determined via a thorough analysis of historical network statistics.
- The use of n-block moving averages for block payouts can reduce the effect of short-term revenue spikes (refer back to Section 6.3), and was also found to mitigate certain attack vectors (Section 6.4). Smaller block windows are able to quickly adapt to changing network state, while longer block windows reduce variance; ultimately, whether volatility is a negative network trait or not is dependent on the specific application's risk profile, and should therefore be a key point of debate for protocol designers.
- A cap on reserve pool withdrawals should also be carefully considered, particularly during the early bootstrapping phase, which could then be gradually relaxed as stability improves. We maintain that some form of a per-block limit should remain permanent throughout the systems lifetime, as it is necessary in disincentivising miners from strategically mining empty blocks and in turn depleting the reserve pool. Our simulated cap of 50% of the projected block payout proved effective in our experiments, although a more extensive long-term analysis could potentially provide deeper insights.

Once the in-transition system has had time to stabilise, further protocol refinements can be undertaken. As noted in Section 3.3, an excessive reserve

pool balance provides no additional utility to the overall ecosystem. Therefore, for example, the protocol could temporarily reintroduce block subsidies, or perhaps subsidise transaction fees for users, in turn rewarding the system stakeholders for honest and fair following of the protocol. The value at which a reserve pool balance is deemed adequate is subject to interpretation and would depend on the system's risk profile. Nonetheless, we reiterate that system parameters should be continuously monitored and recalibrated to best fit the changing landscape of blockchains.

6.6.2 General Fee-Only System Recommendations

Looking beyond the Reserve Pool TFM, blockchains that are relying solely on transaction fees for revenue would likely benefit from the following general practices:

- Even without the use of base fees for transactions, dynamic block size adjustments should become a key feature of fee-only ecosystems. While the particular block size target and limit values are context-dependent, as Kruminis and Navaie, 2022 found, there is no rational reason to artificially restrict the block size limit. If the blockchain application is widely used, rather than allowing congestion, the block size should be allowed to increase in relation to a growing mempool, which would both lower the fees users would pay, and increase the total in fees collected by miners.

- Blockchain governance policies should be clear and transparent, so that any protocol changes are approved based on the benefit to all stakeholders, not just the miners of validators or the network.

- The volatile nature of blockchain systems necessitates the implementation of some form of congestion management scheme, both in terms of validator reward mechanisms, and accessible transaction fee estimation tools for users.

- Given the emergence of new attack vectors in fee-only environments and the insofar limited understanding of such risks, it is critical that blockchain

architects conduct careful analysis and stay up-to-date with the latest academic literature. Continuous experimentation on testnets, followed by gradual mainnet trials, should be used to refine and inform the future security design of fee-only blockchain protocols.

Through gradual protocol transitions, disciplined and careful parameter tuning, vigilant monitoring of key system metrics, and adaptive long-term management, future protocol designers should be able to implement our novel Reserve Pool TFM, or any fee-only model, in a way that preserves network security, sustains miner incentives, and maintains user trust in the system, establishing resilient and economically sustainable fee-only markets in the next generation of blockchains.

Chapter 7

Case Study — BB-FLoC: Blockchain-Based, K-Anonymous Targeted Advertising

Overview

This chapter builds upon the theoretical principles and design recommendations by applying them to a practical case study of "Blockchain-Based Federated Learning of Cohorts" (BB-FLoC). We begin with a presentation of its overall design and architectural components, followed by an assessment of its economic model and security vulnerabilities. Our analysis compares how the system performs under different transaction fee mechanisms, most notably EIP-1559 and the Reserve Pool TFM, in order to assess which design best aligns with BB-FLoC's mission.

7.1 Application Context

Our thesis work established that stable miner incentives are fundamental to the long-term sustainability of blockchain networks; this requirement applies not only to traditional cryptocurrencies, but also to other emerging blockchain applications that

provide non-financial services. As such systems are not cryptocurrencies themselves, their transactions typically represent the transfer of data or some application-specific objects rather than a transfer of coin ownership. Therefore, their individual transactions lack inherent monetary value, failing to generate meaningful economic benefit. Without the ability to mint or distribute native coins in the form of block rewards, non-crypto blockchain platforms cannot produce predictable, protocol-level revenue streams that miners require to secure the network. Because of this, non-crypto fee-only applications are often better deployed as application layers on top of other established blockchains, such as Ethereum, which can provide the underlying security infrastructure required for the application to be viable.

To further ground these dynamics in practice, we draw on BB-FLoC, our previously published research (Kruminis, Navaie, and Ascigil, 2025) on a blockchain-based framework for privacy-preserving targeted advertising. By examining BB-FLoC's design and economic model, we demonstrate how non-cryptocurrency applications must still address the same incentive and security challenges that underpin blockchain sustainability, linking the theoretical findings of earlier chapters to the practical realities shaping the next generation of blockchain applications.

