

Unraveling Ethereum's Mempool: The Impact of Fee Fairness, Transaction Prioritization, and Consensus Efficiency

S M Mostaq Hossain and Amani Altarawneh

*Department of Computer Science, Tennessee Technological University
Cookeville, Tennessee, USA*

Email: {shossain42, aaltarawneh}@tntech.edu

Abstract—Ethereum's transaction pool (mempool) dynamics and fee market efficiency critically affect transaction inclusion, validator workload, and overall network performance. This research empirically analyzes gas price variations, mempool clearance rates, and block finalization times in Ethereum's proof-of-stake ecosystem using real-time data from Geth and Prysm nodes. We observe that high-fee transactions are consistently prioritized, while low-fee transactions face delays or exclusion—despite EIP-1559's intended improvements. Mempool congestion remains a key factor in validator efficiency and proposal latency. We provide empirical evidence of persistent fee-based disparities and show that extremely high fees do not always guarantee faster confirmation, revealing inefficiencies in the current fee market. To address these issues, we propose congestion-aware fee adjustments, reserved block slots for low-fee transactions, and improved handling of out-of-gas vulnerabilities. By mitigating prioritization bias and execution inefficiencies, our findings support more equitable transaction inclusion, enhance validator performance, and promote scalability. This work contributes to Ethereum's long-term decentralization by reducing dependence on high transaction fees for network participation.

Index Terms—Ethereum Mempool, Transaction Inclusion, Fee Market Efficiency, Consensus Latency, Gas Price Dynamics

1. Introduction

Ethereum [1], a leading blockchain platform for smart contracts, relies on the interaction between its execution and consensus layers [2]. The execution layer processes transactions, while the consensus layer ensures agreement on the blockchain's state. A key component in this system is the transaction pool (mempool) [3], where pending transactions compete for inclusion based on gas fees [4]. The efficiency and fairness of this process directly impact transaction inclusion rates, block finalization times, and overall network performance [5]. To improve fee market efficiency, Ethereum introduced EIP-1559 [6], which replaces the first-price auction model with a dynamic base fee that adjusts per block based on congestion and an optional priority fee (tip) for faster inclusion. The base fee is burned, reducing ETH supply, while the tip incentivizes validators. Understanding these dynamics is essential for optimizing Ethereum's scalability and transaction processing.

Despite advancements such as EIP-1559, challenges remain in ensuring both efficient and fair transaction inclusion. The mempool plays a crucial role in shaping network congestion [7], as high-fee transactions tend to be prioritized, often leaving lower-fee transactions pending or excluded. This dynamic raises concerns about Ethereum's fee market fairness and its impact on validator efficiency. While EIP-1559 introduced a more predictable base fee model, fee bidding competition continues to favor high-premium transactions, potentially disadvantaging smaller users [8]. A deeper investigation into how mempool congestion and gas bidding strategies influence network efficiency and fairness is necessary to improve Ethereum's transaction processing.

Existing research has explored Ethereum's fee structures [9] and mempool behavior, focusing on gas price volatility and confirmation times. However, a comprehensive study linking mempool congestion to consensus-layer performance, particularly block finalization and validator workload, remains limited. Additionally, while EIP-1559 has been studied for reducing fee volatility [10], its impact on transaction fairness is still not fully understood. Addressing these gaps is crucial for enhancing Ethereum's scalability and equitable transaction processing. This research makes the following contributions:

- **Analysis of Mempool-Consensus Interaction:** We investigate how mempool congestion and gas price variations impact block finalization times and validator efficiency.
- **Evaluation of Fee Market Fairness Post-EIP-1559:** While EIP-1559 aimed to stabilize gas fees and reduce volatility [10], our findings reveal that fee-based inclusion biases persist and validator efficiency is still constrained during peak congestion. We offer novel data-driven insights that extend beyond prior evaluations of the EIP-1559 mechanism.
- **Empirical Insights for Protocol Optimization:** Our findings provide data-driven recommendations for improving Ethereum's fee mechanisms and transaction prioritization.

By addressing these issues, this study contributes to a more efficient and fair Ethereum transaction processing framework. The paper is structured as follows: Section 2 defines key research questions, Section 3 covers Ethereum's transaction processing and fee mechanisms, and Section 4 reviews related literature. Section 5 details the research framework, Section 6 describes the experi-

mental setup, and Section 7 presents empirical findings. Section 8 interprets results and suggests optimizations, while Section 9 summarizes insights and future research directions.

2. Research Questions

This work investigates Ethereum’s transaction dynamics, fee fairness, and execution-consensus correlation to identify inefficiencies and biases in transaction processing. By analyzing mempool behavior, fee market fairness post-EIP-1559, and transaction inclusion trends, we aim to provide insights for improving network efficiency and equitable access. The key research questions are:

- **RQ-i:** How do Ethereum’s mempool dynamics affect transaction inclusion and network efficiency?
- **RQ-ii:** Does Ethereum’s fee market structure post-EIP-1559 ensure fair transaction processing across different fee levels?
- **RQ-iii:** How does transaction prioritization impact inclusion fairness, and to what extent do high-fee transactions dominate block space?

