Market Quality of Automated Market Makers

by Luke Johnson

Thesis submitted in fulfilment of the requirements for the degree of

Doctor of Philosophy

under the supervision of Tālis J. Putniņš
and Ester Félez-Viñas

University of Technology Sydney
Faculty of Business

March 2026

Certificate of Original Authorship

I, Luke Johnson, declare that this thesis is submitted in fulfilment of the requirements for the award of Doctor of Philosophy, in the Finance Discipline Group at UTS Business School at the University of Technology Sydney.

This thesis is wholly my own work unless otherwise referenced or acknowledged. In addition, I certify that all information sources and literature used are indicated in the thesis.

This document has not been submitted for qualifications at any other academic institution.

This research was supported by an Australian Government Research Training Program (RTP) Scholarship doi.org/10.82133/C42F-K220.

Production Note:
Signature: Signature removed prior to publication.

Date: 9/3/26

Working Papers

Chapters 2-5 of this thesis have each been developed as standalone working papers. Below is a list of the working papers and conference presentations:

1. Félez-Viñas, E., Johnson, L., & Putniņš, T. J. (2025c). Liquidity That Lasts: Learning from AMMs for Market Liquidity. *Working paper* is based on Chapter 2 and has been presented at the following seminars:

   - 2023 UTS Business School doctoral presentation, Sydney, Australia
   - 2024 Digital Finance CRC Pizza and Papers Seminar Series. Sydney, Australia

2. Félez-Viñas, E., Johnson, L., & Putniņš, T. J. (2025a). Can Automated Market Makers Provide Price Discovery? *Working paper* is based on Chapter 3 and has also been circulated under the name *"Active vs Passive: Price Discovery in Automated Market Makers."* This paper has been presented at the following conferences and seminars:

   - 2024 Digital Finance CRC Pizza and Papers Seminar Series. Sydney, Australia
   - 6th Sydney Market Microstructure and Digital Finance Meeting. Sydney, Australia
   - 6th University of Western Australia Blockchain and Cryptocurrency Conference. Perth, Australia
   - 37th Australasian Finance & Banking Conference. Sydney, Australia

3. Chang, I., Johnson, L., & Putniņš, T. J. (2025). Misconduct in Decentralized Exchanges. *Working paper* is based on Chapter 4 and has been presented at the following seminar:

   - 2025 Digital Finance CRC Pizza and Papers Seminar Series. Sydney, Australia

4. Félez-Viñas, E., Johnson, L., & Putniņš, T. J. (2025b). Insider trading in cryptocurrency markets. *Working paper* is based on Chapter 5 and is currently under review at the Journal of Financial and Quantitative Analysis and has been presented at the following conferences and seminars:

- 2023 Digital Finance CRC Pizza and Papers Seminar Series. Sydney, Australia
- 2023 University of Western Australia Blockchain and Cryptocurrency Research Centre seminar. Perth, Australia
- 2024 FMA Asia Pacific Conference. Seoul, South Korea.
- 2024 World Finance Conference. Nicosia, Cyprus

Acknowledgements

I would like to express my deepest gratitude to my academic supervisors, Prof. Tālis Putniņš and Dr. Ester Félez-Viñas, for their invaluable guidance, support, and collaboration throughout this thesis. Their expertise and thoughtful feedback have significantly enhanced the quality and clarity of my work.

I am also thankful to my industry partner, CryptoCompare (CCData, CoinDesk Data), for the opportunity to learn, collaborate, and travel to London. I am particularly grateful to my industry supervisors, Dr. Alex Wendland and Constantine Tsavliris, for their mentorship and support. I extend my appreciation to my colleagues in London, Charles Hayter, Vlad Cealicu, Toby Winterflood, Steve Hewson, Lucas Boyles, Nenko Nenov, Hossam Mahmoud, Jacob Joseph, Joshua de Vos, James Webb, Jamie Sly, Alison Clare, Jimena Leon, and many others for their warm welcome, insights, and friendship.

This work was supported by funding from the Digital Finance CRC, as well as financial assistance from the UTS Graduate Research School and the UTS Business School. I also thank Dr. Kristoffer Glover and A/Prof. Christina Nikitopoulos Sklibosios for their support in their roles as Finance PhD Coordinator.

Chapters of this thesis have benefited from valuable discussions and feedback from Prof. Sean Foley, Prof. Roman Kozhan, A/Prof. Katya Malinova, and Prof. Lin William Cong. I am also grateful to my PhD colleagues and friends for helpful discussions, William Krekel, Connor Arthur, Dustin Weiss, Robert Gaudiosi, Régnier Avice, Atiqur Rasel, Khanh Nguyen, Tra Nguyen, Maximilian Brandl, and Franziska Gmeiner.

Finally, I would like to thank my family. My parents, my brother, and my partner, Boo, for their unwavering support and love. This thesis would not have been possible without them.

Abstract

Automated Market Makers (AMMs) promise to be the next evolution in market design, offering a new approach to trading, enabling passive liquidity provision and automated price discovery governed by mathematical functions. This dissertation examines the market quality of AMMs by analyzing their liquidity, price discovery, and market integrity, assessing whether AMMs will be the next step in market design.

Chapter 2 examines the liquidity characteristics of AMMs. Despite structural differences in market design, we show that standard measures of liquidity—including spreads and depth—apply to both AMMs and Limit Order Books (LOBs). Comparing these markets, we find that AMM liquidity is significantly more stable than LOB liquidity, as AMM orders persist rather than being removed upon execution.

In Chapter 3 we investigate whether AMMs can facilitate efficient price discovery, given that prices are determined by a mathematical function rather than by Liquidity Providers (LPs). Applying a variance decomposition model to separate information from noise in prices, we find that when AMMs are the only market to trade an asset they produce relatively efficient and informative prices. However, when they compete against a LOB trading the same asset, the LOB typically contributes a larger share of the price discovery. AMMs contribution to price discovery increases when they have lower trading costs.

Chapter 4 explores misconduct in Decentralized Exchanges (DEXs) that trade through AMMs. We identify $345 billion in trading volume linked to market misconduct on DEXs—approximately 12% of all DEX volume. We develop a novel market integrity index that tracks four key forms of market misconduct: rug pulls, sandwich attacks, wash trading, and money laundering. Misconduct is widespread but varies by blockchain and DEX design. Despite the scale of misconduct identified, we find that market integrity issues do not appear to deter users from trading on DEXs.

Chapter 5 explores the role that DEXs and AMMs have in the wider market integrity of cryptocurrency markets. We find evidence of systematic insider trading, where individuals use private information to buy coins prior to exchange listing announcements. We identify the specific transactions and wallets (individuals) that consistently trade before announcements, ruling out alternative explanations. Our estimates suggest that 28-48% of cryptocurrency listings involve insider trading, resulting in at least $30 million in trading profits.

Overall, the findings of this thesis suggest AMMs offer a viable market structure, with the potential to improve liquidity resilience in traditional markets. However, the findings of this dissertation underscores the substantial regulatory challenges posed by DEXs, highlighting vulnerabilities to market abuse and the difficulty of enforcement in pseudonymous blockchain environments.

Contents

List of Figures

List of Tables

Acronyms & Abbreviations

AML Anti-money Laundering

AMM Automated Market Maker

BSC Binance Smart Chain

CCCAGG Crypto Coin Comparison Aggregated Index

CEX Centralized Exchange

CL-CPMM Concentrated Liquidity Constant Product Market Maker

CPMM Constant Product Market Maker

DeFi Decentralized Finance

DEX Decentralized Exchange

DOJ Department Of Justice

ESMA European Securities and Markets Authority

HFT High Frequency Trading

KYC Know-Your-Customer

LOB Limit Order Book

LP Liquidity Provider

MEV Maximal Extractable Value

MVDA MVIS® CryptoCompare Digital Assets 100 Index

NFT Non-fungible Tokens

OLS Ordinary Least Squares

PBS Proposer Builder Separation

SEC Securities and Exchange Commission

TVL Total Value Locked

U.S. United States

VAR Vector Autoregression

VECM Vector Error Correction Model

VMA Vector Moving Average

Chapter 1

Introduction

Take a view of the Royal Exchange in London, . . . where the representatives of all nations meet for the benefit of mankind.

Voltaire (1734)

1.1 A brief history of trading and markets

When we think about what makes us human, we often think about the biological features, the way we look, and how we walk. One aspect that is often overlooked is the fact that, as humans, we are the only creature to trade. In the words of Adam Smith (1776), *"Man is an animal that makes bargains: no other animal does this—no dog exchanges bones with another."* We trade because it makes us better off. For an individual, trading enables the movement of resources through space and time and helps facilitate a transfer of risk to those best equipped to bear it. On a broader economic level, trading plays a role in allocating and pooling resources for their most productive uses, it helps generate information to guide decision-making (Levine, 2005). Collectively, these advantages are known as the gains from trade.

How we trade is crucial in realizing the gains from trade. The easier it is to trade and the fewer frictions and costs that traders face, the greater the benefits for both traders and the broader economy. Markets—as a venue which facilitate trade—play a central role in reducing these frictions. Historically, markets have served as hubs that bring participants together, making trade more efficient. Evolutions in trading and markets, driven by advancements in technology, often lead to greater gains from trade.

Early traders, relied on direct barter, exchanging goods without a common medium, such as trading barley for a cow. The first major evolution in trade came with the introduction of money, which served as an intermediary in exchanges. The earliest known currency or form of money is the Mesopotamian shekel, dating back to 2150 BC. Money allowed traders to store the value of perishable goods through time, capturing value that would otherwise be lost.

However, storing and transporting physical money proved impractical. Even today, carrying a pocket full of coins is cumbersome. It is the storage cost associated with physical money that acts as a friction in trading. To reduce these storage costs, the first known forms of paper money and promissory notes emerged during the Song Dynasty in 11th-century China.

In 1531, the Antwerp Bourse became the first purpose-built commodity exchange, setting a model for market design that would persist for the next 500 years. It provided a central location for traders to gather, functioning similar to the open outcry markets of the 20th century. This innovation in markets came as the world was growing rapidly based on the advancements in seafaring. Creating a central place to trade helped reduce the cost associated with finding other traders.

Early markets relied heavily on intermediaries, such as auctioneers, who charged fees on all transactions. The founding of the New York Stock Exchange and the signing of the Buttonwood Agreement of 1792 marked a pivotal shift in market structure. Signed by 24 brokers, the agreement stipulated that they would trade only with one another and capped commissions at 0.25%, effectively reducing reliance on auctioneers and lowering trading costs. The success of the New York Stock Exchange came from its ability to consolidate liquidity into a central market built on trust between brokers.

The most recent advancement in markets has come from innovations in communication technology, particularly with the rise of electronic markets. Today, most markets operate through a Limit Order Book (LOB) model, where traders electronically submit orders specifying their intention to trade at a particular price and quantity. In a LOB market, buyers and sellers are matched based on agreed prices and quantities ordered based on a price time priority. Further advancements in communication technology have led to the rise of algorithmic and high-frequency trading, where traders leverage the speed of execution to profit from inefficiencies in the market.

Automated Market Makers (AMMs) promise to be the next step in the evolution of trading and markets. They offer a new way of thinking about market structure and trading dynamics by facilitating passive liquidity provision and automating price discovery based on a mathematical function.

The focus of this dissertation is to examine whether AMMs will be the next evolution in trading and markets by assessing the market quality of AMMs. Assessing market quality is like assessing the health or performance of a market. We look at a market's liquidity—how cheap and easy it is to trade; price discovery—how well the market incorporates information into prices; and market integrity—how fair the market is for all participants. By analyzing market quality, we can better understand how AMMs can improve modern markets and trading assessing whether AMMs are the next evolution in trading and markets.

1.2 A non-technical introduction to automated market makers

The main innovation of AMMs is the introduction of passive liquidity provision. This design feature reduces the costs of monitoring trading strategies. On blockchains, AMMs have also significantly reduced the cost associated with creating a market, making it feasible to trade assets that would otherwise not be traded. By expanding the scope of what can be traded, AMMs can help unlock new gains from trade.

AMMs are different to traditional markets, introducing a completely new market structure based on a pool of assets instead of orders in a LOB. Take the example of a farmer selling apples at a market. Instead of sitting around all day serving customers, the farmer could create a self-service market by adding "liquidity" in the form of both apples and cash on either side of a balance scale. However, this scale does not measure weight but liquidity (the amount of apples and cash that can be exchanged at this market). This setup frees up the farmer's time to focus on other activities. Customers can buy apples by following a simple rule: keep the scale balanced. To purchase apples, they need to leave cash. The fewer apples left on the scale the more money has been left to keep it balanced. This self-service stand operates like an AMM—where the market automatically manages trades without requiring constant oversight.

Enforcing the AMMs rules is necessary in practice. In the example above, it would be counterintuitive for the farmer to watch his own stand as it would not free his time to be spent elsewhere. Moreover, there would be no incentive for the farmer to provide this service, often for providing this service a small fee is paid for every trade. In Decentralized Finance (DeFi), where AMMs are a popular market design, smart contracts enforce these rules on blockchains, ensuring that trades follow the predefined logic without the need for oversight from any individual. While the mechanics of an AMM may differ in practice (like how prices are defined), the fundamental concept of automated trading and liquidity provision remains the same. For a detailed technical overview of AMMs in DeFi we refer to Chapter 2.

1.3 Market quality

1.3.1 Liquidity

Liquidity in financial markets reflects how much it costs to trade. In a highly liquid market, participants can enter or exit large positions quickly, with minimal impact on prices. Typically, traders face a trade-off between the speed of execution and price: selling quickly often requires accepting a worse price. However, in liquid markets, this trade-off is minimal.

Liquidity is often measured based on three dimensions: width, depth, and immediacy (Harris, 2003). Width refers to the difference between the transaction price and the asset's fundamental value. In LOB markets, this is typically measured by the bid-ask spread, which captures the cost of immediate execution. Depth reflects the volume that can be traded without significantly impacting prices—markets with greater depth can absorb larger trades with minimal price movement. Immediacy refers to the speed at which trades can be executed at prevailing prices, and is closely related to resiliency, which captures how quickly prices revert to equilibrium following a large trade or market shock.