7.2 Overview of the BB-FLoC System

As a potential alternative to third-party cookies, Google, 2019 proposed Federated Learning of Cohorts (FLoC). Despite its name, the system involves no machine learning or artificial intelligence; it is believed that Google intended to incorporate federated learning in future iterations, however, these plans never materialised. In FLoC's design, each user independently applied a locality-sensitive hash (LSH) function to their browsing history to produce a cohort ID, which was then submitted in place of a cookie when visiting participating websites, allowing advertisers to target groups of users with similar interests rather than tracking individuals directly. However, since cohort IDs were calculated independently by each user, there was

no mechanism to enforce a minimum cohort size across the network; users could be assigned to very small cohorts, making them trivially identifiable. Prior attempts to resolve this relied on centralised servers that collected users' raw hashes and reorganised them to satisfy k-anonymity guarantees, introducing new trust and availability risks in the process.

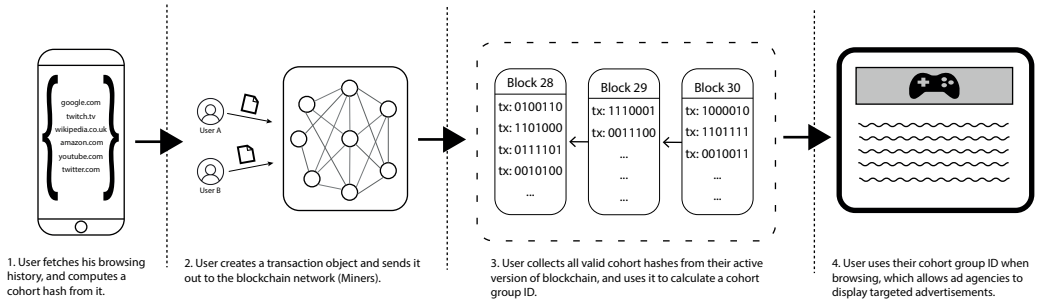

Figure 7.1: The BB-FLoC process for displaying targeted advertisements includes the four steps as illustrated in the picture.

The primary objective of our BB-FLoC system was to offer a privacy-preserving alternative to the conventional third-party cookie system for targeted advertising. The model leveraged blockchain technology to decentralise the formation of user cohort groups with similar browsing patterns while still enforcing k-anonymity guarantees. In other words, users could maintain a minimum cohort group size without relying on a trusted central authority. Figure 7.1 depicts the model's workflow, which involves users, miners, and advertising agencies:

- **Step 1.** The process begins with each user in the BB-FLoC system collecting their browsing history over some predefined period and applying a locality-sensitive hash function on it. This will result in a 50-bit value that represents their browsing activity.

- **Step 2.** Once the hash is computed, a user packages (some part of) this value into a transaction, and submits it to the blockchain network. Miners collect these transaction objects, and include them into blocks.

- **Step 3.** As such blocks are propagated throughout the network, participating users individually verify them, and once accepted, append them to their local chain. With each new block added on top of their chain, users recalculate their cohort group ID by aggregating the current set of valid hashes. Additionally, they apply a cohort reorganisation algorithm that enforces k-anonymity on their cohort group.

- **Step 4.** When a user accesses a BB-FLoC participating website, they submit their cohort group ID in response to the site request. This enables the website to display targeted advertisements to that cohort group, whilst at the same time keeping the individual's browsing history private.

The cohort reorganisation algorithm employed in Step 3 is PrefixLSH (Kruminis, Navaie, and Ascigil, 2025), which works by iteratively partitioning all submitted cohort hashes in the network by their leading bits; first separating those whose first bit is 0 from those whose first bit is 1, then further splitting each group by the next bit, and so on, until any further partition would result in a cohort smaller than the k-anonymity threshold. The final level of partition at which the threshold is still satisfied becomes the user's cohort group ID. Since this entire process relied only on the publicly visible block contents, users could compute their own cohort group ID independently, with no communication with any centralised party required.

To quantify the utility value of a cohort, we defined the following function:

$$U_C(id) = \frac{x}{\log_2(n)}, \tag{7.1}$$

where n is the total number of users in the system and x is the number of bits in the cohort group ID. A longer ID corresponded to a more specialised cohort, and thus a higher utility for advertisers, as ads could be more precisely targeted. The logarithm normalised the effect of user population size, ensuring the function remains meaningful at both small and large scales. Crucially, the utility function was evaluated from the advertisers' perspective and was independent of any individual

user's chosen k value, which was never publicly disclosed. This presents a further decentralisation advantage; rather than relying on a globally enforced privacy threshold, each user could independently set their own k value according to their own personal risk tolerance, placing even more privacy governance back in the hands of the individuals.