These questions guide our empirical analysis and contribute to a deeper understanding of Ethereum’s transaction processing mechanisms and potential optimizations.

3. Background

To analyze Ethereum’s transaction dynamics and their impact on execution-consensus efficiency, we first outline key blockchain components: transaction processing, fee mechanisms, mempool operations, and consensus finality.

3.1. Ethereum Transaction Processing Pipeline

Ethereum transactions progress [11] through multiple stages before finalization. After submission, they enter the mempool, where validators prioritize them based on gas fees and propose blocks for inclusion. The time from submission to finality [12] depends on gas pricing strategies, network congestion, and transaction selection mechanisms. Understanding these processes is crucial for identifying inefficiencies in transaction inclusion and execution delays.

3.2. Ethereum’s Fee Mechanism and EIP-1559

Ethereum initially used a first-price auction model for gas fees, resulting in volatile transaction costs [13]. EIP-1559 [6] introduced a dual-fee model with a dynamically adjusted base fee and an optional priority tip. While this change improved fee predictability, concerns persist regarding fee fairness and the prioritization of high-fee transactions, leading to potential exclusion of lower-fee bids. Analyzing Ethereum’s fee structure is essential to understanding its impact on transaction processing.

3.3. Ethereum Mempool Dynamics

The mempool acts as Ethereum’s transaction buffer, where pending transactions compete for inclusion based

on gas fees [14]. Network congestion and transaction clustering influence clearance rates, while high-fee transactions often bypass lower-fee ones. Investigating how mempool behavior affects validator efficiency and transaction inclusion probability provides insights into network-wide inefficiencies.

3.4. Consensus Mechanism and Block Finalization

With Ethereum’s transition to Proof-of-Stake (PoS), validators now propose and attest to blocks [15]. However, mempool congestion and transaction prioritization strategies can delay block finalization. Studying the correlation between mempool conditions and consensus efficiency helps optimize validator performance and reduce transaction confirmation delays.

3.5. Fee Fairness and Transaction Prioritization

Fairness in Ethereum’s fee market remains a critical issue [16]. Our study examines whether the current fee model equitably processes transactions or systematically favors high-fee bidders. By evaluating fee market efficiency and transaction latency, we assess whether Ethereum’s fee structure promotes fair inclusion or exacerbates accessibility disparities.

With these foundational concepts established, the next section reviews related work on Ethereum’s transaction dynamics, fee mechanisms, and consensus behaviors to position our research within the broader blockchain literature.

4. Related Works

Several studies have analyzed the Ethereum mempool, focusing on transaction inclusion policies and gas fee strategies. Research on mempool congestion and transaction prioritization highlights how gas fees influence the speed and likelihood of transaction inclusion [17]. Furthermore, studies on mempool security vulnerabilities emphasize the role of mempool eviction mechanisms, showing that transaction ordering is often manipulated through MEV (Maximal Extractable Value) techniques [3]. Another study explores alternative lightweight transaction pools for IoT and mobile clients, demonstrating efficiency trade-offs in transaction propagation [18].

Ethereum’s EIP-1559 fee model introduced a dynamic base fee [19] adjustment mechanism to stabilize transaction costs. A research [13] of the EIP-1559 Ethereum fee market shows that this model lowers fee volatility but doesn’t get rid of transaction prioritization biases completely. Another empirical study on Ethereum transaction processing times shows that high gas fees [4] continue to dominate transaction inclusion decisions despite the introduction of base fees [20]. Furthermore, an investigation [21] into large Ethereum blocks highlights how gas fee auctions and block size constraints affect network congestion and transaction processing speeds. Correlating Ethereum’s execution and consensus layers is essential to optimizing validator performance. One study shows how validator selection under PoS affects confirmation times [22], while another uses machine learning