In Chapter 2, we show that traditional measures of liquidity from LOB markets can be applied to AMMs, even though AMMs have a completely different market structure. Traditional limit orders can be thought of as free options to trade at a price extended to other traders (Foucault, Kadan, & Kandel, 2005). AMMs replace this mechanism with prices adjusted automatically based on a deterministic function, providing the option to trade at a price that does not expire when traded against.

The feature of persistent orders in AMMs highlights the potential benefits of AMM liquidity provision. We show that AMM liquidity is more stable compared with traditional

LOB markets, which is particularly beneficial during periods of market stress or turmoil. We also outline several hybrid market designs where AMMs complement traditional market structures, to enhance the resilience of liquidity in modern markets.

1.3.2 Price discovery

Central to the Efficient Market Hypothesis is the idea that asset prices reflect information (Fama, 1970). Price discovery is the process through which information enters prices, and price efficiency refers to how well information is reflected in prices. A key function of financial markets is to produce and incorporate information, which has wide-reaching implications—from allocating capital to its most productive uses, to mitigating adverse selection by providing fair and accurate prices, to guiding corporate decision making through price signals.

In Chapter 3, we investigate whether AMMs can facilitate efficient price discovery. The passive liquidity provision in AMMs means that information enters prices solely through the actions of liquidity demanders. This structure exposes Liquidity Providers (LPs) to adverse selection and raises important questions about the efficiency of AMM pricing when they are the only market to trade an asset.

Through a variance decomposition we measure the amount of information and noise in AMM prices. We estimate the proportion of public, private, and market-wide information in prices for assets trading only on AMMs and for those trading on both AMMs and LOB markets. We find that AMMs can provide relatively efficient price discovery even when they are the only venue to trade an asset.

We also assess whether AMMs can lead price discovery compared to LOBs. When AMMs compete with a LOB trading the same asset, we find that the LOB tends to lead the price discovery process. However, AMM contribution to price discovery increases when they are more liquid or when the LOB market is less liquid. These findings suggest that given sufficient volume AMMs are not only a viable market structure for trading but also contribute to information generation.

1.3.3 Market integrity

Market integrity is often the most overlooked dimension of market quality. Yet, it plays a critical role in fostering trust, encouraging participation, and ensuring the fairness of financial markets. Poor market integrity can distort price discovery and erode liquidity. Despite its importance, empirical research on market integrity remains limited, largely because of the challenges in measuring manipulation and misconduct (Austin, 2017). Many studies rely on prosecuted cases for validation, introducing a prosecution bias.

Decentralized Exchanges (DEXs), built on blockchain technology that leverage AMMs, offer a great opportunity for studying market integrity. The transparency of on-chain data allows researchers and practitioners to detect and quantify forms of misconduct that are otherwise hidden in traditional financial markets.

Chapter 4 investigates the scale and nature of market misconduct in AMM-based DEXs. We identify approximately $345 billion in transaction volume associated with market misconduct in DEXs. Our analysis captures both familiar forms of misconduct, such as wash trading and money laundering, and forms of misconduct unique to AMMs, such as sandwich attacks and rug pulls. We estimate that roughly 12% of total DEX volume is linked to these illicit activities. Our findings suggest that despite the scale of misconduct in DEXs, market integrity issues do not appear to deter users from trading on DEXs.

Chapter 5 examines the role of DEXs and AMMs in the broader integrity of cryptocurrency markets. We study the first prosecuted case of insider trading related to token listing announcements and find evidence of systematic insider trading, supported both by blockchain-based trade patterns and by traditional price run-up metrics consistent with prior literature. We observe that insiders exhibit a strong preference for trading on DEXs, likely due to their permissionless and pseudonymous nature. We find that in response to regulation, insiders do not cease their activity, instead shifting to methods to conceal their behavior.

Together, our findings highlight the regulatory challenges posed by blockchain-based markets. While these platforms remain difficult to police in conventional ways because of the unclear regulatory boundaries and challenges with permissionless access, the availability of transparent, high-quality data enables the detection and monitoring of market misconduct in these markets. However, the findings from Chapter 5 suggest that rather than ceasing illicit behavior, insiders appear to adapt their concealment strategies—emphasizing the need for new surveillance frameworks tailored to DEXs and blockchain-based markets more broadly.

1.4 Thesis outline

The remaining of this thesis is structured around the following research questions:

- Is AMM liquidity more stable? (Chapter 2)
- Can AMMs provide efficient price discovery? (Chapter 3)
- How much market misconduct occurs in DEXs? (Chapter 4)
- Is there insider trading on cryptocurrency listing announcements? (Chapter 5)

Each chapter has its own section dedicated to related literature. Chapter 6 concludes by summarizing the key findings and implications while providing avenues for future research.

Chapter 2

Is Automated Market Maker Liquidity More Stable?

2.1 Introduction

Automated Market Makers (AMMs) have emerged as a core innovation in Decentralized Finance (DeFi), offering a novel approach to trading assets. In contrast to traditional Limit Order Book (LOB) markets, where trades occur by matching buy and sell orders at agreed prices, AMMs use mathematical pricing function and liquidity pools to enable trading. These pools, funded by passive Liquidity Providers (LPs), act as the counterparty to all trades, allowing users to swap one asset for another directly against the pool rather than waiting for a matching order. This fundamental shift in market design raises an important question: Could lessons from AMMs help improve modern financial markets?

A key feature of LOB markets is that liquidity is temporary—once a limit order is filled, it disappears from the book, requiring LPs to constantly replenish their orders to keep the market liquid. In contrast, AMMs enable trading against a standing pool of assets rather than individual orders. In these liquidity pools, each trade adjusts the relative balances of the assets according to a predefined pricing formula. Liquidity isn't removed, it is continuously rebalanced as trades alter the composition of the pool's reserves (Adams et al., 2021; Zhang, Chen, & Park, 2018). This fundamental difference suggests that AMMs may offer more stable and resilient liquidity, particularly during periods of high volatility or market stress, when liquidity is most needed.

In this chapter, we examine the stability of liquidity in AMMs compared to traditional LOB markets. Despite their fundamental differences in design, such as passive liquidity provision, continuous pricing, and risk sharing of LPs in AMMs, AMMs and LOBs share important characteristics that make a meaningful comparison possible using standard liquidity measures. We show that commonly used metrics in LOB markets—such as quoted, effective, and realized spreads, as well as market depth and low-frequency liquidity measures—can be adapted to AMMs, enabling a direct comparison between these new markets and traditional LOB structures.

We use standard liquidity measures to test whether AMM liquidity is more stable than that of LOB markets, aiming to assess whether AMM design can inform improvements in modern market structure. Our findings show that liquidity in AMMs, specifically market depth, is significantly more stable than LOB liquidity. The key difference lies in that AMM liquidity does not need to be constantly replenished, unlike in LOBs where liquidity is removed as orders are executed. However, we find that effective spreads exhibit greater variability in AMMs, likely due to larger average trade sizes, which contribute to higher trading costs.

Moreover, we find that several factors influence the stability of AMM liquidity, including the overall pool size, the specific AMM design, and characteristics of the underlying blockchain infrastructure. More liquid pools tend to exhibit lower variability in liquidity, while larger market activity is associated with more fluctuations in the liquidity of the AMM.

The resilience of AMM liquidity offers promising insights for traditional financial markets, especially during periods of stress. Incorporating AMM-like mechanisms into existing market structures could help improve liquidity resilience and reduce the impact of liquidity shortages in modern markets.

We propose several approaches for integrating AMMs into traditional market settings. First, without altering existing infrastructure, market makers could adopt AMM-style pricing models within LOB markets, dynamically adjusting their order placement to mimic the liquidity provision of AMMs. As market makers' sell orders execute, they replace them with new buy orders, and as buy orders execute, they replace them with new sell orders, mimicking AMM-style liquidity provision. However, this approach still relies on continuous order replenishment. Second, AMMs could be enhanced with limit order functionality, offering greater flexibility in execution strategies while combining the continuous liquidity of AMMs with the benefits of limit orders. However, these markets remain vulnerable to exploitation via blockchain-level information leakage, commonly referred to as Maximal Extractable Value (MEV). Finally, a hybrid approach could introduce AMM-style order types into centralized LOB markets, combining the benefits of both systems while avoiding blockchain-specific risks such as MEV.

Among these solutions, the hybrid model where LOBs incorporate AMM-style orders, offers benefits that go beyond liquidity stability. By introducing a new order type that continuously supplies liquidity, this approach could reshape market microstructure in ways that discourage exploitative High Frequency Trading (HFT) strategies. Prioritizing stability over speed in order execution could help build more robust and resilient financial markets. As financial markets continue to evolve, we propose that integrating AMM-style order types could lead to more stable and resilient liquidity in markets.

This chapter contributes to the ongoing discussion on market design by showing how AMMs can complement and improve traditional market structures. The idea of automated markets dates back to Black (1971), who envisioned financial markets operating without human intermediaries, a vision that closely aligns with the modern implementation of AMMs. More recently, Malinova and Park (2023) argue that transitioning equity markets to AMM-based mechanisms could generate billions in cost savings by reducing transaction costs. Similarly, Foley, O'Neill, and Putniņš (2023) examine a range of markets across asset classes and propose that AMMs may deliver efficiency gains for high-volume, low-volatility assets. We add to this literature by highlighting that AMMs and LOBs can be integrated within a single market structure to capture the strengths of both designs.

2.1.1 Related literature

This chapter builds on the literature measuring liquidity across different market structures. The finance literature has long explored alternative liquidity measures to address data limitations and computational challenges. Goyenko, Holden, and Trzcinka (2009) compare different low-frequency proxies for true liquidity, while Fong, Holden, and Trzcinka (2017) develop computationally efficient liquidity measures that retain accuracy. Hasbrouck (2004) further highlights the constraints posed by data limitations in liquidity measurement. We extend this line of research by developing liquidity metrics tailored to AMMs that facilitate direct comparisons with LOBs, offering a more comprehensive understanding of liquidity dynamics across different market structures.

This chapter contributes to the growing body of research comparing AMMs and LOB markets, highlighting that AMM liquidity is more stable when compared to LOB liquidity. Aoyagi and Ito (2021), Barbon and Ranaldo (2022), and Lehar and Parlour (2025) analyze the differences between the two market designs, focusing on aspects of market quality and efficiency. Furthermore, research on liquidity provision in AMMs by Aquilina et al. (2024), Heimbach, Wang, and Wattenhofer (2021), and Neuder et al. (2021) investigates the incentives and profitability of LPs. Capponi, Jia, and Yu (2023) and Chapter 3 of this dissertation examine the role of AMMs in price discovery.

Finally, this chapter also relates to studies on liquidity resilience in LOB markets. Theoretical models by Foucault, Kadan, and Kandel (2005) and Roşu (2009) explore the dynamics of liquidity provision in order book markets, while empirical studies such as Degryse et al. (2005), Large (2007), and Lo and Hall (2015) examine how LOB liquidity responds to shocks across a range of markets. We add to this literature by introducing a novel market design based on AMMs which could enhance the stability of liquidity in markets.

Measuring liquidity has broad implications across finance, including its role in price efficiency and asset pricing. Amihud and Mendelson (1986) show that illiquidity has a risk premium in asset prices, while Chordia, Roll, and Subrahmanyam (2008) provide evidence to suggest that liquidity affects how quickly and accurately prices incorporate information. These findings motivate the development of reliable liquidity measures that capture trading costs and market depth in AMMs.

Finance literature proposes many liquidity measures to address data limitations and computational constraints. Goyenko, Holden, and Trzcinka (2009) benchmark low-frequency liquidity proxies against high-frequency measures and show that several proxies capture key dimensions of trading costs. Similarly, Fong, Holden, and Trzcinka (2017) identify computationally efficient measures that reflect both percentage trading costs and costs per unit of dollar volume using low-frequency data. AMMs—with their novel liquidity provision—introduce new challenges in measuring liquidity. We show that standard liquidity measures including the effective spread, realized spread, and the Amihud (2002) illiquidity measure extend naturally to AMM markets and keep their economic interpretation. Building on this insight, we develop liquidity metrics tailored to AMMs that allow for direct comparison with other markets, enabling a unified analysis of liquidity across different market structures.

Moreover, this chapter contributes to the growing literature comparing AMMs and LOB market designs. Prior studies examine differences in market quality and efficiency across these trading mechanisms (Aoyagi & Ito, 2021; Barbon & Ranaldo, 2022; Lehar & Parlour, 2025). Barbon and Ranaldo (2022) show that AMMs prices are less efficient than those in centralized exchanges, identifying exchange fees and transaction (gas) costs as the primary frictions impairing price efficiency. Related work by Capponi, Jia, and Yu (2023), as well as Chapter 3 of this dissertation, study the role of AMMs in price discovery and highlights how liquidity conditions and explicit trading costs shape their contribution to the price discovery process.

Parallel literature examines liquidity provision in AMMs. Aquilina et al. (2024), Heimbach, Wang, and Wattenhofer (2021), and Neuder et al. (2021) analyze the incentives and profitability of LPs, showing that although AMMs allow passive liquidity provision, the most profitable LPs actively manage their positions. We build on this research by documenting that AMMs liquidity is more stable than LOB liquidity. This stability helps explain cross-sectional variation in LPs profits, as more stable liquidity exposure increases sensitivity to adverse selection.

This chapter also relates to the literature on liquidity resilience in LOB markets. Theoretical models by Foucault, Kadan, and Kandel (2005) and Roşu (2009) characterize liquidity provision in limit order books as an endogenous outcome of order arrivals and order types, with limit orders supplying liquidity and market orders consuming it. In this setting, liquidity depends critically on traders' willingness to post and maintain limit orders. In contrast, liquidity in AMMs always rests in the pool by design, implying that trade arrival primarily affects the usage of liquidity rather than its immediate availability.

Empirical studies, including Degryse et al. (2005), Large (2007), and Lo and Hall (2015), show that liquidity in LOB markets can respond sluggishly to shocks, with disruptions that persist and impair price efficiency. We build on this literature by studying a market design based on AMMs and showing that its mechanical liquidity provision can enhance liquidity stability relative to LOB markets. This perspective highlights a channel through which alternative market designs improve the stability of liquidity and market resilience.