To evaluate the system, we developed a fully functional BB-FLoC blockchain implementation (available at Kruminis, 2023), and assessed its cohort utility using the MovieLens dataset Harper and Konstan, 2015, the same dataset that was used in Google's original FLoC whitepaper to assess cohort utility. This dataset contains a total of 27 million entries of movies watched and rated by 276,224 users between January 1995 and September 2018. This served as a suitable substitute for browsing history data, as the type of movies watched would reveal users' preference for a particular category/genre, similarly to how users' web browsing habits could reflect their possible interests.

Examining Google's limited published trial results (The Chromium Projects, 2021), we set the minimum cohort size at $k = 2{,}000$ users. In our simulations, after randomly sampling 5,000 cohort hashes from the dataset, our results showed a median cohort group ID length of 9 bits, ranging between 8 and 10 bits, yielding an advertiser utility of 0.498. This closely matched Google's own trial result of 0.549 at 13 bits, with the minor gap attributable to the difference in scale, as Google's trial involved an estimated 13,250,000 users. We further examined how utility varies across different k values. At $k = 1$, the average cohort ID reached 20 bits with a utility of 1.11, whilst at $k = 10{,}000$ it fell to 6.3 bits and a utility of 0.35. Notably, beyond a certain k threshold, utility stabilised rather than declining sharply, confirming that users could substantially strengthen their privacy guarantees with only a marginal reduction in advertiser value.

Another key feature of BB-FLoC is its lightweight mode, which enables users with limited storage or computational capacity to still participate in this process. Accessibility is essential in building a large and diverse user base, however,

the system's functional viability nevertheless remains dependent on the ongoing participation of powerful miners. Because they produce blocks and secure the ledger at their own expense, providing them with worthy economic incentives is essential to the system's continued operation.

7.3 Economic Incentive Challenges

The BB-FLoC paper suggested only a conceptual incentive model in which websites seeking to display targeted advertisements would pay fees to miners for this ability. However, a detailed evaluation of this was intentionally left open, stating that *"We leave the feasibility study of such an economic proceeding for future work"*. This chapter therefore addresses the deferred feasibility analysis.

Since BB-FLoC has no native block reward, and no inherent cryptocurrency attached to it, it is only suited to run as an application layered on top of a larger blockchain platform, such as Ethereum, as a means to provide a way of compensating miners. Nevertheless, the BB-FLoC model still contains a critical economic flaw, the free-rider problem. The system's design supports unrestricted block sizes, and allows for open participation. All the cohort hashes are also publicly available, so as a result, websites or ad agencies can independently run nodes, read the cohort data and deliver targeted ads without contributing anything to the system's infrastructure. With no mechanism to enforce some form of payment, rational participants have no reason to pay anything at all, which would result in zero revenue generated for miners, and in turn a potential collapse of the entire network.

The most practical solution is to introduce some form of barrier to entry; websites should pay a mandatory, protocol-enforced base fee to have their domain be validated for BB-FLoC participation. This could be implemented as a cryptographically verifiable *"participation certificate"*, similarly to how HTTPS certificates function (Cloudflare, 2024). Before a website can request a cohort ID from a user, it must hold a valid certificate which proves their payment of the required base fee. Since

each client verifies the certificates locally against their own version of the blockchain, the system remains decentralised, and thus helps to resolve the *"why pay?"* dilemma.

The existence of the base fee frames our analysis as that of a comparison between base-fee transaction fee mechanisms, namely the EIP-1559 and the Reserve Pool TFMs. As such, we are presented with a new challenge in the form of dynamically adjusting said base fee value. Both the EIP-1559 and the Reserve Pool TFMs update the fee based on whether the previous block's size was above or below the target block size amount; BB-FLoC's unbounded block size makes such an approach inapplicable. However, what we can do, is to redefine the notion of the *'target'*. One approach could be basing it on user activity, where the target value corresponds to some specified number of user transactions per block, thus linking the base fee directly to the network's commercial value. Alternatively, the target could be characterised in terms of the desired network hashrate, similarly to how the difficulty is adjusted in proof-of-work blockchains. Deciding between these approaches involves a trade-off between market responsiveness and security, however, both make the base-fee model viable in BB-FLoC's setting.

7.3.1 Scenario 1: EIP-1559 Implementation

Although by enforcing a mandatory $BASE\_FEE$ we can make EIP-1559 applicable to BB-FLoC, its design is ultimately unsuitable as a way to compensate miners. Under this mechanism, all the $BASE\_FEE$ amounts are burned, and miners only receive *'tip'* values as reward. Since BB-FLoC has an unbounded block size limit, with no competition for block inclusion, there is also no rational reason for paying extra through tips. Consequently, websites will simply just pay the minimum base fee, which is burned, leaving zero to be collected by miners. This distinct feature is crucial when examining the viability of BB-FLoC under the EIP-1559 transaction fee mechanism. With no way to sustain the protocol's security budget, miners have no reason to mine at all, ultimately threatening the network's existence.

