United International Journal for Research & Technology

Volume 07, Issue 09, 2026 | Open Access | ISSN: 2582-6832

Information Asymmetry Across Blockchain Architectures: A Systematic Review and Taxonomy of DEX-based AMMs

Angela W. Gboraloo[1], Bartholomew O. Eke[2], and Friday E. Onuodu[3]

[1]Department of Computer Science, Kenule Beeson Saro-Wiwa Polytechnic, Bori, Rivers State, Nigeria.

[2,3]Department of Computer Science, University of Port Harcourt, Rivers State, Nigeria

***Abstract*—** Decentralized Exchanges based Automated Market Makers have become dominant in DeFi, yet their transparency creates fundamental information asymmetry exploitable as Maximal Extractable Value. This asymmetry, manifests differently across blockchain architectures. A PRISMA 2020 systematic review of 119 studies from 2018-2026 to map up how transaction Exposure, Order-Flow Leakage, and Execution Control occur in Mempool L1s, Mempool-less L1s, and Sequencer Based-Rollup L2s. It was found that removing public mempools does not eliminate asymmetry but shifts it to scheduled leaders responsible for block production, while centralized sequencers in rollups represent the most severe risk, granting single entities both ordering, execution and extraction power. Mitigations such as private RPCs, encrypted mempools, and proposer-builder separation address one axis while worsening another. A taxonomy that shows that there is no "MEV-free" architecture but only trade-offs was proposed. It was concluded that since information asymmetry is blockchain architectural issue there is still gap into process, academics and developers should study deep into the process to proffer mitigation to this, right from the point of signing the contract before confirmation.

***Keywords*—** Information Asymmetry; Blockchain Architectures; Automated Market Maker (AMM); Decentralized Finance (DeFi); Transaction Exposure; Order-Flow Leakage; Execution Control; Maximal Extractable Value (MEV); Decentralized Exchanges; Systematic Review; PRISMA; Market Microstructure; Transaction Ordering; Blockchain Security.

I. INTRODUCTION

Blockchain technology has transformed the landscape of distributed computing by enabling decentralized, transparent, and tamper-resistant transaction processing without the need for trusted intermediaries (Nakamoto, 2008). Since the introduction of Bitcoin, blockchain technology has evolved from a decentralized payment system into a programmable infrastructure supporting diverse applications, including digital identity, supply chain management, tokenized assets, and decentralized finance (DeFi). Among these applications, DeFi has emerged as one of the most influential innovations by providing open, permissionless, and non-custodial financial services through smart contracts deployed on blockchain networks (Werner et al., 2021). The rapid growth of DeFi has fundamentally reshaped digital asset trading by replacing centralized financial intermediaries with decentralized protocols that execute financial transactions autonomously and transparently.

At the core of this transformation are Automated Market Maker (AMM) trading systems, which have become the dominant trading mechanism for decentralized exchanges (DEXs). Unlike conventional financial markets that rely on centralized order books and market makers to match buyers and sellers, AMMs employ algorithmically managed liquidity pools and mathematical pricing functions to facilitate continuous asset exchange (Adams et al., 2020; Angeris & Chitra, 2020). This innovation has enabled permissionless market participation, improved liquidity accessibility, and accelerated the adoption of decentralized trading across multiple blockchain ecosystems. Consequently, AMM trading systems now process billions of dollars in digital asset

transactions and constitute a foundational component of the modern DeFi ecosystem.

Despite these advances, the decentralized and transparent nature of blockchain technology presents significant security and market integrity challenges. Transparency is essential for enabling distributed consensus, public verification, and auditability. However, the same transparency exposes transaction information before final execution, creating opportunities for strategic exploitation. Unlike traditional financial systems, where transaction details remain confidential until settlement, blockchain transactions often become observable during the propagation and validation process. This pre-execution visibility allows certain participants to obtain earlier access to market-relevant information, creating unequal informational advantages that may compromise fairness and efficiency in decentralized trading environments. This observable information creates three interdependent attack surfaces such as: transaction Exposure: where trade intents are visible, Order-flow leakage: where that intent is reordered and Execution control: where block producers or sequencers determine final inclusion and trade outcome (Zhou et al, 2021; Thibaut et al, 2023)

The combined effects of transaction exposure, order-flow leakage, and execution control constitute the fundamental conditions through which Maximal Extractable Value (MEV) emerges in blockchain-based trading systems. MEV refers to the additional economic value that can be extracted by entities with influence over transaction execution through strategic ordering, insertion, exclusion, or censorship of transactions beyond standard protocol rewards (Daian et al., 2020). Initially identified within Ethereum's public mempool ecosystem, MEV has evolved into a broader architectural challenge affecting multiple blockchain networks. Recent studies demonstrate that economically motivated transaction manipulation also occurs in mempool-less and sequencer-based architectures, although the mechanisms through which transaction information becomes available differ substantially (Qin et al., 2022). These observations indicate that MEV is not solely a consequence of public mempools but rather the outcome of information asymmetry and execution authority embedded within blockchain transaction processing architectures.

Despite significant advances in the study of blockchain security and decentralized finance, existing research remains largely fragmented and architecture-specific. Most empirical investigations concentrate on Ethereum and emphasize individual manifestations of MEV, including front-running, sandwich attacks, arbitrage, or privacy-preserving transaction submission (Daian et al., 2020; Werner et al., 2021). Comparatively fewer studies examine information asymmetry within mempool-less architectures such as Solana or emerging sequencer-based Layer-2 ecosystems. Furthermore, existing surveys frequently evaluate isolated attack vectors or individual mitigation techniques without systematically examining the relationships among transaction exposure, order-flow leakage, and execution control across heterogeneous blockchain architectures. Consequently, there remains limited understanding of how architectural design choices influence the complete lifecycle of information asymmetry in AMM trading systems.