2.2 Automated market makers

Most modern financial markets rely on a LOB structure, where traders submit orders specifying the price and quantity at which they are willing to buy or sell. Liquidity in these markets is represented by the set of open orders available for execution, which provide other participants the opportunity to trade at pre-specified prices and quantities. When a trade takes place, the matched orders are executed and subsequently removed from the LOB, causing the associated liquidity to vanish. Market makers, who specialize in providing liquidity, must continuously monitor market conditions and update their quotes to ensure they remain competitive. This process helps prevent adverse selection from stale quotes but requires ongoing monitoring and frequent adjustments of their LOB positions to maintain effective participation.

Although modern electronic LOB markets are generally resilient, they are not without limitations. The price-time priority rule, which grants execution priority based on order price followed by submission time, can trigger latency races among HFTs competing for speed advantages (Budish, Cramton, & Shim, 2015). Moreover, liquidity provision is inherently costly due to risks such as adverse selection and inventory exposure. As a result, markets remain vulnerable to liquidity shocks when market makers temporarily withdraw from the market, requiring liquidity to be replenished (Lo & Hall, 2015).

AMMs represent a fundamentally different market design from traditional order book systems. Rather than facilitating trades through order matching, AMMs enable trading through a mathematical formula (commonly known as the invariant or pricing function). This function automatically determines the price of assets based on their relative quantities in a shared liquidity pool (Zhang, Chen, & Park, 2018). In AMMs, traders swap one asset for another directly against this pool, which is funded by LPs who deposit pairs of tokens into it. Unlike LOBs, where trades involve cash transfers between counterparties, AMM trades occur through token exchanges: one asset is deposited into the pool, and the other is withdrawn, maintaining the pool's balance according to the pricing formula. This automated structure eliminates the need for a matching engine or active counterparties at the moment of trade.

LPs contribute pairs of assets to the pool in exchange for a share of the trading fees generated by trades. A key feature of AMMs is that liquidity is not depleted when a trade occurs. Instead, each trade modifies the pool's asset balances, and in doing so, adjusts the prices according to the pricing formula. This mechanism ensures a self-sustaining pool without the need for LPs to constantly replenish their liquidity, as is required in LOB markets. Traders effectively "replenish" the pool by adding one asset while removing another. This structure allows LPs to provide liquidity passively, without having to monitor or update their positions. The invariant ensures that every successive trade can be executed, though at a price that reflects the new balance of assets in the pool (Adams et al., 2021; Zhang, Chen, & Park, 2018).

LPs share both the risks and rewards of providing liquidity in proportion to their contribution to the pool. They earn a share of the trading fees generated by swaps, but also face the risk of adverse selection that arises when the relative prices of the pooled assets change. This structure contrasts with the LOB model, where market makers post individual orders and bear the full risk of having their quotes hit by better-informed traders.

Although AMMs gained popularity through their use in trading crypto assets within DeFi, their conceptual origins trace back to prediction markets. Early AMMs were developed to support market-making in low-liquidity environments where maintaining LOBs was impractical (Hanson, 2003). These early designs were typically not liquidity-sensitive, meaning they could estimate probability distributions, such as the likelihood of a team winning a match, but were limited in their ability to reflect changing conditions. Later work introduced liquidity-sensitive pricing functions, allowing AMMs to better account for trade size and slippage, thereby extending their applicability beyond niche markets (Othman & Sandholm, 2011; Othman et al., 2013).

AMMs gained popularity in DeFi not only for their simplicity, but also due to their low operational and data requirements compared to LOBs. Implementing a fully on-chain LOB demands significant storage and computational resources to track, match, and cancel orders in real time across a distributed system, something that is costly on a blockchain.

In contrast, AMMs offer a more practical and scalable alternative in blockchain-based markets.[1]

In DeFi, AMMs are most commonly implemented as Constant Product Market Makers (CPMMs), the design used by platforms like Uniswap v2. CPMMs determine prices using a simple mathematical formula (e.g., $x \cdot y = k$), which ensures that the product of the two asset reserves in the pool remains constant after each trade. Building on this framework, Concentrated Liquidity Constant Product Market Makers (CL-CPMMs), as introduced in Uniswap v3, allow LPs to provide liquidity within specific price ranges. This design improves capital efficiency by letting LPs leverage their reserves to earn more fees with less capital, provided prices stay within the specified range.

2.2.1 Constant product market maker

Developed by Uniswap, the CPMM is by far the most widely adopted AMM design (Zhang, Chen, & Park, 2018). The CPMM pools two assets according to the following invariant:

$$L = \sqrt{r_1 r_2}, \tag{2.1}$$

where r_1 and r_2 represent the reserves of the two assets in the pool, and L is the invariant, which can be interpreted as the total units of liquidity.[2] The price P quoted by the pool is determined by the ratio of reserves:

$$P = \frac{r_2}{r_1}, \tag{2.2}$$

where P reflects the current exchange rate between the two assets in the pool.

[1] Appendix 2.1 provides a figure showing the share of volume traded by different AMM designs in DeFi. Given their dominant market share, our analysis focuses specifically on the CPMM and CLCPMM designs.

[2] The CPMM invariant is also commonly expressed as $x \cdot y = k$. Both formulations are equivalent for determining swap prices and tracking market states. The only notable difference is that the expression using L has the advantage of being positively homogeneous, meaning that a 1% increase in reserves results in a 1% increase in L (Othman, 2021).

To execute a swap, a trader adds one asset to the pool and receives an amount of the other asset in return, such that the pricing function remains constant:

$$L = \sqrt{r_1 r_2} = \sqrt{(r_1 + \Delta r_1)(r_2 + \Delta r_2)}. \tag{2.3}$$

Solving for Δr_1 and Δr_2, we can determine the amount exchanged. As shown in Equation 2.4, the swap amount depends on the change in the square root of the price $\Delta\sqrt{P}$ ($\Delta\frac{1}{\sqrt{P}}$):

$$\Delta r_1 = L\Big(\Delta\frac{1}{\sqrt{P}}\Big), \quad \Delta r_2 = L(\Delta\sqrt{P}). \tag{2.4}$$

Fees are charged as a percentage of the asset being added to the pool and are calculated before the swap is executed to ensure accurate compensation for LPs. The fee-adjusted amount entering the pool is given by:

$$\Delta r_f = \Delta r(1 + f), \tag{2.5}$$

where f is the swap fee, Δr is the amount of the asset entering the pool before fees are collected, and Δr_f is the fee adjusted amount used in the swap.

In the CPMM, the invariant L is held constant during each swap. Conversely, when a LP adds or removes assets from the pool, the price P is held constant. This ensures that liquidity can be added or withdrawn without affecting the pool's current price, and that the LPs contribution is proportional to the existing liquidity, preserving the balance defined by the invariant.

Figure 2.1 illustrates the design of a CPMM. The left panel shows the invariant curve as defined in Equation 2.1. Each point on the curve represents a unique combination of reserves that satisfies the constant product condition. When a swap takes place, the trade moves the price along this curve. The direction of movement depends on the asset being added: adding one asset increases its supply in the pool, which reduces its relative price and raises the price of the other, capturing the supply-demand dynamics embedded in the AMM.

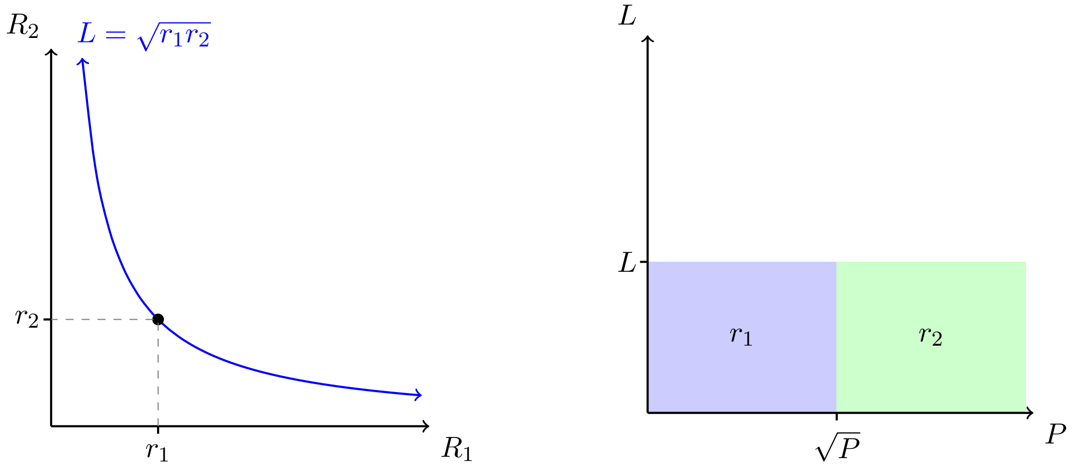

FIGURE 2.1: Visualization of a constant product market maker

This figure provides a visualization of a CPMM. The left panel plots the invariant curve defined by $L = \sqrt{r_1 r_2}$, where r_1 and r_2 represent the reserves of two assets in the pool. Each point along the curve reflects a combination of reserves that satisfies the invariant condition. The right panel shows the relationship between the invariant L and the price P. The region above the current square root price ($\sqrt{P}$) can be thought of as the reserves held in the pool for asset r_2, while the region below the current square root price can be thought of as the reserves held in the pool for asset r_1.

The right panel of Figure 2.1 plots the relation between the invariant L and price P. Each shaded region illustrates the current reserves of the two assets in the pool (r_1 and r_2). Assets held above the current square root price ($\sqrt{P}$) correspond to the reserve of one asset (typically the quote asset, r_2), while those below correspond to the other (typically the base asset, r_1). This structure is similar to an order book: to shift the price, a trader must add one asset and remove the other, similar to consuming liquidity on one side of a traditional LOB. As such, the CPMM implicitly provides quotes on both sides of the current market price.

2.2.2 Concentrated liquidity

A popular extension of the CPMM is the CL-CPMM, introduced in Uniswap v3 (Adams et al., 2021). The CL-CPMM enhances the original CPMM design by allowing LPs to allocate their capital within a specific price range, rather than across the full range of prices $[0, \infty]$. In the standard CPMM, liquidity is spread evenly across all possible prices, including those far from the current market price, many of which are unlikely to be reached. This leads to inefficient use of capital, as much of the liquidity remains idle.

The CL-CPMM addresses this inefficiency by allowing LPs to provide liquidity within a selected price range $[P_a, P_b]$. Within this range, the reserves of the two assets must satisfy the following invariant condition:

$$(r_1 + \frac{L}{\sqrt{P_b}})(r_2 + L\sqrt{P_a}) = L^2, \tag{2.6}$$

where r_1 and r_2 denote the reserves of the two assets in the pool within the price range $[P_a, P_b]$, and L represents the amount of liquidity provided in this range.

Figure 2.2 illustrates the CL-CPMM design. The left panel shows the invariant curve restricted to the price range $[P_a, P_b]$, capturing only the relevant portion of the reserves. For comparison, a second curve, representing the full-range CPMM invariant, depicts the "virtual" position that would exist if liquidity were spread across all prices. The right panel plots the relation between liquidity L and price P. Liquidity is distributed across different price intervals, with each interval representing a separate liquidity position. Similar to the CPMM, reserves of asset r_2 are held above the current square root price $\sqrt{P}$, while reserves of asset r_1 are held below $\sqrt{P}$.

Trade amounts are calculated similarly to the CPMM, using changes in the square root of the price $\sqrt{P}$, as shown in Equation 2.4. However, when trades move the price across multiple price intervals, each with its own level of liquidity (L), the calculations are performed iteratively across each interval. The key difference between the CPMM and CL-CPMM lies in the distribution of liquidity across price intervals. By allowing LPs to concentrate liquidity around the expected trading range, the CL-CPMM improves capital efficiency and increase market depth near the current price.

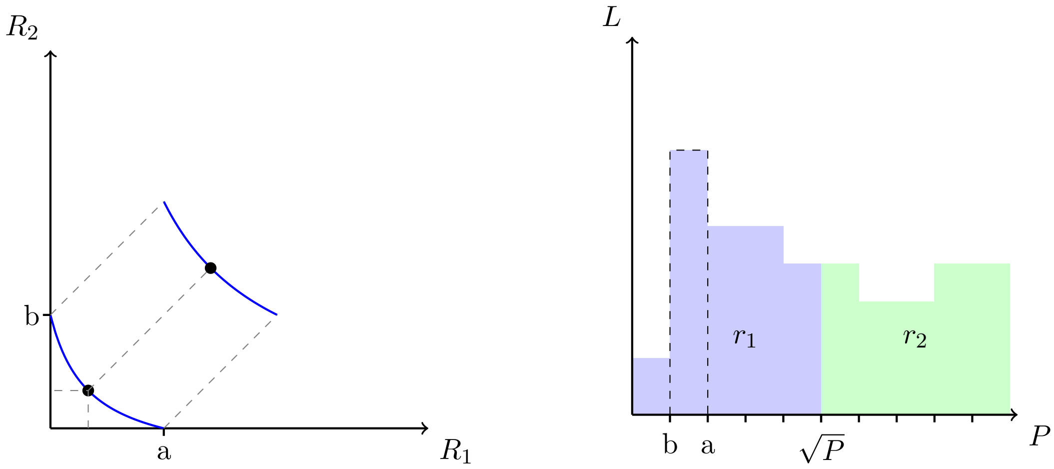

FIGURE 2.2: Visualization of a concentrated liquidity constant product market maker

This figure provides a visualization of a CL-CPMM. The left panel shows the invariant curve plotted within the price range of a to b. The combination of reserves satisfying the invariant condition for a price range is reflected in the curves. Based on real reserves, the invariant $L^2 = (r_1 + \frac{L}{\sqrt{p_b}})(r_2 + L\sqrt{p_a})$ describes the curve ending at prices a and b. The other curve can be thought of as the equivalent invariant L if the market uses a constant product market maker model. Illustrated in the right panel is the relationship between liquidity L and the price P. The region above the current square root price ($\sqrt{P}$) represents the reserves of asset r_2 available in the pool, while the region below the current square root price represents the reserves of asset r_1.