7.3.2 Scenario 2: Reserve Pool TFM Implementation

The Reserve Pool TFM, which also uses a base fee approach, provides a much more sustainable solution, since the revenue of miners does not solely depend on tips. Here, instead of burning the base fees, they are collected into a network-shared reserve pool, from which miners later extract their income. The lack of rational incentive to offer tips remains, but at least here the miners are able to recoup some value for their mining efforts. Although the absolute revenue numbers from Chapter 6.2 would differ with a zero-tip environment, the results still remain applicable for illustrating the effectiveness of our unique transaction fee mechanism:

- **Economic Stability:** The revenue-smoothing mechanism significantly reduces payout variance, with previous simulation results (refer to Chapter 6.2) showing a 77% decrease compared to the EIP-1559 protocol. This stability would undoubtedly carry over to the BB-FLoC system. Additionally, our TFM achieved a 97% block profitability rate, showing that miners can expect a predictable source of income.

- **Enhanced Security:** By getting rid of extreme miner payout outliers, the Reserve Pool TFM is able to remove the primary incentive for opportunistic attacks such as fee-sniping, raising the hashrate threshold required for such exploits to a substantial value of over 28%. Additionally, by pooling together base fees rather than burning them, the mechanism is able to establish a self-sustaining economic loop, where surplus revenue generated during periods of high activity is preserved to secure the network during its more vulnerable periods, further reinforcing the system's long-term resilience.

7.4 Summary and Insights

The case study of BB-FLoC accurately demonstrated the real-life challenges of transitioning to, and applying a fee-only model on a practical blockchain system.

It showcased that without a carefully designed fee mechanism, the security and the long-term sustainability of such systems cannot be guaranteed. As such, we were also able to illustrate that our novel Reserve Pool TFM, developed and validated in this thesis, proved to be a viable and particularly effective solution, providing the economic stability that is necessary for fee-only blockchain systems to operate reliably.

The lessons learned from BB-FLoC can be extended beyond just a single application; they suggest design principles for the overall next generation of decentralised systems. By demonstrating both the vulnerabilities of EIP-1559, and the stability of the Reserve Pool mechanism, this case study points towards a future where fee-only protocols could be not only feasible, but also thriving.

Chapter 8

Conclusions and Future Work

Overview

In this final chapter, we summarise the main contributions of our thesis, highlighting the theoretical and practical insights gained from our study of fee-only blockchain protocols, with particular emphasis on the results obtained from the evaluation of our novel Reserve Pool TFM. We then outline important directions for future research, discussing potential extensions of our work, and the remaining open questions that could further inform the design, security, and stability of fee-only blockchain systems.

8.1 Contributions

This thesis examined a fundamental challenge to the long-term sustainability of blockchain systems: the transition from a dual-incentive model, combining block rewards with transaction fees, to one sustained solely by fees. This shift, referred to as *"The Great Fee Migration"*, poses a systemic risk, as it replaces a reward structure anchored in the stable, predictable income of block rewards with one that is entirely dependent on the volatile, demand-driven nature of user fees. This instability undermines the security budget of the overall network, amplifying existing

vulnerabilities and at the same time also introducing new risks, such as *mining gaps*, where block production becomes economically unattractive, or *fee sniping*, where miners attempt to reorder past blocks to capture high revenue streams. Our work progressed from identifying the impending problem to analysing its implications, and ultimately to designing, validating, and applying our novel solution to existing use-cases.

We began by establishing the economic and security challenges of a fee-only model, empirically examining historical Bitcoin data and quantifying the scale of the looming revenue shortfall. We were able to demonstrate that transaction fees have been persistently insufficient in sustaining network security on their own; in Bitcoin, since its inception in 2009, all the way to 2024, fees have historically accounted for only 3.19% of total miner revenue. Even in their peak period in 2017, fees never exceeded 13% of miner compensation, underscoring the overwhelming reliance on block rewards in dual-incentive systems like Bitcoin, a situation particularly troubling since these same block rewards are designed to be programmatically phased out over time.

This was further extended by a new theoretical perspective of fee-only market dynamics, as presented in one of our published works, **Game-Theoretic Analysis of an Exclusively Transaction-Fee Reward Blockchain System**. The findings in the paper dismantled the previously longstanding notion that miners should artificially restrict block space in an attempt to inflate their revenue from fees. Instead, we showed that miner income is primarily driven by mempool congestion levels, and that dynamic block size adjustments can benefit both users, through reduced transaction fees, and miners, by increasing their total revenue collection.