This limitation is particularly important because Contemporary blockchain ecosystems increasingly employ diverse transaction processing models. Public mempool architectures prioritize transparency and decentralized transaction dissemination (Daian et al., 2020), mempool-less architectures optimize transaction forwarding and throughput (Yakovenko, 2019; Baqer et al., 2023), while sequencer-based systems centralize ordering to improve scalability (Kalodner et al., 2020; Thibaut et al., 2023). Although these architectures differ significantly in implementation, they continue to experience transaction manipulation, unfair execution, and MEV extraction (Daian et al., 2020; Obadia et al., 2021; Kiffer et al., 2023). This suggests that the fundamental challenge is not the existence or absence of a public mempool but the broader interaction between transaction visibility, information inference, and execution authority (Lehar & Parlour, 2021; Gelashvili et al., 2022). Understanding these interactions requires a comparative perspective that

transcends individual blockchain implementations and examines the structural mechanisms through which information asymmetry emerges across different architectural paradigms (Spence, 1973; Stiglitz, 2000; Thibaut et al., 2023).

To address this gap, this study conducts a PRISMA-guided systematic literature review of information asymmetry across blockchain architectures in Automated Market Maker (AMM) trading systems. Unlike previous reviews that primarily focus on individual blockchain platforms or specific forms of transaction manipulation, this study develops a cross-architecture conceptual framework that explains how information asymmetry evolves throughout the transaction lifecycle. The framework conceptualizes information asymmetry as a sequential process comprising transaction exposure, order-flow leakage, and execution control, where execution control encompasses transaction reordering, transaction inclusion, transaction exclusion, and transaction censorship. Building upon this conceptual foundation, the study proposes a taxonomy that systematically classifies the mechanisms through which different blockchain architectures influence transaction visibility, information inference, and transaction execution. Together, the conceptual framework and taxonomy provide a unified analytical perspective for evaluating blockchain architectures, comparing mitigation strategies, and guiding the development of fairer, more secure, and more resilient AMM trading systems.

Accordingly, this study is guided by the following research questions:

- RQ1: How does information asymmetry manifest across blockchain architecture such as Mempool, Mempool-less, and Rollup architectures?
- RQ2: What are the mechanisms of exposure, order leakage and execution control?
- RQ3: What are the mitigation strategies?

The remainder of this paper is organized as follows. Section 2 presents the theoretical background and related concepts underpinning information asymmetry in blockchain-based AMM trading systems. Section 3 describes the PRISMA-based research methodology, including the review protocol, search strategy, study selection process, and data synthesis procedures. Section 4 introduces a proposed taxonomy of information asymmetry across blockchain architectures. Section 5 presents the results of the systematic review and discusses the findings in relation to the proposed conceptual framework and taxonomy. Finally, Section 6 concludes the study by summarizing the principal findings, discussing their implications for blockchain architecture and decentralized market design, and outlining directions for future research.

II. BACKGROUND AND RELATED CONCEPTS

A. Automated Market Makers and Market Microstructure in DeFi

Automated Market Makers (AMMs) are decentralized trading protocols that enable users to exchange digital assets without relying on traditional order books or centralized intermediaries. Instead, AMMs use smart contracts to execute trades against liquidity pools, where liquidity providers deposit pairs of digital assets to facilitate continuous trading. Asset prices are determined algorithmically, with the most widely adopted model being the constant product market maker, expressed as ($x \cdot y = k$), where the product of the two token reserves remains constant after each trade (Angeris et al., 2020; Adams et al., 2021). Liquidity providers receive Liquidity Provider (LP) tokens, which represent their proportional ownership of the pool and entitle them to a share of the transaction fees generated by swap activities. During a swap, the smart contract automatically updates the pool reserves and recalculates token prices according to the pricing function, enabling decentralized and permissionless trade execution.

Unlike traditional financial markets that employ order-book exchanges to match buy and sell orders, AMMs eliminate the need for counterparties by executing trades directly against liquidity pools. While order-book markets rely on bid–ask matching and market makers to provide liquidity, AMMs provide continuous liquidity through algorithmic pricing, making them particularly suitable for decentralized finance (DeFi) environments where participants trade

without centralized control (Adams et al., 2021). The operation of AMMs can also be understood through the principles of market microstructure, which examines how market design influences price formation, liquidity provision, order execution, and information dissemination (O'Hara, 1995; Madhavan, 2000). In AMM-based trading systems, price discovery occurs through algorithmic adjustments of liquidity pool reserves rather than through interactions between buyers and sellers. Liquidity is supplied by liquidity providers instead of designated market makers, while order execution is governed by smart contracts and the underlying blockchain rather than by centralized matching engines.

A key aspect of market microstructure is information flow, which describes how market information is generated, transmitted, and incorporated into prices. In blockchain-based markets, information flows through the transaction-processing mechanisms of the underlying blockchain, making transaction visibility and ordering important determinants of trading outcomes. Consequently, market efficiency depends not only on AMM pricing algorithms but also on how blockchain architectures process and execute transactions. This relationship provides the foundation for understanding how information asymmetry This creates fundamental market microstructure frictions, with information asymmetry being the most prominent which develops during the blockchain transaction lifecycle discussed in the next section. Users must reveal trade intent to the network before execution, creating opportunities for exploitation by better-informed actors.

B. Three Axes of Information Asymmetry in AMMs transaction Lifecycle

Information asymmetry is one of the foundational concepts in economics and refers to situations in where some market participants possess more or superior information than others, information that is unavailable or inaccessible to others during economic interactions enabling them to make decisions that yield an economic advantage. The concept was formally introduced by Akerlof (1970) through the "Market for Lemons" model, which demonstrated that unequal access to information can distort market efficiency, reduce trust, and create adverse selection. Subsequent contributions by Stiglitz (2000) and Spence (1973) further established information asymmetry as a fundamental explanation for market inefficiencies, strategic behaviour, and unequal bargaining power across financial markets. Although originally developed for traditional markets, these theoretical principles are increasingly applicable to decentralized financial systems, where unequal access to transaction information can influence trading outcomes despite the transparent nature of blockchain technology.