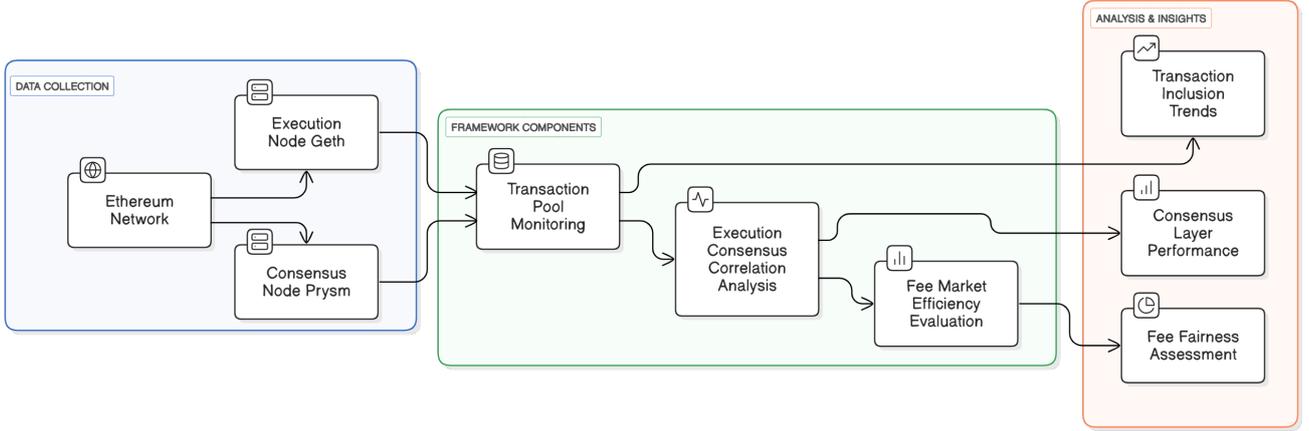


Figure 1. The framework analyzes mempool behavior, execution-consensus correlation, and fee market efficiency using real-time monitoring, statistical modeling, and empirical validation. Key components include transaction pool monitoring, execution-consensus analysis, and fee market evaluation.

TABLE 1. COMPARISON WITH RELATED WORK

Study	Focus	Gap Addressed
Leonardos et al. [13]	EIP-1559 modeling	No live-node validation
Pacheco et al. [20]	Transaction timing study	Lacks fairness or validator insights
Wang et al. [3]	Mempool DoS analysis	No latency or fairness view
Alipanahloo et al. [7]	MEV defense focus	Doesn't generalize to fee fairness
This Work	Fee fairness + validator efficiency	Post-EIP-1559 congestion + inclusion link

to link gas fee predictions with clearance rates [23]. Symbolic fuzzing reveals DoS exploits that manipulate transaction prioritization, raising concerns about mempool resilience [3]. Fee fairness studies show that low-fee transactions are systematically deprioritized, reducing accessibility for smaller users [11].

Ethereum's peer-to-peer (P2P) network plays a critical role in transaction propagation and mempool synchronization [24]. A study on Ethereum's P2P network structure [18] examines client diversity and node discovery inefficiencies, revealing potential bottlenecks that impact transaction broadcasting. Research into Ethereum's blob fee market [25] also sheds light on the complexities of gas price adjustments and their impact on transaction finalization [11]. While our study does not directly model MEV extraction techniques, the observed fee-based inclusion biases and congestion patterns offer insights relevant to MEV behavior. As shown in [7], mempool conditions affect transaction ordering and bundling strategies; our findings on prioritization under fee pressure highlight systemic patterns that MEV actors could exploit or be constrained by.

5. Methodology

The research methodology is presented for analyzing Ethereum's mempool behavior, execution-consensus correlation, and fee market efficiency. The approach integrates real-time transaction monitoring, statistical modeling, and empirical analysis of transaction prioritization.

5.1. Research Framework Overview

Figure 1 illustrates the structure of the framework, which consists of three key components:

Transaction Pool Monitoring: Tracks real-time mempool activity using `txpool` RPCs to evaluate gas fees and congestion.

Execution-Consensus Correlation Analysis: Assesses how congestion and gas fees affect block proposals and latency.

Fee Market Efficiency Evaluation: Measures inclusion probabilities, gas price effects, and fairness deviations.

The research process involves real-time data collection, empirical modeling, and statistical validation to assess network-wide transaction processing efficiency.

5.2. Transaction Pool Monitoring and Mempool Dynamics

The Ethereum mempool is a temporary queue for pending transactions. Clearance rates depend on gas bidding and congestion [4]. We define a proxy metric for the likelihood of a transaction T_{xi} being included in the next block B_t , based on its relative gas price compared to competing transactions. This is not a formal probability definition, but a normalized approximation derived for empirical correlation analysis:

$$P(T_{xi} \in B_t) = \frac{G(T_{xi})}{\sum_{j=1}^N G(T_{xj})} \quad (1)$$

where: $G(T_{xi})$ represents the gas price of transaction T_{xi} , and N is the total number of transactions in the mempool. This proxy allows us to evaluate how relative fee strength influences transaction inclusion likelihood during congestion.

5.3. Execution-Consensus Correlation

Ethereum's Proof-of-Stake consensus relies on efficient validator proposals [2]. Mempool congestion and

bidding increase latency. To estimate the effect of mempool congestion on block finalization time, we use a regression-based empirical model expressed as:

$$T_{final} = \alpha \cdot M_{pend} + \beta \cdot G_{avg} + \gamma \quad (2)$$

This linear model is used to fit observed data where T_{final} is the block finalization time, M_{pend} is the pending transaction count, G_{avg} is the average gas price in the block, and α, β, γ are regression coefficients learned from our dataset.

5.4. Fee Market Efficiency and Transaction Prioritization

Despite EIP-1559’s base fees and tips, prioritization remains fee-driven [26]. We measure fee market efficiency [27] by evaluating transaction inclusion probability and latency variations. To assess inclusion disparities for low-fee transactions, we define a fairness deviation metric based on expected and observed inclusion rates:

$$\Delta F_{fair} = \frac{I_{LP}^{obs} - I_{LP}^{exp}}{I_{LP}^{exp}} \quad (3)$$

Here, I_{LP}^{obs} is the observed fraction of low-fee transactions included in blocks over a given period, while I_{LP}^{exp} represents the expected fraction assuming proportional inclusion relative to their presence in the mempool. Our fairness deviation metric focuses solely on transaction-level inclusion disparities. We do not evaluate validator-side fairness (e.g., block selection diversity, revenue variance), which remains an open area for future exploration.