2.2.3 Automated market makers and limit order books

Although AMMs and LOBs operate using fundamentally different mechanisms to facilitating trades, they can ultimately be treated the same when examining market liquidity. Figure 2.3 illustrates how both systems relate price to quantity across various market designs. Panel (A) shows a traditional LOB, where liquidity is provided at discrete price levels, or ticks, with each representing limit orders queued for execution. In contrast, AMMs use a continuous pricing model. Panel (B) present the CPMM design, where the price-quantity relationship follows a constant curve defined by the liquidity invariant and is only split at the midpoint by the trading fee. This curve depends on the total amount of liquidity L in the pool. Panel (C) shows the CL-CPMM design, which functions similarly to the CPMM but with liquidity L varying across different price intervals, which causes the slope of the curve to differ based on how much liquidity is allocated to each price range.

Another distinction between AMMs and LOBs is how liquidity responds to trading

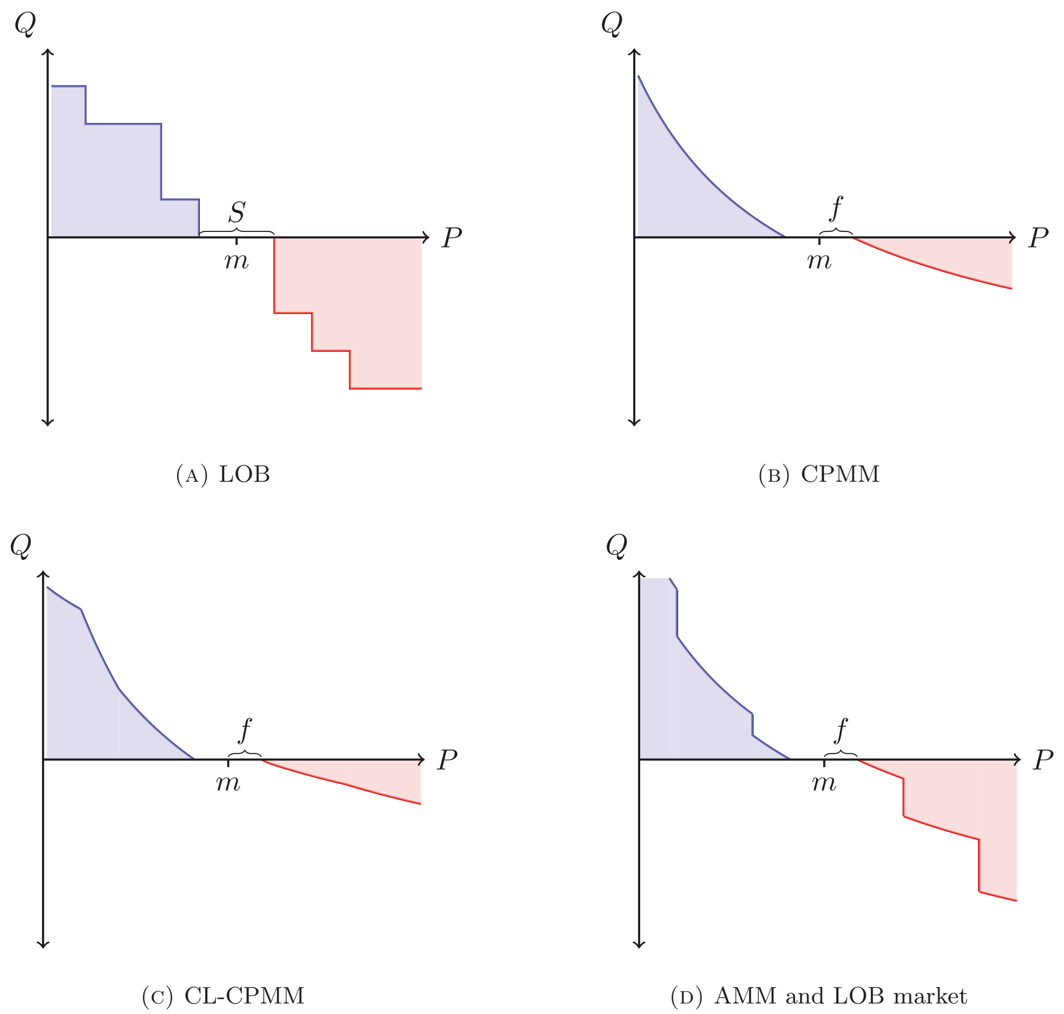

FIGURE 2.3: Comparison of price and trade quantity across different market designs

This figure visualizes the quantity available for trade at different prices under various market designs. Each panel plots price P on the x-axis and quantity Q on the y-axis. The mid-quote price is denoted by m, the bid-ask spread by S, and the AMM trading fee by f. Panel (A) illustrates a LOB, where discrete quantities are available at specific price levels, separated by the spread S. Panel (B) shows a CPMM where the trade quantity is calculated through the invariant $L = \sqrt{xy}$. Panel (C) illustrates a CL-CPMM, where liquidity is concentrated in different price ranges leading to changes in the slope of the tradable amount Q. Panel (D) shows a hybrid design combining an AMM with an LOB. In this panel, liquidity is provided both continuously via the AMM and discretely through the LOB.

activity. In LOBs, liquidity is removed when orders are executed, making it potentially fragile during times of stress. In contrast, AMM liquidity is more persistent as each trade alters the pool's balance but does not deplete liquidity in the same manner. This distinction raises an important question: **Is AMM liquidity more stable?** Liquidity stability—liquidity that remains continuously available —is especially valuable during periods of market stress or heightened volatility, when access to continuous trading is essential. If

AMM liquidity is more resilient, it could offer valuable insights for improving the design of modern markets to foster more stable and resilient liquidity.

Importantly, AMMs and LOBs are not mutually exclusive. Elements of both systems can be combined to improve market resilience. Liquidity can be additive across these systems, with AMM-style continuous liquidity complementing discrete LOB-based order flow. This hybrid approach may enhance market stability, reducing liquidity shocks and improving resilience during turbulent market conditions, Panel (D) of Figure 2.3 highlights the combination of the LOB and AMM market. This combination market has already been implemented by the cryptocurrency exchange Bullish, which integrates on-chain AMM liquidity into its traditional order book to improve execution quality and overall market depth.[3] We do not directly study Bullish, as the joint implementation of AMM and LOB within a single market makes it difficult to isolate the effects attributable specifically to the AMM. Regardless, Bullish markets its exchange as offering deterministic liquidity, consistent with the implementation of an AMM.

We focus our analysis on LOB (order-driven) markets because quote-driven markets are not a natural benchmark for AMMs. In a quote-driven market, a trader receives a single executable price and quantity from a dealer conditional on the requested trade size. Market makers keep full discretion over whether to quote, at what price, and for what size, based on inventory constraints and expected adverse selection. When a trade cannot be quoted or is quoted at prohibitive terms, liquidity effectively withdraws, generating a liquidity shock not indifferent to a LOB market.

In contrast, AMMs post a deterministic pricing function that maps all trade sizes into prices and adjusts mechanically through trading rather than discretionary quote revision. LOB markets remain different from quote-driven markets, where dealers observe order requests before committing to prices. From an empirical perspective, quote-driven markets are also opaque—particularly in cryptocurrency markets—making liquidity difficult to

[3] see https://bullish.com/amm/.

measure. LOBs provide the optimal setting to compare liquidity across different market designs.

2.3 Measuring liquidity in automated market makers

We argue that AMMs and LOBs are comparable in terms of the relationship between price and quantity. This allows us to apply similar liquidity measures across both market designs. In this section, we outline how traditional liquidity metrics used in LOB markets can be adapted to AMMs. Specifically, we consider several well-established measures such as spreads, market depth, and other low-frequency measures of liquidity.

2.3.1 Spreads

Spreads are among the most commonly used measures of liquidity in financial markets, both in academic research and industry practice. Spreads reflect the implicit trading costs faced by market participants. There are three commonly used measures of spread: the quoted spread, the effective spread and the realized spread, each capturing a different aspect of market liquidity and trading frictions.

2.3.1.1 Quoted spread

In a LOB market, the quoted spread is defined as the difference between the best ask and best bid quotes, normalized by the mid-quote. The quoted spread captures the cost of immediacy for an investor submitting a market order and is widely used as a proxy for market tightness: the narrower the spread, the lower the cost to trade (Huang & Stoll, 1996). Following Foucault, Pagano, and Roell (2013), the quoted spread is given by:

$$s = \frac{a - b}{m}, \tag{2.7}$$

where a is the best ask price, b is the best bid price, and the mid-quote is defined as $m = \frac{a+b}{2}$.

In an AMM, the quoted spread is equivalent to the swap fee, which is typically charged as a percentage of the amount being traded. This fee is collected from the asset being added to the liquidity pool and effectively reduces the tradable amount received in return. As a result, the swap fee functions like a spread as it creates a wedge between the amount given and the amount received, altering the pool's reserves and impacting the invariant L and consequently the tradable amount Δr_1.

To approximate a comparable measure of the quoted spread in an AMM, we can simply double the fee f:

$$s = 2f. \tag{2.8}$$

This reflects the total round-trip cost of executing a buy and then a sell (or vice versa) of a infinitesimally small trade in the AMM.

In most AMMs, the fee is constant and defined at the pool level, meaning the quoted spread remains static over time. However, this is not universal. Some AMMs, such as Quickswap v3, implement dynamic fees that adjust in response to market conditions. While the fee at any given point in times remains fixed for that market state, it may vary over time, introducing potential variability in the AMMs quoted spread.

While some studies on AMMs, such as Barbon and Ranaldo (2022), treat the swap fee as an explicit trading cost—analogous to maker-taker fees on centralized exchanges—and gas fees as a variable implicit cost, we argue that the swap fee is more appropriately considered an implicit cost, similar to the quoted spread in LOBs, while gas fees should be treated as an explicit blockchain cost. Gas fees vary across participants and trades; for example, block builders often pay only the base fee without priority fees, particularly when including arbitrage transactions across centralized and decentralized exchanges (Heimbach, Pahari, & Schertenleib, 2024). Treating gas fees as explicit trading costs abstracts from these variations and could lead to conclusions that are specific to certain blockchains rather than reflecting AMMs, which do not have to exist on a blockchain.

Swap fees, in contrast, are embedded in the execution price within the pool. Traders on Decentralized Exchanges (DEXs) interact through front-end interfaces that optimize routing and trade execution across pools and fee tiers, meaning the fee is effectively internalized rather than visible as a separate cost. Moreover, traders cannot choose which token the fee is paid in; it is always taken from the token entering the AMM. This reinforces the interpretation of the swap fee as analogous to a quoted spread. This view is consistent with the Uniswap interface, which reports a "price impact" metric inclusive of the swap fee, reflecting the effective cost of the trade. Treating the swap fee as an implicit cost therefore aligns with the mechanics of AMMs, while treating gas as an explicit cost separates blockchain frictions from trading costs.

2.3.1.2 Effective spread

Unlike the quoted spread, which reflects the static cost of immediacy, the effective spread measures the actual cost a trader incurs when executing a trade. The effective spread captures how far the execution price deviates from the prevailing market price at the time of the trade, reflecting both the quoted spread and any additional slippage due to trade size (Huang & Stoll, 1996). In a LOB market, the effective spread is defined as:

$$s_e = d \cdot \frac{p - m}{m}, \tag{2.9}$$

where d is a trade direction indicator equal to 1 for buyer-initiated trades and -1 for seller-initiated trades, p is the execution price, and m is the mid-quote at the time of the trade (Foucault, Pagano, & Roell, 2013).

We can apply the same principle to AMMs. In this case, the swap price (i.e., the effective execution price) is determined by the change in the pool's reserves, $p = \frac{\Delta r_2}{\Delta r_1}$. The mid-quote depends on the AMM design. In a CPMM the mid-quote is equal to the ratio of reserves $m = \frac{r_2}{r_1}$. In a CL-CPMM, the protocol maintains an up-to-date market price, often tracked as $\sqrt{P}$, which is equivalent to the mid-quote $m = P$.

In AMMs, the effective spread reflects two components: the swap fee (analogous to the quoted spread) and the price impact, which results from the trade moving along the curve. While the fee component is fixed and proportional to trade size, the price impact varies with the depth of the pool and the size of the trade. For smaller trades, the price impact is minimal and the effective spread converges toward the swap fee.

2.3.1.3 Realized spread

The realized spread introduces a time dimension to spread measurement. While the effective spread captures the immediate cost of executing a trade, the realized spread measures the long-term profitability of liquidity provision by comparing the execution price to the market price at a later point in time (Huang & Stoll, 1996). It reflects the extent to which a LP gains or loses value after the trade has occurred. In a LOB, the realized spread is defined as:

$$s_r = d \cdot \frac{p - m_{t+\Delta}}{m}, \tag{2.10}$$

where d is a trade direction indicator equal to 1 for buyer-initiated trades and -1 for seller-initiated trades, p is the execution price, m is the mid-quote at the time of the trade, and $m_{t+\Delta}$ is the mid-quote observed at a future time $t+\Delta$ (Foucault, Pagano, & Roell, 2013). The time delta Δ is typically fixed (e.g., 5 minutes after the trade) but it may also be defined in event time (e.g.,number of trades or, in blockchain-based environments, number of blocks after the trade).

In AMMs, the realized spread can be calculated similarly to the effective spread. The trade price is derived from the change in reserves during the swap, and the future mid-price can be inferred from the pool state at a later time. This measure is particularly useful for evaluating whether LPs earn compensation above and beyond the fee, or if their returns are eroded by price changes following trades, highlighting the information content of AMM trades.

2.3.2 Depth

Market depth is a common measure of liquidity that captures a market's ability to absorb large trades without significant price movements. It reflects the quantity of assets available for trading at or near the current price and plays a critical role in determining price impact (Hasbrouck, 2007).