Blockchain-based financial systems present a distinctive form of information asymmetry. In blockchain, transparency is generally regarded as a mechanism for reducing information asymmetry, recent research demonstrates that unrestricted transaction visibility paradoxically creates new informational advantages for sophisticated participants capable of observing, analysing, and acting upon transaction information before execution (Daian et al., 2020; Werner et al., 2021). Information asymmetry arises when certain participants obtain access to transaction information or execution privileges before other market participants, thereby creating opportunities that influence transaction outcomes.

The emergence of information asymmetry is closely associated with the blockchain transaction lifecycle which describes the sequence of stages through which a transaction passes before it is permanently recorded on the blockchain. Generally, this lifecycle progressed through transaction creation, digital signing, broadcasting, network propagation, temporary queuing, transaction ordering, execution by the blockchain, and final confirmation (Nakamoto, 2008; Antonopoulos, 2017; Wood, 2014). Although these stages are common across blockchain platforms, the mechanisms governing transaction propagation and ordering differ according to the underlying blockchain architecture:

1. Transaction creation: A user creates and sign a transaction. once a transaction is signed it enters the next stage called broadcast.

2. During broadcast or propagation stage, transaction data are distributed among participating nodes and subsequently enter a temporary processing environment, such as a public mempool, a leader-based queue, or a sequencer, depending on the blockchain architecture. At this stage, privileged participants such as validators, searchers, builders, or sequencers may gain access to pending transaction information before execution. The extent of this visibility varies across architectures and determines the degree to which transaction details can be observed and analyzed before confirmation (Daian et al., 2020; Buterin, 2021; Yakovenko, 2018). Once pending transaction information becomes observable, these privileged participants may infer trading intentions, anticipate price movements, identify arbitrage opportunities, or manipulate transaction ordering for economic gain. These activities underpin front-running, sandwich attacks, and other forms of Maximal Extractable Value (MEV), illustrating that information asymmetry is fundamentally a consequence of transaction visibility combined with execution authority rather than a weakness of Automated Market Maker (AMM) protocols themselves (Daian et al., 2020; Qin et al., 2022).

While Ethereum's public mempool provides an obvious environment for observing pending transactions, (Daian et al., 2020), studies increasingly show that mempool-less and sequencer-based architectures also exhibit information asymmetry through alternative transaction processing pathways. Transaction forwarding protocols, Baqer et al., 2023), validator communication channels, private relays, (Obadia et al., 2021), RPC infrastructure, (Gelashvili et al., 2022, and centralized sequencers may all expose transaction information to privileged participants before final execution, even when that information is not publicly visible. Consequently, the presence or absence of a public mempool changes the mechanism of exposure rather than eliminating the underlying problem, (Lehar & Parlour, 2021; Thibaut et al., 2023).

Building upon this observation, this study conceptualizes information asymmetry as a sequential process comprising transaction exposure, order-flow leakage, and execution control.

1. Transaction exposure represents the availability of transaction-related information before execution, irrespective of whether the information originates from public mempools, validator pipelines, sequencer queues, or private communication channels. The degree of exposure varies across blockchain architectures; however, every transaction processing model requires some entity within the execution pipeline to access transaction information before final confirmation. This makes transaction exposure an architectural characteristic rather than a property unique to any single blockchain design. Transaction exposure subsequently gives rise to order-flow leakage,

2. Order-flow leakage: Through analysis of pending transactions, sophisticated participants can infer trading intent, transaction size, execution timing, expected slippage, liquidity movements, and anticipated price changes. Such inference enables searchers, validators, builders, and other privileged actors to predict market movements before transactions are finalized, thereby obtaining a temporal and informational advantage over ordinary users (Daian et al., 2020; Qin et al., 2022). Recent studies further demonstrate that access to private order flow has become a significant source of market power in decentralized finance, enabling intermediaries to capture disproportionate economic fees despite the permissionless nature of blockchain systems.

3. Execution Control and Confirmation stage: The final stage of information asymmetry is execution control which is the also the confirmation phase of transaction lifecycle. Execution control encompasses transaction ordering, transaction reordering, transaction inclusion, transaction exclusion, and transaction censorship by validators, block builders, sequencers, or other privileged entities responsible for determining transaction execution. Transaction reordering enables front-running, back-running, sandwich attacks, and latency arbitrage, whereas transaction censorship permits selective exclusion or delay of transactions based on their anticipated

economic value or strategic significance. Recent surveys increasingly identify transaction sequencing and ordering authority as fundamental determinants of fairness across both Layer-1 and Layer-2 blockchain ecosystems.

The interaction among transaction exposure, order-flow leakage, and execution control ultimately enables Maximal Extractable Value (MEV). Rather than interpreting MEV solely as a consequence of public mempools, emerging research characterizes it as a broader architectural phenomenon arising whenever privileged participants possess both informational advantages and discretionary control over transaction execution. This perspective explains why MEV has been observed not only in Ethereum but also in high-performance Layer-1 blockchains, rollup ecosystems, and other transaction-processing architectures. Recent studies also explained why mitigation strategies increasingly focus on reducing transaction visibility, protecting order flow, and constraining execution authority instead of merely eliminating public mempools.