We conducted a comprehensive evaluation of existing transaction fee mechanisms, assessing their positioning within the theoretical properties of user-incentive compatibility, miner-incentive compatibility, and off-chain agreement proofness. In response to the limitations we identified in our analysis, we proposed our own **Reserve Pool TFM**, a novel mechanism that redirects surplus fees into a network-

shared reserve. This fund acts as a stabilising buffer, ensuring miners receive predictable income even during times of reduced transaction demand, ultimately reinforcing revenue stability in fee-only systems. Theoretical analysis of our Reserve Pool TFM indicates that, under stable network conditions, it successfully achieves the properties of UIC, MIC, and OCA-proofness, further advancing the field of TFM research, and opening new directions for sustainable protocol design.

To empirically evaluate the viability of our proposed design, we also developed our own novel evaluation framework, the **Zero-Block Reward Auction (ZBRA-TFM)** simulator, to the best of our knowledge, being the first comprehensive simulation tool tailored specifically for comparative TFM analysis. This modular, discrete-event simulator enables reproducible experimentation that can isolate away the effects of mechanism design from broader network variability. Leveraging a month-long Bitcoin dataset that encompassed over 12 million transactions, we conducted robust empirical evaluation of the five major TFMs, with particular emphasis on our novel Reserve Pool TFM. The quantitative results were conclusive; the Reserve Pool mechanism reduced miner payout variance by approximately 75% relative to traditional auction methods, attaining a 97% miner profitability rate, all the while also maintaining miner payouts and transaction fees at levels similar to existing mechanisms, demonstrating improved stability without sacrificing competitiveness.

In evaluating the security of TFMs in purely fee-driven environments, we constructed mathematical models that quantify the conditions under which various attacks become rational for miners, placing particular emphasis on fee-sniping strategies. Our analysis demonstrated that the Reserve Pool TFM substantially strengthens network security, increasing the hashrate required for profitable fee-sniping attacks to 28%, compared to 5% in 2nd-Price Auctions, 18% in EIP-1559, and 13% for 1st-Price Auction TFMs. This enhanced security profile stems from the Reserve Pool's ability to suppress extreme payout outliers, mitigating opportunistic behaviour and in general establishing a more stable reward environment that

incentivises honest mining strategies.

Ultimately, the practical applicability of our theoretical and empirical findings was demonstrated through the case study of **BB-FLoC: Blockchain-Based, K-Anonymous Targeted Advertising**. There, we highlighted the ability of the Reserve Pool TFM to provide a sustainable economic model in non-financial blockchain applications, where the more conventional fee mechanisms can typically be impractical to implement. BB-FLoC's requirement for a base fee made auction-based TFMs impractical for our analysis, and at the same time EIP-1559 was rendered ineffective, since its burnt fees and lack of incentive for tipping produced zero revenue for miners. By contrast, our Reserve Pool mechanism was the only TFM able to preserve economic viability, also ensuring stable and predictable payouts for its miners.

Collectively, these contributions delivered a unified theoretical framework, a practical implementation platform, and an empirical validation methodology for designing the next generation of sustainable fee-only blockchain systems. Our work not only identified and quantified the critical challenges facing blockchain networks that transition away from block rewards, but also provided concrete solutions for it through a novel mechanism design, rigorous evaluation methodologies, and practical deployment guidelines, advancing the field of fee-only blockchains towards greater resilience and economic sustainability.

8.2 Future Work

While this thesis established a solid foundation for understanding and designing transaction fee mechanisms in fee-only blockchain systems, several important directions for future research still remain that could significantly extend and strengthen our contributions.

First, our evaluations focused exclusively on PoW consensus mechanisms, particularly using Bitcoin as the underlying case study. A critical next step would

involve extending our research into PoS systems such as Ethereum. Additionally, Ethereum's 2022 transition from PoW to PoS provides for a unique experimental baseline; as a major blockchain with comprehensive historical records, it would allow for a rigorous empirical investigation into the impact that changing a consensus scheme would have on validator economics, network security, and fee dynamics.

While our security analysis primarily focused on fee-sniping attacks, there also exists a wider array of sophisticated threats that are feasible in fee-only blockchains. Future research work could involve comprehensively modelling *undercutting attacks*, where miners deliberately fork the chain while leaving some fees unclaimed in an attempt to incentivise miners to mine upon their fork, or *whale transactions*, where a miner forks the blockchain before a large transaction is confirmed and uses the recovered funds as fees to incentivise others to extend the new fork. Moreover, collusion between miners or mining pools in a fee-only environment represents a critical research gap, since coordinated strategic behaviour could substantially impact the effectiveness of any transaction fee mechanism, and the attack space associated with it. The current ZBRA-TFM simulator could be extended to model malicious behaviours of miners, and in turn simulate these complex attack scenarios, enabling for more realistic and robust security evaluations.