In a LOB, depth is typically measured as the cumulative volume of resting buy and sell orders within a certain percentage band around the mid-price (e.g., $\pm 1\%$). This metric indicates how much liquidity is available close to the current price, helping assess how resilient the market is to trading pressure. Low depth implies higher slippage and trading costs, while, deeper markets allow for larger trades to be executed with minimal price disruption, reducing trading costs and improving overall market efficiency.

Formally, depth in a LOB for a given percentage deviation δ from the midpoint is defined as:

$$Y(\delta) = \sum_{p=m}^{m(1\pm\delta)} \sum_{i=1}^{j} y_{i,p}, \tag{2.11}$$

where $y_{i,p}$ is the quantity of order i at price level p, and the summation includes all orders from the midpoint m up to a percentage price change of δ (Foucault, Pagano, & Roell, 2013).

In a CPMM, depth can be expressed as the amount of an asset that can be traded before the price moves by a given percentage deviation δ. This is calculated based on the invariant and swap pricing formula in Equation 2.4. Specifically, the depth at a percentage deviation δ from the current price P is:

$$Y(\delta) \equiv |\Delta r_1| = \left| L\left(\frac{1}{\sqrt{P \pm \delta}} - \frac{1}{\sqrt{P}}\right)\right|, \tag{2.12}$$

where L is the liquidity in the pool, and Δr_1 is the input asset required to move the price by δ.

In CL-CPMMs, the same logic applies, but with a key difference: liquidity L is not constant across price ranges. Instead, it varies by different price ranges depending on where LPs have concentrated their capital. To compute depth in these markets, we must iteratively calculate the tradable quantity across each price range up to the desired price bound. For each step, the corresponding L is used until the full percentage deviation δ is reached, similar to how a swap is calculated in these markets.

2.3.3 Other measures of liquidity

Low frequency measures offer an approach to estimate market liquidity using low frequency data on prices and volumes. Due to their reliance on standard variables, such as returns and traded volume, these measures translate easily to AMMs. In this chapter, we consider the widely used Amihud (2002) illiquidity ($ILLIQ$) measure:

$$ILLIQ = \frac{1}{T}\sum_{t=1}^{T}\frac{|return_t|}{\$volume_t}. \tag{2.13}$$

This measure captures the price sensitivity to trading volume, where higher values indicate lower liquidity. That is, large price movement associated with low trading volume suggest a more illiquid market.

With the new market structure introduced by AMMs in DeFi, additional liquidity metrics have emerged that are specific to these protocols and do not have direct applications in other market designs such as the LOB. One widely used metric is the Total Value Locked (TVL) in the liquidity pool.

TVL measures the total dollar value of the assets held in a liquidity pool. In AMMs TVL is directly related to market depth. TVL is commonly used to compare the scale and liquidity of different DeFi protocols, notably by platforms like Defillama who use

TVL to highlight significant economic contribution to platforms or projects.[4] In AMMs, TVL reflects the size of the reserves available for trading, which determines how sensitive prices are to trade size. A higher TVL indicates the AMM has more reserves, therefore greater available liquidity, lower price impact, and better execution quality for traders. Equation 2.14 defines the TVL as the sum of the dollar value of all reserves in the pool:

$$TVL = \sum \$r_t, \tag{2.14}$$

where $\$r_t$ denotes the dollar value of reserve r at time t. While TVL is not directly comparable to traditional spread or depth measures, it provides an intuitive benchmark of relative liquidity across AMM pools.

2.4 Data

To evaluate whether AMM liquidity is more stable than in traditional LOB markets, we calculate liquidity measures across Binance, Uniswap, and Quickswap over the period from June 2020 to March 2023.[5]Our dataset includes Uniswap v2 on Ethereum, as well as Uniswap v3 on Ethereum, Polygon, Arbitrum, and Optimism blockchains. We also include data from Quickswap v2, a direct fork of Uniswap v2 operating on the Polygon blockchain, and Quickswap v3 on Polygon, which adopts a concentrated liquidity model similar to Uniswap v3 but incorporates a dynamic fee mechanism that adjusts based on volatility, trading volume, and liquidity to optimize returns for LPs (Volosnikov et al., 2021).

We collect Binance trade data from the Binance Data Vision platform, while we gather Binance tick data from Tardis.dev. AMM data is collected directly from blockchain nodes associated with the relevant networks. The dataset comprises data on 27 of the most actively traded AMM pairs across the Ethereum, Polygon, Arbitrum, and Optimism blockchains. Appendix 2.3 provides a full list of the Binance trading pairs and the AMM

[4]see https://defillama.com/.

[5]Uniswap v2 was released in May 2020.

pools with their corresponding smart contract addresses across the different blockchains. Where applicable, we use wrapped versions of native assets (e.g., Wrapped ETH for ETH) to ensure consistency across AMM and LOB trading environments.

To focus on meaningful levels of liquidity, we exclude trading pairs with TVL below $100 thousand.[6] All liquidity measures are calculated at an hourly frequency and subsequently aggregated to obtain daily measures of liquidity variance.

Table 2.1 presents summary statistics for both AMM and LOB trading pairs and highlights several key differences between the two market structures. LOBs exhibit higher trading volumes and a substantially larger number of transactions, consistent with the high frequency, small trade size activity. In contrast, AMMs display larger average trade sizes and less frequent trades, suggesting that there are fewer but larger transactions.

Across both market types, roughly 10% of observations correspond to days with less than $500 thousand in daily trading volume, and the median daily trading volume in each market is approximately $1 million. Despite similar median volumes, AMMs execute far fewer trades than LOBs, with some AMM pairs recording fewer than 100 trades per day.

We find that, on average, quoted spreads are narrower in LOB markets, reflecting competition among market makers and their ability to adjust both prices and quantities in response to order flow. In contrast, AMMs offer only a limited set of discrete fee tiers, and this restricts the range of possible quoted spreads. As a result, LOB market makers can quote tighter spreads with more control over the exposure in the market.

At the same time, AMMs exhibit significantly greater market depth, implying stronger price resilience for large trades, as liquidity is continuously available in the pool. While LOB liquidity can replenish rapidly under normal conditions, it depends on ongoing, active quoting, whereas AMM liquidity appears always resting.

[6] The $100 thousand threshold is chosen to exclude periods of very low trading activity, where liquidity in the AMM is not be meaningfully comparable to the LOB. This threshold could be adjusted depending on market conditions and the goal of the research.

TABLE 2.1: Summary statistics of automated market maker and limit order book trading pairs

This table presents the summary statistics for the LOB (Panel A) and AMM (Panel B) trading pairs based on daily observations for the sample period between June 2020 and March 2023. *Volume* is the value of trade volume in millions of dollars in the day. *Trades* are the number of thousands of trades for the day. *Mean Trade* is the average of the hourly average trade size in dollars for the day. *Volatility* is the squared daily mid-quote return. The *Half Quoted Spread* for the AMM is equal to the liquidity pool fee. The Half Quoted Spread for the LOB is the hourly time weighted average of half the quoted spread for the day. *Effective Spread* is the average of the hourly volume weighted effective spread for the day. *Realized Spread* is the average of the hourly volume weighted realized spread for the day. *Depth*$_{[1\%]}$ is the hourly average value of orders in thousands of dollars between -1% and 1% of the mid-quote price in the day. *Amihud ILLIQ* is the hourly average Amihud (2002) illiquidity measure.

Panel A. Limit Order Book Summary Statistics					
	Mean	Standard Deviation	p10	Median	p90
Volume ($)	278	1,003	0	1	534
Trades	241.3	928.1	0.4	4.1	560.1
Volatility (%)	0.26	1.28	0.00	0.05	0.53
Mean Trade ($)	549	555	88	330	1,511
Half Quoted Spread (%)	0.12	0.12	0.00	0.08	0.28
Effective Spread (%)	0.11	0.12	0.01	0.08	0.26
Realized Spread (%)	0.15	0.15	0.02	0.11	0.35
Depth$_{[1\%]}$ ($)	1.27	3.69	0.00	0.04	3.42
Amihud ILLIQ	2.08	4.75	0.00	0.17	6.15

Panel B. Automated Market Maker Summary Statistics					
	Mean	Standard Deviation	p10	Median	p90
Volume ($)	15	70	0	1	28
Trades	1.1	3.2	0.0	0.1	2.7
Volatility (%)	0.25	1.67	0.00	0.04	0.47
Mean Trade ($)	15,067	39,940	319	2,644	33,131
Half Quoted Spread (%)	0.34	0.26	0.05	0.30	1.00
Effective Spread (%)	0.74	11.91	0.15	0.44	1.21
Realized Spread (%)	0.55	13.10	0.11	0.44	1.21
Depth$_{[1\%]}$ ($)	8.86	36.50	0.00	0.06	6.77
Amihud ILLIQ	2.67	202.73	0.00	0.15	4.04

To assess the stability of liquidity in AMMs, we calculate the intraday variation in liquidity by taking the log difference between the minimum and maximum hourly values of each liquidity measure within a day. Table 2.2 presents the summary statistics for the daily variability of the quoted and effective spread, 1% depth, and the ILLIQ measure.

As expected, variation in the quoted spread is observed only in Quickswap v3, which uses a dynamic fee model. In contrast, AMMs with fixed-fee pools (such as Uniswap v2, v3, and Quickswap v2) exhibit no intraday variation in quoted spreads. This underscores one of the key structural differences between AMM platforms.

Across the remaining liquidity measures—effective spread, depth, and ILLIQ—the average intraday variation is consistently lower in AMMs compared to LOBs. These

TABLE 2.2: Summary statistics of automated market maker and limit order book liquidity variance measures

This table presents the summary statistics for the variance of liquidity measures from LOB (Panel A) and AMM (Panel B) trading pairs based on daily observations for the sample period between June 2020 and March 2023. Δ^{min}_{max} represents the log difference between the min and max values in the day. The *Half Quoted Spread* for the AMM is equal to the liquidity pool fee. The *Half Quoted Spread* for the LOB is the hourly time weighted average of half the quoted spread for the day. *Effective Spread* is the hourly volume weighted effective spread. *Depth*$_{[1\%]}$ is the hourly average value of orders in thousands of dollars between -1% and 1% of the mid-quote price in the day. *ILLIQ* is the hourly average Amihud (2002) illiquidity measure.

Panel A. Limit Order Book Summary Statistics					
	Mean	Standard Deviation	p10	Median	p90
Δ^{min}_{max}Half Quoted Spread	0.75	0.52	0.27	0.64	1.32
Δ^{min}_{max}Effective Spread	1.76	0.84	0.94	1.64	2.71
Δ^{min}_{max}Depth$_{[1\%]}$	3.77	1.76	1.63	3.61	6.12
Δ^{min}_{max}ILLIQ	4.67	1.60	3.08	4.45	6.46

Panel B. Automated Market Maker Summary Statistics					
	Mean	Standard Deviation	p10	Median	p90
Δ^{min}_{max}Half Quoted Spread	0.01	0.11	0.00	0.00	0.00
Δ^{min}_{max}Effective Spread	1.09	0.93	0.24	0.84	2.19
Δ^{min}_{max}Depth$_{[1\%]}$	0.43	0.72	0.02	0.10	1.20
Δ^{min}_{max}ILLIQ	3.93	1.70	1.93	3.85	6.01

results suggest that AMM liquidity is more stable throughout the trading day than LOB liquidity. This finding provides empirical support for the argument that AMMs provide more resilient liquidity.

2.5 Is automated market maker liquidity more stable?

To statistically test whether AMM liquidity is more stable than that of LOB markets, we estimate the following regression model:

$$LiquiditySpread_{i,t} = \alpha + AMMDummy_{i,t} + Controls_{i,t} + \epsilon_{i,t}, \tag{2.15}$$

where the dependent variable, $LiquiditySpread_{i,t}$ represents the log difference between the minimum and maximum hourly values of key liquidity measures (effective spread, 1% depth and the ILLIQ measure) within a given day. The key independent variable, $AMMDummy_{i,t}$, is a dummy variable equal to 1 for AMM trading pairs (both CPMM and CL-CPMM) and 0 for LOB pairs. The CPMM and CL-CPMM variables can be interpreted as how much more (or less) stable liquidity is in these market designs compared to LOB. This specification allows for direct comparison of intraday liquidity variability across market designs. We include controls for trading volume, volatility, and other liquidity measures such as spread and depth to account for market conditions that may influence liquidity variability. We exclude the quoted spread from the analysis, as most AMMs use a fixed fee per pool, resulting in no variation in this measure over time.

Table 2.3 presents the results from Equation 2.15. We find that market depth is significantly more stable in AMM markets relative to LOBs. This suggests that AMMs are better able to maintain liquidity during volatile periods, when liquidity is needed most, supporting the idea that AMM-style liquidity provision could enhance the resiliency of traditional markets.

The ILLIQ measure further highlights differences between AMM designs. Liquidity in CPMM pools is significantly more stable than in LOBs, whereas CL-CPMM pools exhibit no statistically significant difference. This suggests that CL-CPMM markets, by concentrating liquidity in narrow price bands, behave more like LOBs in terms of liquidity dynamics. The relative stability of ILLIQ in CPMMs likely reflects the even distribution of liquidity across prices, whereas CL-CPMMs mirror LOB-style depth.

For the effective spread, we observe greater variability in AMMs compared to LOBs, as outlined by Regression 1 in Table 2.3. This variability likely stems from structural differences in trading behavior, particularly the larger trade sizes in AMMs, which generate larger price impact and wider effective spreads. In AMMs, continuous pricing means that trades often “walk the book”, in LOB terms, causing the effective spread to reflect not just execution costs, but also traders’ willingness to tolerate slippage.