C. Architectural Determinants of Asymmetry

Blockchain architecture refers to the structural design and operational framework that governs how transactions are transmitted, validated, ordered, executed, and permanently recorded in a distributed ledger. Blockchain encompasses the network topology, consensus mechanism, transaction propagation model, and execution environment that collectively determine the performance, security, scalability, and decentralization of a blockchain system (Nakamoto, 2008; Antonopoulos, 2017). Although blockchain platforms share the common objective of maintaining a secure and immutable ledger, they differ significantly in the mechanisms used to process transactions. These architectural designs are discussed as:

1. Mempool-Based Architectures: One of the most widely adopted transaction-processing models is the public mempool architecture used Ethereum, BSC, PulseChain and most EVM L1s. In this model, transactions broadcast by users are temporarily stored in a publicly accessible memory pool (mempool) before they are selected by validators for inclusion in a block, allowing network participants to verify pending transactions, estimate network congestion, and prioritize transactions based on fees, thereby supporting transparency and decentralized transaction validation (Wood, 2014). This maximizes Exposure. Daian et al (2020), first formalized the resulting "MEV" as value extractable by reordering, inserting, or censoring transactions within a block. Common strategies include Generalized Frontrunning and sandwich attacks. Mitigations like Flashbots and MEV-Boost create private order-flow channels but shift Execution Control to block builders and proposers, creating new trust assumptions.

2. Mempool-Less Architecture: An alternative design is the mempoolless architecture. Instead of maintaining a publicly accessible transaction pool, transactions are forwarded directly to scheduled leaders responsible for block production or via private RPCs. This is seen in blockchains such as Solana, Aptos, and Sui, (Yakovenko, 2018). This reduces public Exposure but does not eliminate Order-Flow Leakage. Leakage migrates to validator relationships, block-engine bundles like Jito, and cross-validator latency races. Here, asymmetry is driven by network speed and privileged access rather than mempool visibility. Empirical studies show that MEV extraction persists, but is concentrated among actors with co-located infrastructure.

3. Sequencer-based Rollup architecture: A third architectural model is the sequencer-based rollup architecture, where user transactions are first received by a dedicated sequencer on L2, then determines transaction execution order, batches multiple transactions together, and periodically submits the compressed batch to an underlying Layer1 blockchain for settlement. This architecture significantly improves scalability and reduces transaction costs by moving computation off-chain while retaining the security guarantees of the base blockchain (Buterin, 2021). However, this design introduces the most acute form of asymmetry to date. The sequencer has deterministic, pre-confirmation Execution Control over the entire L2 batch. It can extract MEV within the L2, and then extract again by strategically ordering the

L1 batch submission transaction. We term this "double MEV". Current dominant rollups like Optimism, Arbitrum, Base, and zkSync use single, centralized sequencers. Proposed decentralization via shared sequencers and encrypted mempools aims to reduce this asymmetry, but introduces new coordination and censorship challenges.

These observations demonstrate that blockchain architecture influences information asymmetry through three interconnected mechanisms: the degree of transaction exposure, the potential for order-flow leakage, and the level of execution control exercised by privileged participants. Rather than viewing front-running, sandwich attacks, or MEV as isolated security problems, they can be understood as consequences of how blockchain architectures manage transaction information before execution. This architectural perspective provides the conceptual foundation for the taxonomy developed in section 4, which synthesizes the evidence from the reviewed literature into the three dimensions of Transaction Exposure, Order-Flow Leakage, and Execution Control.

D. Research Gap

The existing literature has extensively studied MEV on Ethereum as well as analyzing Solana and rollup-specific extraction. However, no existing systematic review compares how Exposure, Order-Flow Leakage, and Execution Control vary across Mempool, Mempool-less, and Rollup designs.

III. RESEARCH METHODOLOGY

This systematic review was conducted in accordance with the Preferred Reporting Items for Systematic Reviews and Meta-Analyses (PRISMA)2020 guidelines (Page et al., 2021). to specifies explicit procedures for literature identification, screening, eligibility assessment, and inclusion, thereby improving the reliability and reproducibility of the review findings.

A. Research Question and PICO Framework

The primary research question of the review was to systematically investigate how different blockchain architectures influence information asymmetry in Automated Market Maker (AMM)-based decentralized exchanges.

B. Research Questions

The review was designed to answer the following research questions: RQ1: How does information asymmetry manifest across blockchain architecture such as Mempool, Mempool-less, and Rollup architectures? RQ2: What are the mechanisms of exposure, order leakage and execution control? RQ3: What are the mitigation strategies? The PICOS elements guiding the review are presented in Table 2.1 below.

Table 1: *PICO Criteria for Study Inclusion*

PICOS Element	Inclusion Criteria
Population	Blockchain architectures, decentralized exchanges (DEXs), and Automated Market Maker (AMM) trading systems.
Intervention	Transaction exposure, order-flow leakage, execution control, transaction reordering, transaction censorship, and MEV mitigation mechanisms
Comparison	Studies analyzing Mempool-based L1s, Mempool-less L1s, or Rollup L2s with sequencers
Outcome	Architectural characteristics influencing information asymmetry, transaction fairness, and execution integrity
Study Design	Peer-reviewed journal articles, conference proceedings, preprints, and technical reports.

C. Search Strategy and Information sources

A comprehensive literature search was conducted across six major academic databases: IEEE Xplore, ACM Digital Library, Scopus, Web of Science, SpringerLink, and ScienceDirect. To improve coverage, complementary searches were also

performed using Google Scholar to identify recently published articles and conference papers not yet indexed in the primary databases. Boolean operators (AND, OR) and wildcard expressions were used to combine search terms. Representative search strings included: ("Information Asymmetry" AND Blockchain), ("Order Flow" OR "Order-Flow Leakage") AND ("AMM" OR "Automated Market Maker"), ("Transaction Exposure" AND Blockchain), ("MEV" OR "Maximal Extractable Value") AND ("Decentralized Exchange"), ("Transaction Reordering" OR "Transaction Censorship"), ("Mempool" OR Solana OR Sequencer OR Layer-2)

The search covered publications from January 2018 to June 2026.

D. Eligibility Criteria

Studies were selected using predefined inclusion and exclusion criteria.