Our research also did not consider the impact of *Maximal Extractable Value* (MEV) with respect to miner payout analysis. MEV has become increasingly important in modern blockchain systems that support smart contracts, where validators can extract additional value through specific transaction inclusion, exclusion, and ordering strategies. Future research should investigate the effect that MEV would have on the economics of fee-only systems, and in turn whether transaction fee mechanisms like our Reserve Pool TFM can maintain their stabilising properties when MEV extraction is considered.

8.3 Final Remarks

In closing, the transition to fee-only blockchain ecosystems represents both an inevitable evolution, and a substantial challenge; without carefully designed transaction fee mechanisms, the economic foundations that incentivise honest participation and secure decentralised networks may collapse under high volatility and strategic miner behaviour. This thesis demonstrated that no existing transaction fee design can fully satisfy the theoretical properties of an ideal TFM, necessitating empirically validated solutions; by introducing the Reserve Pool TFM and rigorously evaluating it through comprehensive simulation experiments, security modelling, and case studies, we have demonstrated a pathway towards sustainable, stable, and secure fee-only blockchains. As blockchains continue to underpin critical financial and commercial application contexts, mechanisms of this kind will be essential in ensuring their long-term viability, enabling such systems to fulfil their promises of robust, trustless, and economically secure decentralisation.

References

Aave (2020). *Aave – Open Source Liquidity Protocol*. https://aave.com/. Accessed: 2025-08-27.

Bahrani, Maryam, Pranav Garimidi, and Tim Roughgarden (2023). *Transaction Fee Mechanism Design with Active Block Producers*. https://arxiv.org/abs/2307.01686. Accessed: 2025-08-01. arXiv: 2307.01686 [cs.GT].

Bao, Claire (May 2024). "Mitigating Undercutting Attacks: A Study on Mining and Transaction Fee Behavior". Master of Engineering in Electrical Engineering and Computer Science thesis. Massachusetts Institute of Technology.

Basu, Soumya et al. (2019). "Towards a Functional Fee Market for Cryptocurrencies". In: *SSRN Electronic Journal*. DOI: 10.2139/ssrn.3318327.

Beaconcha.in (2025). *Ethereum Burn Dashboard - Beaconcha.in*. https://beaconcha.in/burn. Accessed: 2025-08-05.

Bitcoin Developers (2015). *BIP 0141: Segregated Witness (Consensus layer)*. https://github.com/bitcoin/bips/blob/master/bip-0141.mediawiki. Accessed: 2025-05-22.

Bitcoin Optech (2024). *Fee sniping*. https://bitcoinops.org/en/topics/fee-sniping/. Accessed: 2025-08-12.

BitInfoCharts (2023). *BitInfoCharts*. https://bitinfocharts.com/. Accessed: 2025-08-09.

Blockchain.com (2025). *Blockchain.com Explorer — Charts*. https://www.blockchain.com/explorer/charts. Accessed: 2025-06-06.

Carlsten, Miles et al. (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. Vienna, Austria: Association for Computing Machinery, pp. 154–167. ISBN: 9781450341394. DOI: 10.1145/2976749.2978408. URL: https://doi.org/10.1145/2976749.2978408.

Chung, Hao and Elaine Shi (2023). "Foundations of Transaction Fee Mechanism Design". In: *Proceedings of the 2023 ACM-SIAM Symposium on Discrete Algorithms, SODA 2023, Florence, Italy, January 22-25, 2023*. Ed. by Nikhil Bansal and Viswanath Nagarajan. SIAM, pp. 3856–3899. DOI: 10.1137/1.9781611977554.ch150. URL: https://doi.org/10.1137/1.9781611977554.ch150.

Cloudflare (2024). *What is an SSL Certificate?* https://www.cloudflare.com/en-gb/learning/ssl/what-is-an-ssl-certificate/. Accessed: 2025-08-31.

CoinMarketCap (2025). *CoinMarketCap: Cryptocurrency Prices, Charts and Market Capitalizations*. https://coinmarketcap.com/. Accessed: 2025-08-18.

DappRadar (2025). *NFT Sales Rankings (Monthly)*. https://dappradar.com/rankings/nft/sales?range=month. Accessed: 2025-08-28.

Digiconomist (2025). *Bitcoin Energy Consumption Index*. https://digiconomist.net/bitcoin-energy-consumption. Accessed: 2025-08-27.

Dimitri, Nicolas (2019). "Transaction Fees, Block Size Limit, and Auctions in Bitcoin". In: *Ledger* 4, pp. 1–14. DOI: 10.5195/ledger.2019.145. URL: https://ledger.pitt.edu/ojs/ledger/article/view/145.