Next, we examine what determinants how stable liquidity is in AMM markets. While the prior analysis shows that AMM liquidity is generally more stable than LOB liquidity, it is also important to understand what drives this stability. Several factors may influence it: First, higher volatility may reduce liquidity stability as LPs adjust or withdraw positions

TABLE 2.3: Regressions to assess if liquidity is more stable in automated market makers

This table reports coefficient estimates from OLS regressions using daily observations to examine if liquidity is more stable in AMMs. The dependent variable $\Delta_{max}^{min}ES$ in Regression 1 is the log difference between the min and max values of the Effective Spread in the day. The dependent variable $\Delta_{max}^{min}Depth_{1\%}$ in Regression 2 is the log difference between the min and max values of the $\text{Depth}_{[1\%]}$ in the day. The dependent variable $\Delta_{max}^{min}ILLIQ$ in Regression 3 is the log difference between the min and max values of the Amihud (2002) illiquidity measure in the day. The independent variables of focus are $CPMM$ and $CLCPMM$, which are dummy variables that equal to 1 if the pair trades on each AMM design (CPMM, CL-CPMM), respectively, otherwise 0. $Volume$ is the value of trade volume in millions of dollars in the day. $Volatility$ is the squared daily mid-quote return. $Depth_{[1\%]}$ is the hourly average value of orders in thousands of dollars between -1% and 1% of the mid-quote price in the day. The $HalfQuotedSpread$ for the AMM is equal to the liquidity pool fee. The Half Quoted Spread for the LOB is the hourly time weighted average of half the quoted spread for the day. All explanatory variables are in natural log form. The sample comprises 27 trading pairs during the period June 2020 to March 2023. t-statistics are reported in parentheses. Significance at the 10%, 5%, and 1% levels is indicated by *, **, and ***, respectively.

	(1)	(2)	(3)
Dependent variable:	$\Delta_{max}^{min}ES$	$\Delta_{max}^{min}Depth_{1\%}$	$\Delta_{max}^{min}ILLIQ$
$CPMM$	0.028***	−4.175***	−1.008***
	(3.16)	(-376.53)	(-43.40)
$CLCPMM$	0.195***	−3.327***	0.014
	(21.53)	(-295.73)	(0.55)
$Volume$	0.168***	0.107***	0.379***
	(76.69)	(39.43)	(65.81)
$Volatility$	−0.003***	0.040***	−0.027***
	(-2.90)	(28.63)	(-9.35)
$Depth_{1\%}$	−0.238***	−0.119***	−0.319***
	(-106.83)	(-43.37)	(-51.37)
$HalfQuotedSpread$	−0.278***	0.248***	0.224***
	(-92.78)	(66.96)	(28.73)
Intercept	−0.137***	5.880***	3.540***
	(-4.65)	(161.43)	(41.19)
Pair Effects	Y	Y	Y
Observations	84,225	83,881	54,446
Adjusted R^2	36.4%	71.0%	16.8%

Significance: * $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$

to avoid losses (Aoyagi & Ito, 2021). Similarly, trading volume may reduce stability in CL-CPMMs, where large movements in order flow can push prices outside the range of existing LP positions. Third, the level of total liquidity itself may play a role as deeper pools are likely less sensitive to individual trades, resulting in less variation in liquidity. Another factor to consider is market design: CPMMs distribute liquidity uniformly, while CL-CPMM pools, which concentrate liquidity within specific price intervals, mimic discrete LOB behavior and may therefore exhibit greater variability. Finally, blockchain-specific characteristics, such as block times and gas fees, can affect how frequently LPs adjust their positions, potentially impacting how stable liquidity is in AMMs.

To formally test these relationships, we estimate the following regression model on AMM pairs only:

$$LiquiditySpread_{i,t} = \alpha + CPMM_i + Volume_{i,t} + Volatility_{i,t} + Liquidity_{i,t} + \epsilon_{i,t}, \quad (2.16)$$

where the dependent variable, $LiquiditySpread_{i,t}$ represents the log difference between the minimum and maximum hourly values of key liquidity measures (effective spread, 1% depth and the ILLIQ measure) within a given day. Here, the $CPMM_i$ is a dummy variable equal to 1 if the pair follows a CPMM model and 0 if it uses a CL-CPMM design. Controls include trading volume, volatility, and liquidity measures, such as the depth and spread.

Table 2.4 presents the results of the regression analysis based on Equation 2.16. We find that liquidity is more stable in CPMMs than in CL-CPMMs. We also find that liquidity stability in AMMs decreases with trading volume, suggesting that greater trading activity leads to more variability in liquidity. While volatility is negatively related with the variation in effective spread and ILLIQ, the effect size is modest. In contrast, depth variability increases with volatility. LPs in CL-CPMMs set their liquidity positions across price ranges. Price changes can change how liquidity is distributed around the current price, leading to higher variability in depth.

We also find that pools with greater depth exhibit more stable liquidity across all measures. This supports the view that deeper pools not only reduce trading costs but also provide more stable liquidity. Furthermore, we find that lower-fee pools show more stability in effective spreads and ILLIQ, though fees appear unrelated to depth stability.

Combined, these findings reinforce the view that liquidity stability in AMMs is influenced by both market activity and the underlying liquidity design. Pools with larger liquidity reserves and lower trading costs offer more stable liquidity. This implies that constraining liquidity provision can enhance the stability of liquidity, but does so at the cost of limiting LPs ability to manage adverse selection risks, shifting risk from liquidity variability to potential losses for LPs.

TABLE 2.4: Regressions to assess the determinants of liquidity stability in automated market makers

This table reports coefficient estimates from OLS regressions using daily observations to examine the determinants of liquidity stability in AMMs. The dependent variable $\Delta_{max}^{min}ES$ in Regression 1 is the log difference between the min and max values of the Effective Spread in the day. The dependent variable $\Delta_{max}^{min}Depth_{1\%}$ in Regression 2 is the log difference between the min and max values of the Depth$_{[1\%]}$ in the day. The dependent variable $\Delta_{max}^{min}ILLIQ$ in Regression 3 is the log difference between the min and max values of the Amihud (2002) illiquidity measure in the day. $CPMM$ is a dummy variables that equal to 1 if the pair trades on a CPMM, otherwise 0. $Volume$ is the value of trade volume in millions of dollars in the day. $Volatility$ is the squared daily mid-quote return. $Depth_{[1\%]}$ is the hourly average value of orders in thousands of dollars between -1% and 1% of the mid-quote price in the day. The $HalfQuotedSpread$ for the AMM is equal to the liquidity pool fee. All explanatory variables are in natural log form. The sample comprises 27 trading pairs that trade on AMMs during the period June 2020 to March 2023. t-statistics are reported in parentheses. Significance at the 10%, 5%, and 1% levels is indicated by *, **, and ***, respectively.

	(1)	(2)	(3)
Dependent variable:	$\Delta_{max}^{min}ES$	$\Delta_{max}^{min}Depth_{1\%}$	$\Delta_{max}^{min}ILLIQ$
$CPMM$	−0.145***	−0.778***	−0.545***
	(-20.94)	(-124.16)	(-21.62)
$Volume$	0.210***	0.112***	0.491***
	(88.40)	(52.20)	(76.10)
$Volatility$	−0.002*	0.022***	−0.025***
	(-1.86)	(21.02)	(-8.05)
$Depth_{1\%}$	−0.261***	−0.120***	−0.379***
	(-112.14)	(-57.12)	(-55.57)
$HalfQuotedSpread$	−0.457***	−0.003	−0.179***
	(-114.22)	(-0.70)	(-13.39)
Intercept	−1.230***	0.891***	0.039
	(-42.20)	(33.89)	(0.35)
Pair Effects	Y	Y	Y
Observations	65,056	65,496	40,823
Adjusted R^2	44.3%	24.1%	25.5%

Significance: * $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$

Transaction fees (gas) are a key friction for LPs. High fees make frequent updates to pool positions costly. On blockchains like Ethereum, where transactions are expensive, this cost creates a trade-off: LPs may choose to leave their positions unchanged rather than actively adjusting them in response to price movements. As a result, high transaction fees may contribute to more stable liquidity, since LPs are less frequently adjusting their positions. Besides fees, other blockchain-specific factors, such as block times, and transaction ordering, may affect how stable liquidity is in AMMs.

We further investigate whether blockchain-specific factors contribute to the variation in AMM liquidity stability. Differences in network infrastructure, such as transaction costs and block times, may influence how frequently LPs adjust their positions and how often traders execute orders. These frictions can affect the overall stability of liquidity provision in AMMs.

To examine these effects, we augment the regression model in Equation 2.16 by including blockchain-specific dummy variables and gas fees:

$$LiquiditySpread_{i,t} = \alpha + ChainDummy_{i,t} + Gas_{i,t} + Controls_{i,t} + \epsilon_{i,t} \tag{2.17}$$

where the key independent variable, $ChainDummy_{i,t}$, is a set of dummy variables identifying the blockchain on which each AMM pool operates. Ethereum serves as the baseline and take value 0, with the different dummy variables being equal to 1 for either Polygon, Arbitrum, or Optimism. $Gas_{i,t}$ captures the daily average gas cost on the respective blockchain of the AMM pool. We also control for trading volume, volatility, and liquidity.

Table 2.5 presents the regression results based on Equation 2.17. We find that liquidity is less stable in AMM pools deployed on Arbitrum and Optimism compared to those on Ethereum. This effect is particularly pronounced for the variation in depth, indicating that liquidity on these Layer 2 blockchains is more prone to fluctuations. In contrast, effective spreads on Polygon show lower variance, suggesting that trading costs are more stable on this blockchain. However, depth varies more on Polygon relative to Ethereum, which may reflect heightened sensitivity of liquidity provision to trading activity or LP repositioning on this network.

We also find that gas fees are positively associated with variation in both the effective spread and depth. This suggests that higher transaction costs may constrain users' ability to update liquidity positions in response to price movements, resulting in less frequent adjustments and greater variability in liquidity measures. Conversely, variance in ILLIQ decreases as gas fees increase, implying that elevated transaction costs could suppress trading activity, which can lead to more stable liquidity in the AMM. Overall, these results indicate that blockchain-specific characteristics play an important role in determining liquidity stability in AMMs, underscoring the need to account for network-level costs when evaluating AMM performance. For AMM design, these findings show that blockchain

Table 2.5: Regressions to assess blockchain determinants of liquidity stability in automated market makers

This table reports coefficient estimates from OLS regressions using daily observations to examine the blockchain based determinants of liquidity stability in AMMs. The dependent variable $\Delta_{max}^{min} ES$ in Regression 1 is the log difference between the min and max values of the Effective Spread in the day. The dependent variable $\Delta_{max}^{min} Depth_{1\%}$ in Regression 2 is the log difference between the min and max values of the $\text{Depth}_{[1\%]}$ in the day. The dependent variable $\Delta_{max}^{min} ILLIQ$ in Regression 3 is the log difference between the min and max values of the Amihud (2002) illiquidity measure in the day. *Arbitrum*, *Optimism*, and *Polygon* are dummy variables that are equal to 1 if the pair trades on the respective blockchain, otherwise 0. *Gas* is the daily average gas price for the on-chain transactions. *Volume* is the value of trade volume in millions of dollars in the day. *Volatility* is the squared daily mid-quote return. All explanatory variables are in natural log form. The sample comprises 27 trading pairs that trade on AMMs during the period June 2020 to March 2023. t-statistics are reported in parentheses. Significance at the 10%, 5%, and 1% levels is indicated by *, **, and ***, respectively.

	(1)	(2)	(3)
Dependent variable:	$\Delta_{max}^{min} ES$	$\Delta_{max}^{min} Depth_{1\%}$	$\Delta_{max}^{min} ILLIQ$
Arbitrum	0.126***	0.293***	0.811***
	(9.16)	(26.54)	(23.46)
Optimism	0.062**	0.800***	0.359***
	(2.08)	(33.20)	(5.05)
Polygon	−0.194***	0.284***	0.012
	(-10.02)	(18.35)	(0.26)
Gas	0.010***	0.041***	−0.050***
	(4.23)	(21.72)	(-8.71)
Volume	0.045***	0.018***	0.243***
	(31.23)	(15.43)	(66.14)
Volatility	0.024***	0.028***	0.014***
	(17.74)	(25.65)	(4.54)
Intercept	0.976***	0.884***	−0.311***
	(25.87)	(29.17)	(-3.24)
Pair Effects	Y	Y	Y
Observations	65,060	65,496	40,827
Adjusted R^2	14.5%	7.8%	19.7%

Significance: * $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$

characteristics such as transaction fees, ordering, and speed can affect liquidity dynamics on AMMs

2.6 Learning from automated market makers

We find that AMM liquidity is more stable than LOB liquidity. This stability arises from the design of AMMs, where liquidity is not depleted after a trade. In contrast, LOB liquidity is removed when orders are executed, requiring market makers to continuously replenish the order book. This replenishment is not guaranteed, particularly during periods of volatility or market turmoil (Large, 2007). The persistence of AMM liquidity suggests that integrating AMM-style mechanisms into modern markets could enhance the resilience and stability of liquidity provision.

Several approaches exist for incorporating AMM principles into existing market structures, each with unique benefits and limitations.

One approach involves market makers adopting AMM-inspired strategies when placing limit orders within the existing LOB market. Market makers could simulate an AMM pricing curve with individual orders in a LOB. As new trades and quotes arrive in the market, the market maker can dynamically submit and adjust their limit orders. This approach requires no changes to existing exchange infrastructure, but relies on algorithmic order management to approximate the behavior of an AMM.

While replicating an AMM with limit orders may improve the stability of liquidity under normal conditions. Market makers must actively monitor and update their order with every trade. In times of heightened market activity, emulating an AMM with many individual orders in a LOB may lead to market infrastructure slowing further because of the increase in messages required to update the AMM orders.[7] If the market makers cannot effectively replicate the AMM because of slow market infrastructure (a likely result of an increase in orders) the market maker may be left exposed to adverse selection and latency games. Despite the potential ability of AMM-inspired market making to enhance resilience of liquidity to some extent, it falls short of fully replicating the passive and continuous liquidity provisioning inherent in AMM designs.

Another approach centers on a full transition to AMM-based trading systems. Trading tokenized assets on existing blockchain-based AMMs could offer the benefits of continuous liquidity while preserving many of the desirable features of limit orders. Recent innovations such as Uniswap v4 "hooks" allow for customized trading logic within AMMs, enabling the creation of dynamic strategies. Hooks can introduce limit order-like behavior for LPs

[7]Any trading system has a finite throughput, and increasing the number of messages can slow down processing. Trading platforms are engineered to handle millions of messages per second, with design efforts focused on mitigating congestion and maintaining consistent latency, including reducing tail latency (Alizadeh et al., 2012). As message rates approach system limits, delays can propagate, potentially slowing order execution and affecting market quality.

in AMMs where liquidity is removed once a trade is executed, emulating a limit order. This creates the potential for a hybrid execution model, where traders benefit from both persistent AMM liquidity and discretionary order placement within a DEX.