5.5. Statistical Validation and Empirical Analysis

We validate our findings using regression analysis (gas price vs. inclusion), time-series models (congestion vs. latency), and a fairness index (observed vs. expected inclusion). These models reveal Ethereum’s prioritization behavior and suggest fairness and efficiency optimizations.

6. Experimental Setup

To evaluate Ethereum’s mempool behavior, consensus correlation, and fee efficiency, we built an experimental setup using a live node. It captures transaction events, validator metrics, and gas fee changes for comprehensive analysis.

6.1. Node Setup and Data Collection

We deployed a full Ethereum execution node (Geth v1.13.10) [28] and a consensus layer validator node (Prysm v5.1.2) [29] to track transaction flow and validator decision-making. The `txpool` RPC methods `txpool.content`, `txpool.inspect` provided live insights into pending transactions, gas price distribution, and mempool clearance rates, while `eth.gasPrice` monitored real-time fee adjustments. The Prysm beacon chain API (`/eth/v1/beacon/blocks`) recorded

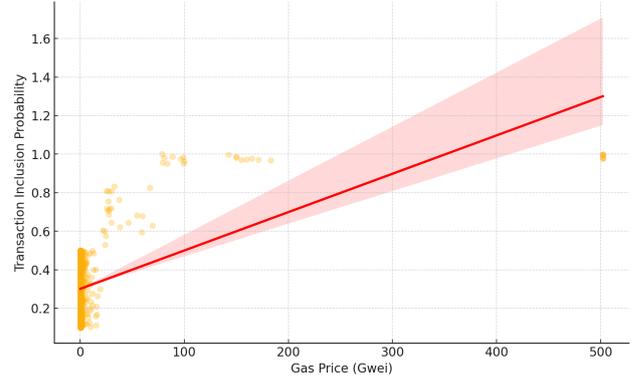


Figure 2. The plot shows a positive correlation between gas price and inclusion probability. Higher fees lead to near-certain inclusion, while low-fee transactions face uncertainty and delays.

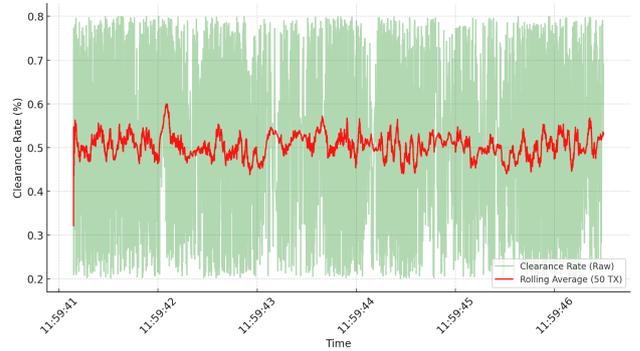


Figure 3. Mempool clearance rate over time, with raw fluctuations and a red 50-transaction rolling average. The x-axis shows wall-clock time; data is indexed by block height. Y-axis is limited to 20–80% for clarity.

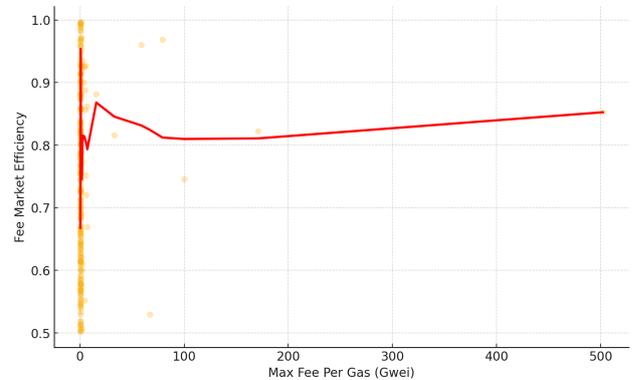


Figure 4. The plot illustrates the relationship between max fee per gas and fee market efficiency, quantified using the fairness deviation metric defined in Equation 3.

block timestamps and proposer efficiency, enabling correlation between mempool congestion and validator workload. The nodes ran on a dedicated system (**Ubuntu 22.04, 16-core CPU, 64GB RAM, 2TB NVMe SSD**) to handle high-throughput transaction monitoring.