Ethereum Foundation (July 2021). *London Mainnet Announcement*. https://blog.ethereum.org/2021/07/15/london-mainnet-announcement. Accessed: 2026-02-24.

Etherscan (2020). *Ethereum transaction 0xca8f8c*. https://etherscan.io/tx/0xca8f8c315c8b6c48cee0675677b786d1babe726773829a588efa500b71cbdb65. Accessed: 2026-02-24.

— (2025). *Ethereum Network Utilization Chart*. https://etherscan.io/chart/networkutilization. Accessed: 2025-05-29.

Gafni, Yotam and Aviv Yaish (2022). "Greedy Transaction Fee Mechanisms for (Non-)myopic Miners". In: *Annual Conference of the Israeli Chapter of the Game Theory Society*.

Ganesh, Aadityan, Clayton Thomas, and S. Matthew Weinberg (2024). *Revisiting the Primitives of Transaction Fee Mechanism Design*. https://arxiv.org/abs/2410.07566. Accessed: 2025-08-01. arXiv: 2410.07566 [cs.GT].

Garimidi, Pranav, Lioba Heimbach, and Tim Roughgarden (May 2025). *Transaction Fee Mechanism Design for Leaderless Blockchain Protocols*. https://arxiv.org/abs/2505.17885. Accessed: 2025-08-01. DOI: 10.48550/arXiv.2505.17885.

Gong, Tiantian et al. (2022). "Towards Overcoming the Undercutting Problem". In: *Financial Cryptography and Data Security - 26th International Conference, FC 2022, Grenada, May 2-6, 2022, Revised Selected Papers*. Ed. by Ittay Eyal and Juan A. Garay. Vol. 13411. Lecture Notes in Computer Science. Springer, pp. 444–463. DOI: 10.1007/978-3-031-18283-9_22. URL: https://doi.org/10.1007/978-3-031-18283-9_22.

Google (2019). *Federated Learning of Cohorts (FLoC)*. https://github.com/google/ads-privacy/blob/master/proposals/FLoC/FLOC-Whitepaper-Google.pdf. Accessed: 2025-08-01.

Harper, F. Maxwell and Joseph A. Konstan (Dec. 2015). "The MovieLens Datasets: History and Context". In: *ACM Transactions on Interactive Intelligent Systems* 5.4. ISSN: 2160-6455. DOI: 10.1145/2827872. URL: https://doi.org/10.1145/2827872.

Hashrate Index (2025). *Bitcoin and Crypto ASIC Miner Profitability – Rigs*. https://hashrateindex.com/rigs. Accessed: 2025-08-11.

Hoenicke, Jochen (2024). *Johoe's Mempool Size Statistics*. https://jochen-hoenicke.de/queue/. Accessed: 2025-06-06.

Kapengut, Elie and Bruce Mizrach (2023). "An Event Study of the Ethereum Transition to Proof-of-Stake". In: *Commodities* 2.2, pp. 96–110. ISSN: 2813-2432.

DOI: 10.3390/commodities2020006. URL: https://www.mdpi.com/2813-2432/2/2/6.

Kiayias, Aggelos et al. (July 2024). "Tiered Mechanisms for Blockchain Transaction Fees". In: *Mathematical Research for Blockchain Economy – Proceedings of the 5th International Conference on Mathematical Research for Blockchain Economy (MARBLE 2024)*. DOI: 10.48550/arXiv.2304.06014. URL: https://www.marble-conference.org/marble2024-cfp.

Kruminis, Edvinas (2022). *Blockchain Model*. https://github.com/ekruminis/blockchain-model. Accessed: 2025-08-19.

— (2023). *BB-FLoC: Blockchain-Based Federated Learning of Cohorts*. https://github.com/ekruminis/floc-blockchain. Accessed: 2025-08-01.

— (2025). *ZBRA-TFM Simulator*. https://github.com/ekruminis/ZBRA-TFM-simulator. Accessed: 2025-08-03.

Kruminis, Edvinas, Onur Ascigil, and Keivan Navaie (2025). "A Transaction Fee Mechanism for Fee-Only Blockchains: The Reserve Pool Approach". Manuscript in preparation.

Kruminis, Edvinas and Keivan Navaie (2022). "Game-Theoretic Analysis of an Exclusively Transaction-Fee Reward Blockchain System". In: *IEEE Access* 10, pp. 5002–5011. DOI: 10.1109/ACCESS.2022.3140921.

Kruminis, Edvinas, Keivan Navaie, and Onur Ascigil (Aug. 2025). "BB-FLoC: A Blockchain-Based Targeted Advertisement Scheme with k-Anonymity". In: *Distributed Ledger Technologies: Research and Practice* 4.3. DOI: 10.1145/3672404. URL: https://doi.org/10.1145/3672404.