However, full adoption of AMM systems on decentralized infrastructure comes with several challenges. Blockchain-based trading is subject to issues such as MEV, high gas fees during periods of network congestion, and regulatory concerns on market integrity European Securities and Markets Authority (2023b). These frictions currently limit the feasibility of deploying pure AMM systems to help improve traditional markets. However, tackling these challenges would unlock the potential benefits of these markets.

A potentially more practical solution may lie in a hybrid market structure that integrates AMM orders directly within the traditional centralized LOB markets. In this model, a new order type in the LOB could replicate the AMM liquidity. These "*AMM orders*" can be thought about as limit orders that can be both a buy or a sell, depending on the current mid-quote price of the market. If the current price is above (below) the AMM order, the AMM order represents a buy (sell) order. When the AMM order is matched against another order, instead of the order being removed, it is replaced with the other asset in the asset trading pair, the buy order becomes a sell order and vice versa, replicating an AMM. This mechanism ensures that liquidity from the AMM orders is persistent post-trade. Introducing a mechanism that reflects a fee for the AMM would act like a spread, meaning all trades need to cross the spread towards the AMM order to be executed. Bullish, a centralized cryptocurrency exchange, offers a live example of a similar market structure that resembles this hybrid model.

However, merging AMM and LOB designs presents practical challenges. LOBs operate with discrete price ticks, while AMMs rely on continuous price curves. This discrepancy creates a mismatch in how prices and quantities are expressed. A simple solution to this problem is to use a different pricing function, the constant sum, within each tick. Over discrete price intervals within each interval the constant sum emulates a LOB and is closely related to a CPMM, such that when discrete price intervals are crossed, the markets are equivalent.[8] Some evidence suggests that larger discrete ticks improve liquidity, leading some to view continuous AMM pricing as undesirable (Foley, Meling, & Ødegaard, 2023).

Another important design consideration is how execution priority of AMM orders relative to traditional limit orders should be determined. Should AMM orders receive priority within a tick due to their continuous nature and passive liquidity provision? And how should execution be allocated across multiple AMM orders at the same price level? One solution could involve proportional execution, where AMM orders share fill allocations at the current price based on their liquidity contributions. Alternatively, exchanges could assign priority to AMM orders only when traditional depth is insufficient, reserving their role for backstopping liquidity rather than replacing the order book entirely.

An implementation of an AMM order would need to include a price, quantity, and a fee/spread. The price in the AMM order would not represent the trade price like a limit order, instead it would represent a fair price for the order. The price for the order that can be traded in the LOB would be based on the fair price of the order adjusted by the fee or spread. The fee/spread for the AMM order could be as a percentage of price or quoted in a number of price ticks.

Take an example of a single tick AMM order, if the mid-quote price in the order book is \$10 and there is an AMM order for a single quantity is placed at the fair price of \$9, with a spread of 1%, because the current mid-quote price is above the price of the AMM order, the price for the buy order in the order book is placed at \$8.91.[9] After some time of trading, if the price has moved below the fair price of \$9 to \$8, the AMM order would be traded against for its full quantity, rebalancing to a sell order at \$9.09.[10] Partial order fills where some part of the AMM order is trade but not in full, creates a market position where liquidity is added on both sides of the market with the fair price. In the example above, this would be the AMM order, as both a buy order for \$8.91 and a sell order for

[8] The difference between a constant sum and constant product market is based on the pricing curve within a price interval, the effective price for the constant product emulates the geometric mean, while the constant sum is traded at the price of the interval. The geometric mean of the price at the start and the end of the price range is equivlent to the mid price of this interval.

[9] calculated as $9(1 - 0.01)$.

[10] calculated as $9(1 + 0.01)$.

\$9.09. Combining multiple AMM orders across different prices gives a liquidity surface like what is common in concentrated liquidity AMMs.

The potential benefits of AMM orders go beyond improving the stability of liquidity. By modeling liquidity provision in as a tradeoff between the cost of monitoring markets against the cost of adverse selection (Khomyn & Putniņš, 2021). Introducing AMM orders may significantly reduce the cost of monitoring markets, meaning LPs can accept higher adverse selection costs. For informed traders, LPs accepting higher adverse selection costs leads to greater profit from their information, further leading to greater information production.

More broadly, hybrid models that combine AMM-style liquidity with discrete LOB mechanics open up opportunities to rethink market structure around resilience and efficiency, rather than just speed and fragmentation. As market participants increasingly prioritize robustness over marginal latency gains, these designs may offer a path forward for evolving both centralized and decentralized trading platforms.

2.7 Conclusion

This chapter demonstrates that AMMs can be analyzed using frameworks traditionally applied to LOB markets highlighting that, despite their structural differences, liquidity in both systems can be measured in comparable ways. Our findings show that AMM liquidity is significantly more stable than LOB liquidity. This stability stems from the fact that AMM liquidity is not removed after each trade, unlike in LOB markets where executed orders must be actively replenished by market makers. As a result, AMMs provide more persistent liquidity, particularly during periods of heightened volatility or market stress.

The resilience of AMM liquidity has important implications for the future of market design. By incorporating AMM-style mechanisms into traditional financial markets, exchanges may address one of the most persistent challenges in LOB markets: the disappearance of liquidity when it is needed most. AMMs offer a path toward a more self-sustaining and robust liquidity model, with lower operational demands on LPs and less reliance on continuous market-making activity.

We explore several potential pathways for integrating AMM principles, including algorithmic market-marking strategies, hybrid LOB-AMM order types, and programmable AMM protocols that support limit order functionality. Each option presents a different balance between infrastructure complexity, flexibility, and resilience.

These findings are especially relevant for exchanges, policymakers, and platform designers seeking to future-proof markets against volatility, high-frequency trading distortions, and liquidity dry-ups. As financial systems evolve toward tokenization, 24/7 trading, and increasingly automated execution environments, the principles underpinning AMM liquidity may prove essential for building next-generation markets.

While further research is needed to refine these hybrid models and address implementation challenges, our study highlights a promising path forward. Incorporating AMM-style liquidity mechanisms into traditional market designs could help build more robust liquidity infrastructures for the future of financial markets.

Appendix 2.1 Automated market maker design dominance

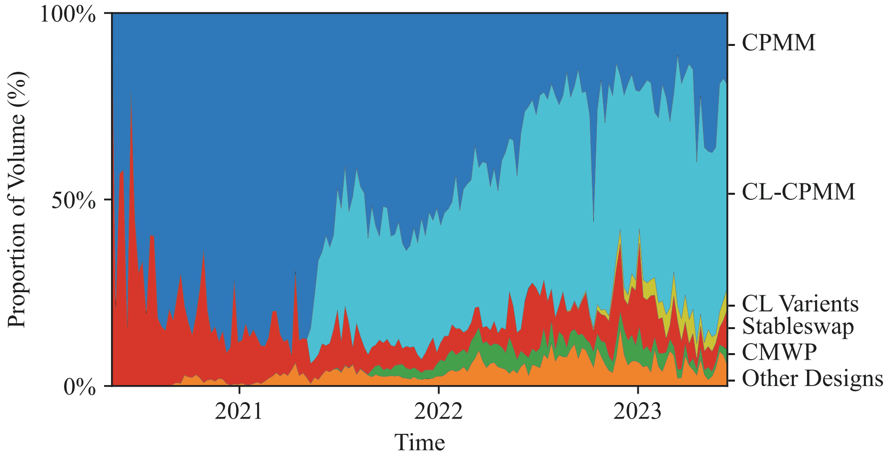

FIGURE A2.1: Dominance of automated market maker designs by volume

This figure plots the proportion of total daily volume for different AMM designs between May 2020 and June 2023. Each shaded region represents a different AMM design. The volume includes all the trade volume across different blockchains where the DEX has been deployed. The constant mean weight pools (CMWP) group includes the DEXs Balancer, Beethoven X and Osmosis DEX. The Stableswap group includes the DEXs Curve, Ellipsis Finance, Saddle Finance and Pancakeswap's Stableswap. The CPMM group includes the DEXs Uniswap Uniswap v2, Pancakeswap, Sushiswap, Quickswap, Spiritswap, Spookyswap, Apeswap, VVS Finance, Ubeswap, Trader Joe, Sunswap, Biswap, Fraxswap, Honeyswap, Klayswap and Solarbeam. The CL-CPMM group includes the DEXs Uniswap v3, Sushiswap v3, Pancakeswap v3, and Zyberswap v3. The concentrated liquidity (CL) variants group contains the DEXs Trader Joe Uniswap v2.1, Camelot v3, Kyberswap Elastic and Quickswap v3. The other design group contains the DEXs DoDo, Velodrome, Bancor, Clipper and Solidly.

Appendix 2.2 Correlation tables

TABLE A2.1: Correlation of automated market maker and limit order book variables

This table presents the correlation matrix of variables used in regressions containing AMM and LOB markets. $\Delta_{max}^{min}ES$ is the log difference between the min and max values of the Effective Spread in the day. $\Delta_{max}^{min}Depth_{1\%}$ is the log difference between the min and max values of the $\text{Depth}_{[1\%]}$ in the day. $\Delta_{max}^{min}ILLIQ$ is the log difference between the min and max values of the Amihud (2002) illiquidity measure in the day. $CPMM$ and $CLCPMM$ are dummy variables that equal to 1 if the pair trades on each AMM design (CPMM, CL-CPMM), respectively, otherwise 0. The $HalfQuotedSpread$ for the AMM is equal to the liquidity pool fee. The Half Quoted Spread for the LOB is the hourly time weighted average of half the quoted spread for the day. $Depth_{[1\%]}$ is the hourly average value of orders in thousands of dollars between -1% and 1% of the mid-quote price in the day. $Volatility$ is the squared daily mid-quote return. $GasPrice$ is the daily average gas price in dollars for the on-chain transactions. $Volume$ is the value of trade volume in millions of dollars in the day. All continuous variables are in natural log form.

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
$\Delta_{max}^{min}ES$ (1)	1.00									
$\Delta_{max}^{min}Depth_{1\%}$ (2)	0.25	1.00								
$\Delta_{max}^{min}ILLIQ$ (3)	0.24	0.23	1.00							
CPMM (4)	0.01	-0.41	-0.22	1.00						
CL-CPMM (5)	-0.26	-0.28	0.07	-0.62	1.00					
$HalfQuotedSpread$ (6)	-0.35	-0.32	-0.11	0.26	0.21	1.00				
$Depth_{1\%}$ (7)	-0.20	-0.04	0.09	-0.28	0.27	-0.39	1.00			
$Volatility$ (8)	0.11	0.11	0.08	0.07	-0.10	0.03	-0.08	1.00		
$GasPrice$ (9)	0.19	0.02	-0.04	-0.00	0.00	0.17	-0.04	0.11	1.00	
$Volume$ (10)	0.10	0.14	0.22	-0.13	-0.06	-0.60	0.81	0.14	0.03	1.00

TABLE A2.2: Correlation of automated market maker variables

This table presents the correlation matrix of variables used in regressions with only AMM markets. $\Delta_{max}^{min}ES$ is the log difference between the min and max values of the Effective Spread in the day. $\Delta_{max}^{min}Depth_{1\%}$ is the log difference between the min and max values of the $\text{Depth}_{[1\%]}$ in the day. $\Delta_{max}^{min}ILLIQ$ is the log difference between the min and max values of the Amihud (2002) illiquidity measure in the day. $CPMM$ and $CLCPMM$ are dummy variables that equal to 1 if the pair trades on each AMM design (CPMM, CL-CPMM), respectively, otherwise 0. The $HalfQuotedSpread$ for the AMM is equal to the liquidity pool fee. $Depth_{[1\%]}$ is the hourly average value of orders in thousands of dollars between -1% and 1% of the mid-quote price in the day. $Volatility$ is the squared daily mid-quote return. $GasPrice$ is the daily average gas price in dollars for the on-chain transactions. $Volume$ is the value of trade volume in millions of dollars in the day. *Arbitrum*, *Optimism*, and *Polygon* are dummy variables that are equal to 1 if the pair trades on the respective blockchain, otherwise 0. All continuous variables are in natural log form.

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)	(13)
$\Delta_{max}^{min}ES$ (1)	1.00												
$\Delta_{max}^{min}Depth_{1\%}$ (2)	0.08	1.00											
$\Delta_{max}^{min}ILLIQ$ (3)	0.24	0.18	1.00										
CPMM (4)	0.15	-0.39	-0.17	1.00									
CL-CPMM (5)	-0.15	0.39	0.17	-1.00	1.00								
$HalfQuotedSpread$ (6)	-0.35	-0.12	-0.22	0.14	-0.14	1.00							
$Depth_{1\%}$ (7)	-0.25	0.09	0.17	-0.33	0.33	-0.33	1.00						
$Volatility$ (8)	0.14	0.12	0.09	0.10	-0.10	0.06	-0.08	1.00					
$GasPrice$ (9)	0.19	0.02	-0.04	-0.00	0.00	0.17	-0.04	0.11	1.00				
$Volume$ (10)	0.04	0.12	0.35	-0.07	0.07	-0.42	0.83	0.15	0.03	1.00			
Arbitrum (11)	-0.04	0.12	0.08	-0.24	0.24	-0.17	0.09	-0.05	0.15	-0.02	1.00		
Optimism (12)	-0.03	0.08	0.05	-0.13	0.13	-0.11	0.03	-0.02	-0.38	-0.00	-0.05	1.00	
Polygon (13)	-0.20	-0.00	0.05	0.04	-0.04	-0.16	0.06	-0.06	-0.77	0.03	-0.13	-0.07	1.00

Appendix 2.3 Automated market maker pool smart contract addresses

Table A2.3: Smart contract addresses
This table reports each Binance trading pair and the corresponding AMM pool used for comparison of liquidity measures.