6.2. Testing and Transaction Tracking

We analyzed transaction inclusion and latency by broadcasting transactions with varying gas fees. A Python script automated submission, mempool logging, and block

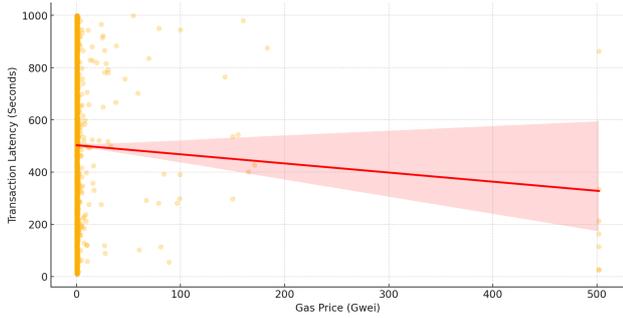


Figure 5. The plot shows that higher gas prices reduce transaction latency, meaning transactions with higher fees are confirmed faster, though other network factors also influence delays.

tracking over time. We also used time-series analysis to examine how gas price fluctuations and congestion affect proposal times and processing efficiency.

6.3. Experimental Evaluation Metrics

We assess Ethereum’s mempool efficiency, fairness, and execution-consensus correlation using four key metrics: Mempool Clearance Rate (transaction mining speed), Inclusion Probability (likelihood based on gas price), Fee Market Efficiency (impact of fee variation on performance), and Transaction Latency (delay from submission to inclusion). Together, they offer a comprehensive view of Ethereum’s processing and fee dynamics.

7. Result and Analysis

This section presents results from our empirical evaluation of Ethereum’s mempool behavior, execution-consensus correlation, and fee fairness. We analyze inclusion probabilities, clearance rates, fee market efficiency, and how gas prices affect transaction latency.

7.1. Transaction Inclusion Probability vs. Gas Price

Figure 2 shows a strong positive correlation between gas price and transaction inclusion probability. Low-fee transactions exhibit high variability, while those above 75 Gwei reach near 100% inclusion. Below 20 Gwei, transactions often remain pending or delayed. A logistic trend emerges, indicating a threshold where transactions shift from low to high priority. This supports the view that EIP-1559 does not eliminate prioritization bias, and low-fee transactions remain disadvantaged. The observed correlation aligns with the inclusion proxy in Eq. 1. While a linear regression is shown for simplicity, the actual trend follows a logistic curve, with diminishing returns at high gas prices.

7.2. Mempool Clearance Rate Over Time

Figure 3 shows fluctuations in mempool clearance rates over time. The raw data (green) and rolling average (red), computed over 30 blocks, highlight medium-term congestion patterns while smoothing short-term volatility. Clearance rates range between 40% and 60%, with

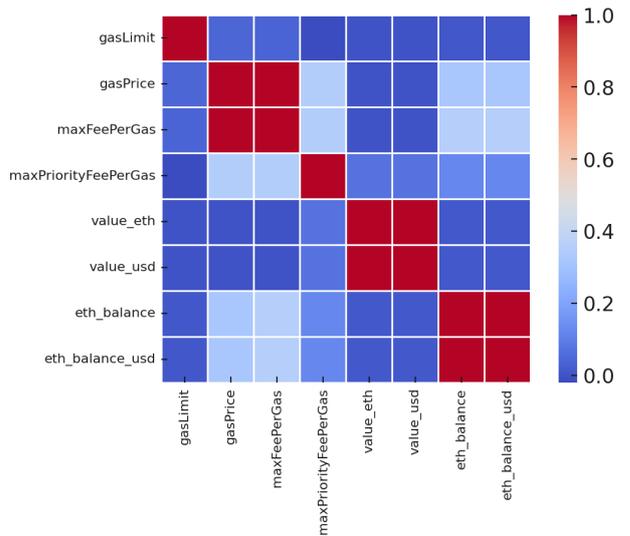


Figure 6. Correlation heatmap showing strong ties between Gas Price and Max Fee, and between ETH and USD. A moderate link between Mempool Size and Finalization Time suggests congestion affects validation speed.

occasional spikes. Although the x-axis shows wall-clock time, data is indexed by block height; the y-axis is bounded from 20% to 80% for clarity. High congestion (low clearance rates) correlates with transaction delays, illustrating how fee volatility impacts performance. Validator behavior appears stable, but transaction processing is affected by dynamic congestion. These results (see Eq. 2 and Fig. 6) confirm that Ethereum’s mempool load varies significantly, impacting inclusion probabilities.

7.3. Max Fee Per Gas vs. Fee Market Efficiency

The third analysis, illustrated in Figure 4, explores the relationship Figure 4 shows how fee market efficiency, measured using the fairness deviation metric ΔF_{fair} (Eq. 3), varies with max fee per gas. Efficiency rises sharply at low max fees, then plateaus around 80% for higher fees. At very low fee levels, efficiency fluctuates, indicating unpredictable inclusion. Above 50 Gwei, efficiency stabilizes, suggesting diminishing returns. Thus, high fees do not guarantee better performance, reinforcing that Ethereum’s fee market operates best within an optimal range. A semi-log x-axis may offer better resolution of low-fee behavior and is worth exploring in future work.