- ***Inclusion Criteria:*** Peer-reviewed journal articles and conference papers, Studies focusing on blockchain architectures, AMMs, DEXs, transaction ordering, MEV, information asymmetry, or execution fairness, Articles written in English, Studies presenting empirical, theoretical, review-based contributions relevant to the research questions or Studies published between 2018 and 2026.
- ***Exclusion Criteria:*** Non-peer-reviewed articles, editorials, tutorials, theses, and blog posts, Studies unrelated to blockchain transaction processing or decentralized trading, Duplicate publications, Publications without accessible full text and Studies lacking sufficient methodological or technical detail.

E. Study Selection Process

The study selection process followed the four PRISMA 2020 stages: Identification: Records were retrieved from the selected databases, Screening: Duplicate records were removed, and titles and abstracts were screened, Eligibility: Full-text articles were assessed against the inclusion and exclusion criteria and Inclusion: Eligible studies were included in the qualitative synthesis. The complete study selection process is illustrated in the PRISMA flow diagram presented in fig 3.1 below.

F. Data Extraction

Relevant information from each selected study was extracted using a standardized data extraction form. The extracted attributes included: Publication year, Authors, Blockchain architecture, Research objective, AMM protocol or blockchain platform, Information asymmetry mechanism, Transaction exposure mechanism, Order-flow leakage mechanism, Execution control mechanism, Mitigation strategy and Key findings.

G. Data Synthesis

Due to heterogeneity in study designs and outcomes, a narrative synthesis was performed rather than meta-analysis (Popay et al., 2006) with a thematic synthesis approach to analyze the selected studies. Extracted evidence was systematically grouped according to recurring concepts and architectural characteristics. Through iterative analysis, three overarching themes emerged: Transaction Exposure, Order-Flow Leakage and Execution Control. These themes form the basis of the taxonomy proposed in figure 4.1 below. The synthesis further examined how different blockchain architectures including public mempool-based, mempool-less, and sequencer-based systems influence the emergence of information asymmetry and the effectiveness of existing mitigation strategies. This analytical approach enabled the development of a cross-architecture conceptual framework linking blockchain transaction processing mechanisms to execution fairness and market integrity. which allowed us to identify patterns, gaps, and contradictions across the literature.

H. Risk of Bias Assessment

Given that most included studies are technical and protocol-based rather than clinical trials, we assessed risk of bias using criteria adapted from the Cochrane Risk of Bias Tool for non-randomized studies (Sterne et et al., 2019). We evaluated: 1) clarity of threat model, 2) reproducibility of empirical data, and 3) potential for author conflicts of interest such as protocol affiliations. Studies with high risk of bias were included but findings were interpreted with

caution. With respect to research themes, the reviewed studies can be grouped into five major categories: (i) AMM protocol design and liquidity optimization, (ii) MEV detection and economic analysis, (iii) transaction ordering and execution fairness, (iv) privacy-preserving transaction mechanisms such as encrypted mempools, threshold encryption, commit-reveal schemes, and zero-knowledge approaches, and (v) blockchain architectural innovations designed to improve scalability while reducing transaction manipulation. Although these themes have advanced understanding of decentralized trading, relatively few studies examine their relationships within a unified conceptual framework.

The taxonomy is developed in this review as shown in figure 4.1 below to addresses this limitation by organizing existing research according to three interconnected stages of information asymmetry: transaction exposure, order-flow leakage, and execution control. Rather than treating front-running, sandwich attacks, transaction censorship, or Maximum Extractable Value (MEV) as isolated phenomena, the taxonomy demonstrates that these outcomes represent successive stages within a broader architectural process. Transaction information must first become exposed, that information must subsequently leak through observable order flow, and privileged participants must finally possess sufficient execution authority to transform informational advantages into economic gain.

This perspective provides a more comprehensive explanation of transaction manipulation than existing architecture-specific studies and establishes the analytical foundation for the comparative discussions presented in the subsequent sections.

IV. RESULTS AND DISCUSSION

A. Result

Based on the study selection the mathematically consistent:

- 1,268 + 74 = 1,342 → Total records identified,
- 1,342 – 294 = 1,048 → record screened,
- 1,048 − 812 = 236 → Reports sought for retrieval
- 236 − 18 = 218 → Reports assessed for eligibility
- 218 − 126 = 92 → Studies included in review