Liao, Kevin and Jonathan Katz (2017). "Incentivizing Blockchain Forks via Whale Transactions". In: *Financial Cryptography and Data Security – FC 2017 International Workshops, WAHC, BITCOIN, VOTING, WTSC, and TA, Sliema, Malta, April 7, 2017, Revised Selected Papers*, pp. 264–279. ISBN: 978-3-319-70277-3. DOI: 10.1007/978-3-319-70278-0_17.

Lin, Feilong et al. (2018). "A Sustainable Reward Mechanism for Block Mining in PoW-Based Blockchain". In: *2018 5th International Conference on Information, Cybernetics, and Computational Social Systems (ICCSS)*, pp. 156–161. DOI: 10.1109/ICCSS.2018.8572469.

mempool.space (2023). *Bitcoin transaction: 0301e0480b374b32851a9462db29dc19fe830a7f7d7a88b81612b9d42099c0ae*. https://mempool.space/tx/0301e0480b374b32851a9462db29dc19fe830a7f7d7a88b81612b9d42099c0ae. Accessed: 2025-07-30.

Myerson, Roger B. (1981). "Optimal Auction Design". In: *Mathematics of Operations Research* 6.1, pp. 58–73. URL: https://cramton.umd.edu/market-design-papers/myerson-optimal-auction-design.pdf.

Nakamoto, Satoshi (2008). *Bitcoin: A Peer-to-Peer Electronic Cash System*. https://bitcoin.org/bitcoin.pdf. Accessed: 2025-08-27.

Nisan, Noam and Amir Ronen (1999). "Algorithmic mechanism design". In: *Proceedings of the thirty-first annual ACM symposium on Theory of computing*, pp. 129–140.

PoundSterlingLive (2023). *Bitcoin (BTC) to Pound Sterling Exchange Rate History (2023)*. https://www.poundsterlinglive.com/crypto-currency/bitcoin-to-pound-sterling-history-2023. Accessed: 2025-08-09.

Roughgarden, Tim (2021a). "Transaction Fee Mechanism Design". In: *CoRR* abs/2106.01340. arXiv: 2106.01340. URL: https://arxiv.org/abs/2106.01340.

— (2021b). *Transaction Fee Mechanism Design for the Ethereum Blockchain: An Economic Analysis of EIP-1559*. https://arxiv.org/abs/2012.00854. Accessed: 2025-08-01. arXiv: 2012.00854 [cs.GT].

Sky Mavis (2018). *Axie Infinity – A Web3 Gaming Revolution*. https://axieinfinity.com/. Accessed: 2025-08-28.

Tang, Michael and Alex Zhang (2023). *Transaction Fee Mining and Mechanism Design*. https://arxiv.org/abs/2302.06769. Accessed: 2025-08-01. arXiv: 2302.06769 [cs.GT].

The Chromium Projects (2021). *FLoC Origin Trial & Clustering*. https://www.chromium.org/Home/chromium-privacy/privacy-sandbox/floc/. Accessed: 2025-08-01.

Tsabary, Itay and Ittay Eyal (2018). "The Gap Game". In: *Proceedings of the 2018 ACM SIGSAC Conference on Computer and Communications Security*. CCS '18. Toronto, Canada: Association for Computing Machinery, pp. 713–728. ISBN: 9781450356930. DOI: 10.1145/3243734.3243737. URL: https://doi.org/10.1145/3243734.3243737.

Uniswap Labs (2018). *Uniswap Interface*. https://app.uniswap.org/. Accessed: 2025-08-27.

Worldometer (2025). *GDP by Country*. https://www.worldometers.info/gdp/gdp-by-country/. Accessed: 2025-08-18.

Wu, Ke, Elaine Shi, and Hao Chung (2024). "Maximizing Miner Revenue in Transaction Fee Mechanism Design Under a Known-h-Honest-Users Assumption". In: *15th Innovations in Theoretical Computer Science Conference (ITCS 2024)*. Vol. 287. Leibniz International Proceedings in Informatics (LIPIcs). Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 98:1–98:23. DOI: 10.4230/LIPIcs.ITCS.2024.98.

Yao, Andrew Chi-Chih (2018). "An Incentive Analysis of some Bitcoin Fee Designs". In: *International Colloquium on Automata, Languages and Programming*.

Zhang, Gerui, Xiongfei Zhao, and Yain-Whar Si (2023). *A Comparative Analysis on Volatility and Scalability Properties of Blockchain Compression Protocols*. https://arxiv.org/abs/2303.17643. Accessed: 2025-08-01. arXiv: 2303.17643 [cs.CR].