7.4. Transaction Latency vs. Gas Price

Figure 5 presents the relationship between transaction latency and gas price. The negative regression trend indicates that higher gas prices lead to lower transaction latency, meaning transactions with higher fees are confirmed faster. Low-fee transactions show high latency variance, with some confirming quickly while others face delays. Higher gas prices reduce latency, confirming the priority of high-fee transactions. The shaded confidence interval indicates that factors beyond gas price, such as block space competition, also impact transaction delays.

7.5. key observations from the correlation heatmap

Finally, The correlation heatmap in Figure 6 reveals key relationships among transaction metrics. Gas Price and Max Fee Per Gas exhibit a strong positive correlation, indicating that higher gas prices correspond to higher max fees, aligning with Ethereum’s fee mechanism. ETH and USD values are highly correlated, as USD value is derived from ETH based on market rates. Additionally, Mempool Size and Block Finalization Time show a moderate correlation, suggesting that larger mempools contribute to longer block processing times, likely due to network congestion.

8. Discussion and Improvement

This section discusses how our experimental findings address the key research questions and highlights potential improvements to Ethereum’s transaction processing and fee fairness mechanisms. To address the research questions, the three takeaways (*RQ-i*, *RQ-ii* and *RQ-iii*) for the research questions have been discussed.

Take-Away *RQ-i*

Our analysis shows that mempool congestion significantly impacts transaction clearance and network efficiency. This persists despite EIP-1559’s base-fee adjustment, which alone fails to resolve fee-based inequality or optimize validator performance. Clearance rates range from 40% to 60%, with congestion causing delays and variable processing efficiency. Higher gas fees reduce latency, reinforcing fee-based over first-come, first-served prioritization.

Take-Away *RQ-ii*

Ethereum’s fee market still favors high-fee transactions. Gas prices below 20 Gwei face delays or exclusion, while those above 75 Gwei see near 100% inclusion. EIP-1559 does not eliminate prioritization biases. Fee market efficiency stabilizes around 80% for higher fees, suggesting excessive fees offer little execution advantage.

Take-Away *RQ-iii*

Low-fee transactions face delays and uncertainty, while high-fee ones dominate inclusion. Prioritization mechanisms create systemic disadvantages, and fairness deviations persist as efficiency plateaus beyond a threshold. Although Ethereum-specific, these congestion-driven patterns and fee biases likely extend to other PoS blockchains using mempool-based selection (e.g., Avalanche, Cosmos) [30]. Our fairness metric and congestion-aware strategies thus hold broader relevance.

Improving Ethereum’s mempool efficiency, fairness, and validator performance requires protocol refinements. Dynamic fee adjustments can smooth processing during

congestion, while encrypted mempools help prevent MEV-driven reordering by keeping transactions private until inclusion [31]. Reserving block space for low-fee transactions prevents exclusion without harming efficiency [7]. Addressing out-of-gas issues [32] and using tools like eTainter [33] to detect gas-related bugs improve execution reliability. These steps enhance fairness, reduce congestion, and align validator incentives with network stability.

9. Conclusion

This research investigated Ethereum’s transaction pool dynamics, the interaction between the execution and consensus layers, and the fairness implications of gas fee prioritization. Through real-time data collection from Geth and Prysm nodes, we empirically evaluated transaction inclusion trends, mempool clearance rates, and block finalization delays in the context of EIP-1559. Our findings confirm that Ethereum’s transaction inclusion remains heavily fee-driven: high-fee transactions are consistently prioritized, while low-fee transactions face delays or exclusion, especially during periods of congestion. Despite the stabilizing intentions of EIP-1559, our analysis reveals that fee-based disparities persist and that mempool congestion significantly impacts validator efficiency and finalization latency. Moreover, we show that excessively high fees do not guarantee better outcomes, highlighting inefficiencies in the current fee market. To address these limitations, we propose congestion-aware fee adjustment mechanisms, reserved block slots for low-fee transactions, and robust handling of out-of-gas vulnerabilities in smart contracts to prevent execution failures. These contributions extend current understanding of Ethereum’s fee dynamics and validator performance bottlenecks, offering data-driven insights for improving scalability and fairness. By mitigating systemic prioritization bias and enhancing execution reliability, this work supports a more inclusive, efficient, and decentralized Ethereum network—principles that are equally applicable to other Proof-of-Stake blockchain ecosystems.

10. Acknowledgments

The authors acknowledge the use of AI tools, including ChatGPT for text refinement, summarization, and assistance in structuring content, as well as Eraser AI for figure design and illustration. AI-assisted content was used to enhance clarity and presentation, while all technical content and conclusions were independently verified by the authors.