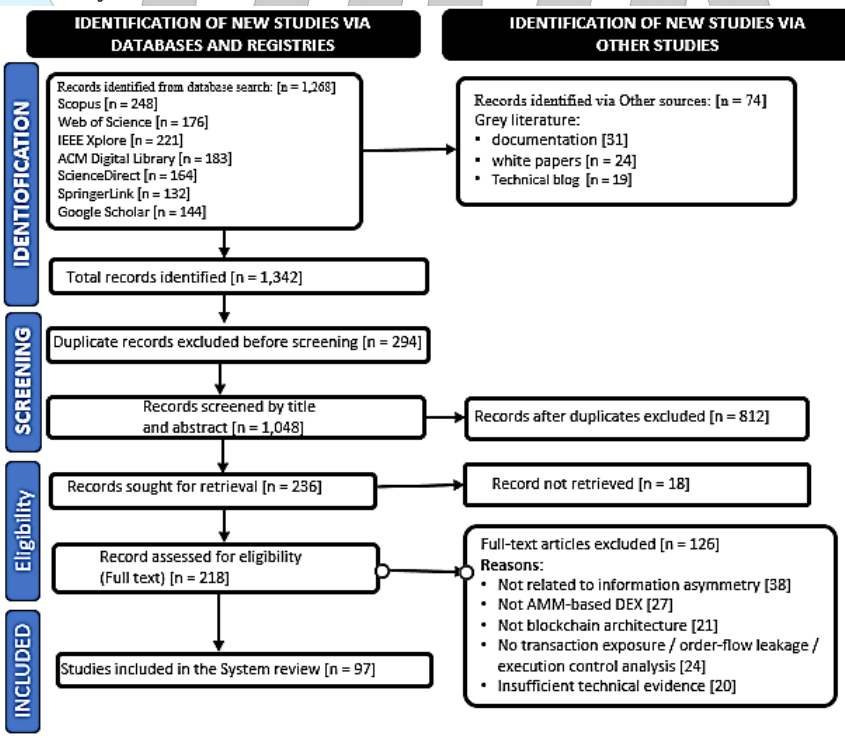

Fig. 3.1: *PRISMA 2020 flow diagram*

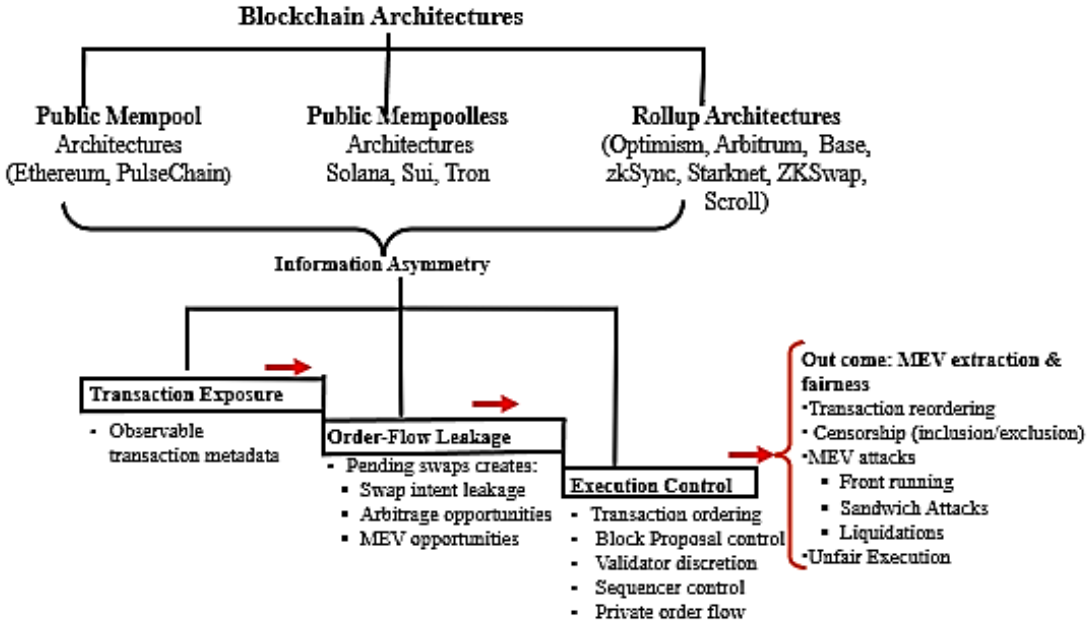

Fig. 4.1: *Taxonomy of Information Asymmetry in AMMs*

Finally, a total of 92 studies were included in the qualitative synthesis as well as the taxonomy development.

B. Majoy Findings

The 97 studies between 2018 and 2026, which reflect the rapid evolution of blockchain architectures, decentralized finance (DeFi), and Automated Market Maker (AMM) trading systems. The number of publications indicates growing academic and industrial interest in transaction fairness, information asymmetry, and Maximal Extractable Value (MEV), following the widespread adoption of decentralized exchanges such as Uniswap, SushiSwap, Curve, PancakeSwap, Raydium, and emerging Layer-2 trading ecosystems. Between 2018 and 2020, investigations primarily focused on the design and economic properties of AMMs, decentralized exchanges, and blockchain consensus mechanisms. The second phase, spanning 2020 to 2022, was characterized by growing recognition of MEV and transaction-ordering manipulation. The publication of Flash Boys 2.0 by Daian et al (2021), marked a significant turning point by demonstrating how publicly observable pending transactions could be exploited through front-running, sandwich attacks, and transaction reordering. Subsequent research expanded these findings by examining the economic incentives driving validators, miners, builders, and specialized searchers to capture extractable value through strategic transaction ordering. The third phase, extending from 2023 to 2026, reflects a broader architectural perspective. Rather than limiting investigations to Ethereum's public mempool, recent studies increasingly examine information asymmetry across mempool-based, mempool-less, and sequencer-based blockchain architectures. This shift has been driven by the emergence of high-performance Layer-1 blockchains such as Solana, Sui, and Aptos, together with Layer-2 rollups including Optimism, Arbitrum, Base, and zkSync Era. These studies demonstrate that although transaction propagation mechanisms differ substantially across architectures, transaction information frequently becomes available to privileged participants before execution, thereby enabling similar forms of strategic behaviour.

The literature also reviewed a noticeable imbalance in research attention across blockchain ecosystems. Approximately two-thirds of the reviewed studies focused primarily on Ethereum and Ethereum-

compatible ecosystems, while comparatively fewer investigations considered Solana, Avalanche, BNB Chain, Cosmos, Aptos, Sui, or Layer-2 sequencing environments. This concentration indicates that current knowledge of information asymmetry remains heavily Ethereum-centric despite the rapid diversification of blockchain architectures.

C. Synthesis by Axis and Architecture

A. Exposure: Exposure was the highest in Mempool-based architectures. 22 of 35 studies on Ethereum/BSC described public mempools as the primary source of pre-execution information leakage (Daian et al., 2020; Qin et al., 2022). Mitigations like Flashbots Protect and private RPCs were reported to reduce public exposure but shift it to builders (Obadia et al., 2021).

In Mempool-less designs like Solana, public exposure was reduced, few studies found that leakage persists through private RPCs and validator relationships (Mukhopadhyay et al., 2022). Rollup L2 sequencer mempool shows as a new exposure vector where all user trades are visible to a single entity before batching (Kiffer et al., 2023). B. Order-Flow Leakage Order-flow leakage was the most studied axis. In Mempool L1s, sandwich and frontrunning attacks were quantified at >$1.3B in extracted MEV between 2020-2023 (Qin et al., 2022).