References

- [1] V. Buterin *et al.*, “A next-generation smart contract and decentralized application platform,” *white paper*, vol. 3, no. 37, pp. 2–1, 2014. [Online]. Available: https://www.weusecoins.com/assets/pdf/library/Ethereum_white_paper-a_next_generation_smart_contract_and_decentralized_application_platform-vitalik-buterin.pdf
- [2] H. Feng, Y. Hu, Y. Kou, R. Li, J. Zhu, L. Wu, and Y. Zhou, “SlimArchive: A lightweight architecture for ethereum archive nodes,” in *2024 USENIX Annual Technical Conference (USENIX ATC 24)*. Santa Clara, CA: USENIX Association, Jul. 2024, pp. 1257–1272. [Online]. Available: <https://www.usenix.org/conference/atc24/presentation/feng-hang>

- [3] Y. Wang, Y. Tang, K. Li, W. Ding, and Z. Yang, "Understanding ethereum mempool security under asymmetric DoS by symbolized stateful fuzzing," in *33rd USENIX Security Symposium (USENIX Security 24)*. Philadelphia, PA: USENIX Association, Aug. 2024, pp. 4747–4764. [Online]. Available: <https://www.usenix.org/conference/usenixsecurity24/presentation/wang-yibo>
- [4] B. K. Meister and H. C. Price, "Gas fees on the ethereum blockchain: from foundations to derivative valuations," *Frontiers in Blockchain*, vol. 7, p. 1462666, 2024. [Online]. Available: <https://www.frontiersin.org/journals/blockchain/articles/10.3389/fbloc.2024.1462666>
- [5] D. Grandjean, L. Heimbach, and R. Wattenhofer, "Ethereum proof-of-stake consensus layer: Participation and decentralization," in *Financial Cryptography and Data Security. FC 2024 International Workshops*, J. Budurushi, O. Kulyk, S. Allen, T. Diamandis, A. Klages-Mundt, A. Bracciali, G. Goodell, and S. Matsuo, Eds. Cham: Springer Nature Switzerland, 2025, pp. 253–280. [Online]. Available: https://doi.org/10.1007/978-3-031-69231-4_17
- [6] V. Buterin, E. Conner, R. Dudley, M. Slipper, I. Norden, and A. Bakhta, "Eip-1559: Fee market change for eth 1.0 chain," Ethereum Improvement Proposal, 2019, accessed: 2025-03-14. [Online]. Available: <https://eips.ethereum.org/EIPS/eip-1559>
- [7] Z. Alipanahloo, A. S. Hafid, and K. Zhang, "Maximum extractable value (mev) mitigation approaches in ethereum and layer-2 chains: A comprehensive survey," *IEEE Access*, vol. 12, pp. 185 212–185 231, 2024. [Online]. Available: <https://doi.org/10.1109/ACCESS.2024.3514375>
- [8] A. Kiayias, E. Koutsoupias, P. Lazos, and G. Panagiotakos, "Tiered mechanisms for blockchain transaction fees," 2023. [Online]. Available: <https://arxiv.org/abs/2304.06014>
- [9] T. Roughgarden, "Transaction fee mechanism design," *SIGecom Exch.*, vol. 19, no. 1, p. 52–55, Jul. 2021. [Online]. Available: <https://doi.org/10.1145/3476436.3476445>
- [10] G. A. Pierro and H. Rocha, "The influence factors on ethereum transaction fees," in *2019 IEEE/ACM 2nd International Workshop on Emerging Trends in Software Engineering for Blockchain (WETSEB)*, 2019, pp. 24–31. [Online]. Available: <https://doi.org/10.1109/WETSEB.2019.00010>
- [11] M. Pacheco, G. A. Oliva, G. K. Rajbahadur, and A. E. Hassan, "What makes ethereum blockchain transactions be processed fast or slow? an empirical study," *Empirical Software Engineering*, vol. 28, no. 2, p. 39, 2023. [Online]. Available: <https://doi.org/10.1007/s10664-022-10283-7>
- [12] S. Tikhomirov, "Ethereum: State of knowledge and research perspectives," in *Foundations and Practice of Security*, A. Imine, J. M. Fernandez, J.-Y. Marion, L. Logrippo, and J. Garcia-Alfaro, Eds. Cham: Springer International Publishing, 2018, pp. 206–221. [Online]. Available: https://doi.org/10.1007/978-3-319-75650-9_14
- [13] S. Leonardos, B. Monnot, D. Reijbergen, E. Skoulakis, and G. Piliouras, "Dynamical analysis of the eip-1559 ethereum fee market," ser. AFT '21. New York, NY, USA: Association for Computing Machinery, 2021, p. 114–126. [Online]. Available: <https://doi.org/10.1145/3479722.3480993>
- [14] X. Lyu, M. Zhang, X. Zhang, J. Niu, Y. Zhang, and Z. Lin, "An empirical study on ethereum private transactions and the security implications," 2022. [Online]. Available: <https://arxiv.org/abs/2208.02858>
- [15] C. T. Nguyen, D. T. Hoang, D. N. Nguyen, D. Niyato, H. T. Nguyen, and E. Dutkiewicz, "Proof-of-stake consensus mechanisms for future blockchain networks: Fundamentals, applications and opportunities," *IEEE Access*, vol. 7, pp. 85 727–85 745, 2019. [Online]. Available: <https://doi.org/10.1109/ACCESS.2019.2925010>