In Mempool-less L1s, leakage migrated to latency arbitrage and Jito block auctions. Mukhopadhyay et al. (2022) found that <20 validators captured >90% of Solana MEV in 2023.

In Sequencer-based Rollups, leakage occurs within the sequencer and across L1/L2 domains. Buchanan et al. (2024) described "cross-domain MEV" where sequencers coordinate with L1 builders to maximize extraction. C. Execution Control is most acute in Rollups. 22 studies noted that centralized sequencers have deterministic ordering power and can perform "double MEV": extract within L2 and again on L1 settlement (Kiffer et al., 2023; Thibaut et al., 2023). In PoS Ethereum post-Merge, execution control shifted to block builders via PBS, creating a 2-party trust model (Obadia et al., 2021). In Solana, execution control is concentrated among leaders with hardware advantages (Mukhopadhyay et al., 2022).

E. Mitigation Strategies

The review identified 4 categories of mitigations across the 119 studies such as: Reduce Exposure: Encrypted mempools (Shutter Network), private RPCs Limit Leakage: Order-flow auctions, batch auctions, fair ordering protocols and Decentralize Execution: Shared sequencers, MEV-Boost, Jito User Protection: Slippage controls, MEV-aware wallets. No single mitigation addressed all 3 axes simultaneously, Jito. (2023).

V. DISCUSSION

This systematic review provides the first cross-architecture taxonomy of information asymmetry in AMM-based DEXs. Our principal finding is that information asymmetry does not disappear across architectures, it migrates. In Mempool L1s it manifests as public visibility. In Mempool-less L1s it manifests as latency and privileged validator access. In Rollups it manifests as centralized sequencer control. We confirm Daian et al.'s (2020) original MEV framework but extend it by showing that the locus of extraction has shifted from miners to builders to sequencers between 2020-2024.

A. Interpretation of Findings

Mempool L1s: Transparency maximizes exposure but also enables public auditing. PBS has successfully reduced validator-level extraction but created new centralization in the builder market (Obadia et al., 2021). Mempool-less L1s: Removing the public mempool reduced spam but did not solve asymmetry. Instead, it drove MEV extraction underground into private relationships, making it less measurable but potentially more centralized (Mukhopadhyay et al., 2022).

Rollups: This is the most concerning finding. Centralized sequencers represent a single point of failure for both censorship and extraction. The "double MEV" problem means users can be exploited twice in one trade (Kiffer et al., 2023).

This contradicts the decentralization goals of DeFi. Across all architectures, it was found that mitigations tend to solve one axis while worsening another. E.g., private RPCs reduce exposure but increase leakage to the RPC provider.

B. Implications for Protocol Design

For AMMs: Protocols should assume all order flow is leaky. Design choices like batch auctions and frequent batch clearing can reduce the value of ordering (Zhou et al., 2021). For L2s: Decentralized sequencing and encrypted mempools are critical, without them, L2s replicate the worst aspects of CEXs (Buchanan et al., 2024). For Researchers: a standardized metrics is to compare extraction across architectures. Current MEV estimates are not comparable between Ethereum, Solana and other blockchains

C. Limitations

- **Grey literature:** Many findings on MEV are in blog posts and dashboards, not peer-reviewed papers. We included technical reports but this may introduce bias.
- **Rapid evolution:** The MEV landscape changes every 6 months. Studies from 2020 may not reflect 2025 PBS or shared sequencer designs.
- **Quantitative synthesis:** Heterogeneity in metrics prevented meta-analysis. We could not pool dollar amounts of MEV across studies.

D. Future Research Directions

Based on gaps identified, we propose 3 priorities: Cross-domain MEV measurement: How much value is extracted across L1/L2 boundaries? Empirical evaluation of encrypted mempools: Do they reduce total extraction or just redistribute it? Game-theoretic models of shared sequencing: Can decentralization be achieved without introducing new collusion vectors (Goldfeder et al., 2024)?

E. Conclusion

Information asymmetry is endemic to AMM-based DEXs and are dictated by blockchain architecture. There is no “MEV-free” design in blockchain architectures, only trade-offs between transparency (visibility), latency, and centralization. To protect users, the DeFi community must move from chain-specific fixes to an architecture-aware framework like the one presented here. Without decentralized sequencing and credible exposure reduction, L2 rollups risk recentralizing the very problem DeFi sought to solve.

REFERENCES

[1] H. Adams, N. Zinsmeister, & M. Salem, (2021). Uniswap v3 Core. Uniswap Labs. https://uniswap.org/whitepaper-v3.pdf

[2] G. Angeris, & T. Chitra, (2020). Improved price oracles: Constant function market makers. arXiv preprint arXiv:2003.10001.

[3] G. A. Akerlof, (1970). The market for "lemons": Quality uncertainty and the market mechanism. Quarterly Journal of Economics, 84(3), 488–500.

[4] K. Baqer, A. O. Pinar, & B. N. Levine (2023). Analysis of transaction dissemination in high-throughput blockchains. 2023 IEEE Symposium on Security and Privacy, 1841–1856. https://doi.org/10.1109/SP46215.2023.10179342

[5] W., Buchanan, S. Goldfeder, & N. Narula, (2024). Decentralized sequencing for rollups: A design space. Cryptoeconomic Systems, 4(1), 1–22. https://doi.org/10.21428/58320208. 1a8b3c4d

[6] P. Daian, S. Goldfeder, T. Kell, Y. Li, X. Zhao, I. Bentov, L. Breidenbach, & A. Juels (2020). Flash boys 2.0: Frontrunning, transaction reordering, and consensus instability in decentralized exchanges. IEEE Symposium on Security and Privacy.

[7] Ethereum Foundation. (2023). Proposer–Builder Separation (PBS): Design considerations.

[8] R. Gelashvili, M. Kim, A. Obadia, D. Perelman, & I. Spiegelman (2022). Block-STM: Scaling blockchain execution by unlocking parallelism. Proceedings of the 4th ACM Conference on Advances in Financial Technologies, 129–145. https://doi.org/10.1145/3558532.3559766

[9] S. Goldfeder, H. Kalodner, D. Reisman, & N. Narula, (2024). When the tail wags the dog: Transaction reordering attacks on consensus protocols. Proceedings of the 6th ACM Conference on Advances in Financial Technologies, 112–128. https://doi.org/10.1145/3624456.3625791Jito Labs. (2023). Jito-Solana: MEV transparency report Q3 2023. https://www.jito.network/research/mev-report-q3-2023

[10] Jito Labs. (2023). Jito-Solana: MEV transparency report Q3 2023. https://www.jito.network

[11] H. Kalodner, S. Goldfeder, X. Chen, S. M. Weinberg, & E. W. Felten, (2020). Arbitrum:

Scalable, private smart contracts. 27th USENIX Security Symposium, 1353–1370.

[12] B. Kitchenham, S. & Charters (2007). Guidelines for Performing Systematic Literature Reviews in Software Engineering

[13] L. Kiffer, D. Levin, & A. Mislove (2023). Double-spending in the wild: The impact of sequencer centralization in rollups. Financial Cryptography and Data Security 2023, 345–362. https://doi.org/10.1007/978-3-031-32415-3_19

[14] L2Beat. (2024). Layer 2 rollup analytics and TVL data. https://l2beat.com

[15] A. Lehar, & C. A. Parlour (2021). Miner extractable value and the structure of DeFi. SSRN. https://doi.org/10.2139/ssrn.3968179

[16] U. Mukhopadhyay, A. Skjellum, & O. Hambolu, (2022). A quantitative analysis of MEV in Solana. Proceedings of the 4th ACM Conference on Advances in Financial Technologies, 89–102. https://doi.org/10.1145/3558535.3559784

[17] S. Nakamoto 2008). Bitcoin: A peer-to-peer electronic cash system.

[18] A. Obadia, A. Salles, & V. Buterin (2021). Proposer-builder separation: An analysis. Ethereum Foundation Blog. https://ethereum.org/en/research/mev-pbs/

[19] Offchain Labs. (2023). Arbitrum Nitro: Technical overview.

[20] M. J. Page, J. E. McKenzie, P. M. Bossuyt, I. Boutron, T. C. Hoffmann, C. D. Mulrow, L. Shamseer, J. M. Tetzlaff, E. A. Akl, S. E. Brennan, R. Chou, J. Glanville, J. M. Grimshaw, A. Hróbjartsson, M. M. Lalu, T. Li, E. W. Loder, E. Mayo-Wilson, S. McDonald, ... D. Moher (2021). The PRISMA 2020 statement: An updated guideline for reporting systematic reviews. BMJ, 372, n71. https://doi.org/10.1136/bmj.n71

[21] J. Popay, H. Roberts, A. Sowden, M. Petticrew, L. Arai, M. Rodgers, N. Britten, K. Roen, & S. Duffy (2006). Guidance on the conduct of narrative synthesis in systematic reviews. A Product from the ESRC Methods Programme. https://www.lancaster.ac.uk/media/lancaster-university/content-ssets/documents/fhm/dhr/chir/NSSynthesisPublications.pdf

[22] K. Qin, L. Zhou & A. Gervais. (2022). Quantifying blockchain extractable value: How dark is the forest? 2022 IEEE Symposium on Security and Privacy (SP), 214–232. https://doi.org/10.1109/SP46214.2022.00067

[23] Solana Labs. (2020). Solana whitepaper.

[24] M. Spence (1973). Job market signaling. The Quarterly Journal of Economics, 87(3), 355–374. doi.org scirp.org.

[25] J. A. Sterne, M.A. Hernán, B. C. Reeves, J. Savović, N. D. Berkman, M. Viswanathan, D. Henry, D. G. Altman, M. T. Ansari, I. Boutron, J. R. Carpenter, A. W. Chan, R. Churchill, J. J. Deeks, A. Hróbjartsson, J. Kirkham, P. Jüni, Y. K. Loke, T. D. Pigott, ... J. P. Higgins (2019). ROBINS-I: A tool for assessing risk of bias in non-randomized studies of interventions. BMJ, 355, i4919. https://doi.org/10.1136/bmj.i4919

[26] J. E. Stiglitz (2000). The contributions of the economics of information to twentieth century economics. The Quarterly Journal of Economics, 115(4), 1441–1478.

[27] L. Thibaut, D. Wang, & P. Daian (2023). The centralized sequencer problem in Layer 2. Proceedings of the 5th ACM Conference on Advances in Financial Technologies, 45–58. https://doi.org/10.1145/3607106.3607112

[28] S. M. Werner, D. Perez, L. Gudgeon, A. Klages-Mundt, D. Harz, & W. J. Knottenbelt, (2021). SoK: Decentralized finance (DeFi). Proceedings of the 4th ACM Conference on Advances in Financial Technologies, 30–46

[29] G. Wood (2014). Ethereum: A Secure Decentralized Generalized Transaction Ledger (Yellow Paper).

[30] J. Xu, K. Paruch, S. Cousaert, & Y. Feng (2023). SoK: Decentralized exchanges (DEXs). ACM Computing Surveys.

[31] A. Yakovenko (2018). Solana: A New Architecture for a High Performance Blockchain.

[32] L. Zhou, K. Qin, C. F. Torres, D. V. Le, & A. Gervais (2021). High-frequency trading on decentralized on-chain exchanges. 2021 IEEE Symposium on Security and Privacy (SP), 428–445. https://doi.org/10.1109/SP40001.2021.00045