- [16] A. Orda and O. Rottenstreich, "Enforcing fairness in blockchain transaction ordering," *Peer-to-peer Networking and Applications*, vol. 14, no. 6, pp. 3660–3673, 2021. [Online]. Available: <https://doi.org/10.1007/s12083-021-01165-8>
- [17] A. Hausman, "Streamlining transaction management: Exploring ethereum mempool apis," *Hausman Marketing Letter*, March 2024, published on March 1, 2024. [Online]. Available: <https://tinyurl.com/ypy49m65>
- [18] H. B. Haq, T. Ahmad, A. Buriro, and S. Ullah, "Neonpool: Reimagining cryptocurrency transaction pools for lightweight clients and iot devices," 2024. [Online]. Available: <https://arxiv.org/abs/2412.16217>
- [19] M. Cambus, F. Kuhn, S. Pai, and J. Uitto, "Time and space optimal massively parallel algorithm for the 2-ruling set problem." Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2023. [Online]. Available: <https://drops.dagstuhl.de/entities/document/10.4230/LIPIcs.DISC.2023.11>
- [20] M. Pacheco, G. Oliva, G. K. Rajbahadur, and A. Hassan, "Is my transaction done yet? an empirical study of transaction processing times in the ethereum blockchain platform," *ACM Trans. Softw. Eng. Methodol.*, vol. 32, no. 3, Apr. 2023. [Online]. Available: <https://doi.org/10.1145/3549542>
- [21] P. Ocheja, M. Cortes-Goicoechea, T. Mohandas-Daryanani, B. Flanagan, H. Ogata, J. L. Munoz, and L. Bautista-Gomez, "An analytical study of large blocks on ethereum," in *Proceedings of the 2024 6th Blockchain and Internet of Things Conference*, ser. BIOTC '24. New York, NY, USA: Association for Computing Machinery, 2024, p. 120–127. [Online]. Available: <https://doi.org/10.1145/3688225.3688240>
- [22] F. Javed and J. Mangues-Bafalluy, "An empirical smart contracts latency analysis on ethereum blockchain for trustworthy inter-provider agreements," 2025. [Online]. Available: <https://arxiv.org/abs/2503.01397>
- [23] V. C. Oliveira, J. Almeida Valadares, J. E. A. Sousa, A. Borges Vieira, H. S. Bernardino, S. Moraes Villela, and G. Dias Goncalves, "Analyzing transaction confirmation in ethereum using machine learning techniques," *SIGMETRICS Perform. Eval. Rev.*, vol. 48, no. 4, p. 12–15, May 2021. [Online]. Available: <https://doi.org/10.1145/3466826.3466832>
- [24] J. Luo, "Unveiling ethereum's p2p network: The role of chain and client diversity," 2025. [Online]. Available: <https://arxiv.org/abs/2501.16236>
- [25] L. Heimbach and J. Milionis, "The early days of the ethereum blob fee market and lessons learnt," 2025. [Online]. Available: <https://arxiv.org/abs/2502.12966>
- [26] A. A. Zarir, G. A. Oliva, Z. M. J. Jiang, and A. E. Hassan, "Developing cost-effective blockchain-powered applications: A case study of the gas usage of smart contract transactions in the ethereum blockchain platform," vol. 30, no. 3, Mar. 2021. [Online]. Available: <https://doi.org/10.1145/3431726>
- [27] S. M. Werner, P. J. Pritz, and D. Perez, "Step on the gas? a better approach for recommending the ethereum gas price," 2021. [Online]. Available: <https://arxiv.org/abs/2003.03479>
- [28] T. G. team, "Go ethereum (geth)," GitHub repository, 2025, accessed: 2025-03-14. [Online]. Available: <https://github.com/ethereum/go-ethereum>
- [29] P. Labs and O. Labs, "Prysm: Ethereum consensus client," GitHub repository, 2025, accessed: 2025-03-14. [Online]. Available: <https://github.com/prysmaticlabs/prysm>
- [30] Q. Zhou, H. Huang, Z. Zheng, and J. Bian, "Solutions to scalability of blockchain: A survey," *IEEE Access*, vol. 8, pp. 16 440–16 455, 2020. [Online]. Available: <https://doi.org/10.1109/ACCESS.2020.2967218>
- [31] E. R. Community, "Encrypted mempools for fairness and efficiency," *CryptoRank*, 2024, accessed: 2025-03-14. [Online]. Available: <https://tinyurl.com/mry9csjd>
- [32] M. M. Ahmed, M. Lim, X. Liu, S.-S. Ho, and N. Ivanov, "Quantifying the risk of out-of-gas conditions in ethereum smart contracts," in *2025 IEEE 4th International Conference on AI in Cybersecurity (ICAIC)*, 2025, pp. 1–8. [Online]. Available: <https://doi.org/10.1109/ICAIC63015.2025.10848774>
- [33] A. Ghaleb, J. Rubin, and K. Pattabiraman, "etainter: detecting gas-related vulnerabilities in smart contracts," in *Proceedings of the 31st ACM SIGSOFT International Symposium on Software Testing and Analysis*, ser. ISSTA 2022. New York, NY, USA: Association for Computing Machinery, 2022, p. 728–739. [Online]. Available: <https://doi.org/10.1145/3533767.3534378>