Pair	Market	Chain	Fee (%)	Pair Address
AAVEETH	Uniswap v2	Ethereum	0.30	0xDFC14d2Af169B0D36C4EFF567Ada9b2E0CAE044f
AAVEETH	Uniswap v3	Ethereum	0.30	0x5aB53EE1d50eeF2C1DD3d5402789cd27bB52c1bB
AAVEETH	Uniswap v3	Ethereum	1.00	0x1353fE67fFf8f376762b7034DC9066f0bE15a723
AAVEETH	Uniswap v3	Polygon	0.30	0x2aCeda63B5e958c45bd27d916ba701BC1DC08F7a
APEETH	Uniswap v2	Ethereum	0.30	0xb011EEaab8bF0c6DE75510128dA95498E4b7e67F
APEETH	Uniswap v3	Ethereum	1.00	0xF79fC43494ce8a4613Cb0b2a67A1b1207fD05D27
APEETH	Uniswap v3	Ethereum	0.30	0xAc4b3DacB91461209Ae9d41EC517c2B9Cb1B7DAF
APEETH	Uniswap v3	Ethereum	0.05	0x52DB2848b59AA96dE51c758b3538C7cE656fCdB8
BADGERBTC	Uniswap v2	Ethereum	0.30	0xcD7989894bc033581532D2cd88Da5db0A4b12859
BADGERBTC	Uniswap v3	Ethereum	0.30	0xe15e6583425700993bd08F51bF6e7B73cd5da91B
BATETH	Uniswap v2	Ethereum	0.30	0xB6909B960DbbE7392D405429eB2b3649752b4838
BATETH	Uniswap v3	Ethereum	0.30	0xAE614a7a56cB79c04Df2aeBA6f5dAB80A39CA78E
BATETH	Uniswap v3	Ethereum	1.00	0xb66776014d002479aCE38EC16e604e465851A33B
BTCUSDC	Uniswap v2	Ethereum	0.30	0x004375Dff511095CC5A197A54140a24eFEF3A416
BTCUSDC	Uniswap v3	Ethereum	0.05	0x9a772018FbD77fcD2d25657e5C547BAfF3Fd7D16
BTCUSDC	Uniswap v3	Ethereum	0.30	0x99ac8cA7087fA4A2A1FB6357269965A2014ABc35
BTCUSDC	Uniswap v3	Ethereum	1.00	0xCBFB0745b8489973Bf7b334d54fdBd573Df7eF3c
BTCUSDC	Quickswap	Polygon	0.30	0xf6a637525402643b0654a54bead2cb9a83c8b498
BTCUSDC	Quickswap v3	Polygon		0xa5cd8351cbf30b531c7b11b0d9d3ff38ea2e280f
BTCUSDT	Uniswap v2	Ethereum	0.30	0x0DE0Fa91b6DbaB8c8503aAA2D1DFa91a192cB149
BTCUSDT	Uniswap v3	Ethereum	0.30	0x9Db9e0e53058C89e5B94e29621a205198648425B
BTCUSDT	Uniswap v3	Ethereum	0.05	0x56534741CD8B152df6d48AdF7ac51f75169A83b2
BTCUSDT	Uniswap v3	Ethereum	1.00	0x5a59E4e647A3Acc42b01715F3A1D271c1f7e7aeb
CHZUSDT	Uniswap v2	Ethereum	0.30	0x31ac548E59565Fdd78604A47E1571Ef68C80E9f5
CHZUSDT	Uniswap v3	Ethereum	0.30	0xB0F4a77Bde7fEE134265307C5CC19abfF0ba409B
CRVETH	Uniswap v2	Ethereum	0.30	0x3dA1313aE46132A397D90d95B1424A9A7e3e0fCE
CRVETH	Uniswap v3	Ethereum	0.30	0x919Fa96e88d67499339577Fa202345436bcDaf79
CRVETH	Uniswap v3	Ethereum	1.00	0x4c83A7f819A5c37D64B4c5A2f8238Ea082fA1f4e
CVPETH	Uniswap v2	Ethereum	0.30	0x12D4444f96C644385D8ab355F6DDf801315b6254
CVPETH	Uniswap v3	Ethereum	1.00	0x860CaF97d5dA97145004cB45a4bF458686FD90E5
ENJETH	Uniswap v2	Ethereum	0.30	0xE56c60B5f9f7B5FC70DE0eb79c6EE7d00eFa2625
ENJETH	Uniswap v3	Ethereum	0.30	0xe16Be1798F860bC1EB0FEb64cD67Ca00AE9b6E58
ETHBTC	Uniswap v3	Arbitrum	0.05	0x2f5e87C9312fa29aed5c179E456625D79015299c
ETHBTC	Uniswap v3	Arbitrum	0.30	0x149e36E72726e0BceA5c59d40df2c43F60f5A22D
ETHBTC	Uniswap v3	Arbitrum	1.00	0x99dFc0126ED31E0169fc32dB6B89adF9FeE9a77e
ETHBTC	Uniswap v2	Ethereum	0.30	0xBb2b8038a1640196FbE3e38816F3e67Cba72D940
ETHBTC	Uniswap v3	Ethereum	0.01	0xe6ff8b9A37B0fab776134636D9981Aa778c4e718
ETHBTC	Uniswap v3	Ethereum	0.05	0x4585FE77225b41b697C938B018E2Ac67Ac5a20c0
ETHBTC	Uniswap v3	Ethereum	0.30	0xCBCdF9626bC03E24f779434178A73a0B4bad62eD
ETHBTC	Uniswap v3	Ethereum	1.00	0x6Ab3bba2F41e7eAA262fa5A1A9b3932fA161526F
ETHBTC	Uniswap v3	Optimism	0.05	0x85C31FFA3706d1cce9d525a00f1C7D4A2911754c
ETHBTC	Uniswap v3	Optimism	0.30	0x73B14a78a0D396C521f954532d43fd5fFe385216
ETHBTC	Quickswap v3	Polygon		0xac4494e30a85369e332bdb5230d6d694d4259dbc
ETHBTC	Uniswap v3	Polygon	0.05	0x50eaEDB835021E4A108B7290636d62E9765cc6d7
ETHBTC	Uniswap v3	Polygon	0.30	0xfe343675878100b344802A6763fd373fDeed07A4
ETHDAI	Uniswap v3	Arbitrum	0.05	0x31Fa55e03bAD93C7f8AFfdd2eC616EbFde246001
ETHDAI	Uniswap v3	Arbitrum	0.30	0xA961F0473dA4864C5eD28e00FcC53a3AAb056c1b
ETHDAI	Uniswap v2	Ethereum	0.30	0xA478c2975Ab1Ea89e8196811F51A7B7Ade33eB11
ETHDAI	Uniswap v3	Ethereum	0.01	0xD8dEC118e1215F02e10DB846DCbBfE27d477aC19
ETHDAI	Uniswap v3	Ethereum	0.05	0x60594a405d53811d3BC4766596EFD80fd545A270
ETHDAI	Uniswap v3	Ethereum	0.30	0xC2e9F25Be6257c210d7Adf0D4Cd6E3E881ba25f8
ETHDAI	Uniswap v3	Ethereum	1.00	0xa80964C5bBd1A0E95777094420555fead1A26c1e
ETHDAI	Uniswap v3	Optimism	0.05	0x95d9D28606ee55De7667f0F176eBfc3215CFD9C0

Continued on next page

Table A2.3 – Continued from previous page

Pair	Market	Chain	Fee (%)	Pair Address
ETHDAI	Uniswap v3	Optimism	0.30	0x03aF20bDAaFfB4cC0A521796a223f7D85e2aAc31
ETHDAI	Uniswap v3	Optimism	1.00	0x815AE7bF44DDa74ED9274377eD711eFC8B567911
ETHUSDC	Uniswap v3	Arbitrum	0.05	0xC31E54c7a869B9FcBEcc14363CF510d1c41fa443
ETHUSDC	Uniswap v3	Arbitrum	0.30	0x17c14D2c404D167802b16C450d3c99F88F2c4F4d
ETHUSDC	Uniswap v3	Arbitrum	1.00	0x7e5E4a3F855f19cC1a45b9eFF1c8B2419036CE85
ETHUSDC	Uniswap v2	Ethereum	0.30	0xB4e16d0168e52d35CaCD2c6185b44281Ec28C9Dc
ETHUSDC	Uniswap v3	Ethereum	0.01	0xE0554a476A092703abdB3Ef35c80e0D76d32939F
ETHUSDC	Uniswap v3	Ethereum	0.05	0x88e6A0c2dDD26FEEb64F039a2c41296FcB3f5640
ETHUSDC	Uniswap v3	Ethereum	0.30	0x8ad599c3A0ff1De082011EFDDc58f1908eb6e6D8
ETHUSDC	Uniswap v3	Ethereum	1.00	0x7BeA39867e4169DBe237d55C8242a8f2fcDcc387
ETHUSDC	Uniswap v3	Optimism	0.05	0x85149247691df622eaF1a8Bd0CaFd40BC45154a9
ETHUSDC	Uniswap v3	Optimism	0.30	0xB589969D38CE76D3d7AA319De7133bC9755fD840
ETHUSDC	Uniswap v3	Optimism	1.00	0x94ad9A19126ebb02Dda874237E5820fD4943f5dE
ETHUSDC	Quickswap	Polygon	0.30	0x853ee4b2a13f8a742d64c8f088be7ba2131f670d
ETHUSDC	Quickswap v3	Polygon		0x55caabb0d2b704fd0ef8192a7e35d8837e678207
ETHUSDC	Uniswap v3	Polygon	0.05	0x45dDa9cb7c25131DF268515131f647d726f50608
ETHUSDT	Uniswap v3	Arbitrum	0.05	0x641C00A822e8b671738d32a431a4Fb6074E5c79d
ETHUSDT	Uniswap v3	Arbitrum	0.30	0xc82819F72A9e77E2c0c3A69B3196478f44303cf4
ETHUSDT	Uniswap v3	Arbitrum	1.00	0x58039203442C9F2A45D5536bd021a383C7f3035C
ETHUSDT	Uniswap v2	Ethereum	0.30	0x0d4a11d5EEaaC28EC3F61d100daF4d40471f1852
ETHUSDT	Uniswap v3	Ethereum	0.05	0x11b815efB8f581194ae79006d24E0d814B7697F6
ETHUSDT	Uniswap v3	Ethereum	0.30	0x4e68Ccd3E89f51C3074ca5072bbAC773960dFa36
ETHUSDT	Uniswap v3	Ethereum	1.00	0xC5aF84701f98Fa483eCe78aF83F11b6C38ACA71D
ETHUSDT	Uniswap v3	Optimism	0.30	0xdd0c6bAe8Ad5998c358b823df15A2A4181Da1b80
ETHUSDT	Quickswap v3	Polygon		0x9ceff2f5138fc59eb925d270b8a7a9c02a1810f2
ETHUSDT	Uniswap v3	Polygon	0.30	0x4CcD010148379ea531D6C587CfDd60180196F9b1
FTMETH	Uniswap v2	Ethereum	0.30	0x1ffC57cAda109985aD896a69FbCEBD565dB4290e
FTMETH	Uniswap v3	Ethereum	0.30	0x64652315D86f5dfAE30885FBD29D1da05b63ADD7
FTMETH	Uniswap v3	Ethereum	1.00	0x3B685307C8611AFb2A9E83EBc8743dc20480716E
FUNETH	Uniswap v2	Ethereum	0.30	0x05B0c1D8839eF3a989B33B6b63D3aA96cB7Ec142
FUNETH	Uniswap v3	Ethereum	0.30	0xDA827fe99ADb2643D80fd30f750b8b96321d7726
FUNETH	Uniswap v3	Ethereum	0.01	0x29F0096512B4af1d689c1a11A867A6e707a8DcDe
GALAETH	Uniswap v2	Ethereum	0.30	0xbe19C32B4CD202407e8eEB73e4E2949438461Ae3
GALAETH	Uniswap v3	Ethereum	0.30	0xf8a95b2409C27678A6d18D950c5d913D5C38aB03
GALAETH	Uniswap v3	Ethereum	1.00	0x59f2044fc191f758fD3680478514244ACF253D48
GRTETH	Uniswap v2	Ethereum	0.30	0x2e81eC0B8B4022fAC83A21B2F2B4B8f5ED744D70
GRTETH	Uniswap v3	Ethereum	0.30	0x0e2c4bE9F3408E5b1FF631576D946Eb8C224b5ED
GRTETH	Uniswap v3	Ethereum	1.00	0x46add4B3F80672989b9A1eAF62caD5206F5E2164
LINKETH	Uniswap v3	Arbitrum	0.30	0x468b88941e7Cc0B88c1869d68ab6b570bCEF62Ff
LINKETH	Uniswap v2	Ethereum	0.30	0xa2107FA5B38d9bbd2C461D6EDf11B11A50F6b974
LINKETH	Uniswap v3	Ethereum	0.05	0x5d4F3C6fA16908609BAC31Ff148Bd002AA6b8c83
LINKETH	Uniswap v3	Ethereum	0.30	0xa6Cc3C2531FdaA6Ae1A3CA84c2855806728693e8
LINKETH	Uniswap v3	Ethereum	1.00	0x3A0f221eA8B150f3D3d27DE8928851aB5264bB65
LINKETH	Quickswap	Polygon	0.30	0x5ca6ca6c3709e1e6cfe74a50cf6b2b6ba2dadd67
LINKETH	Uniswap v3	Polygon	0.30	0x3e31AB7f37c048FC6574189135D108df80F0ea26
LRCETH	Uniswap v2	Ethereum	0.30	0x8878Df9E1A7c87dcBf6d3999D997f262C05D8C70
LRCETH	Uniswap v3	Ethereum	0.30	0xe1d92f1De49CAEC73514f696FEa2a7d5441498E5
MANAETH	Uniswap v2	Ethereum	0.30	0x11b1f53204d03E5529F09EB3091939e4Fd8c9CF3
MANAETH	Uniswap v3	Ethereum	0.30	0x8661aE7918C0115Af9e3691662f605e9c550dDc9
MANAETH	Uniswap v3	Ethereum	1.00	0xD9ED2b5F292a0E319a04e6C1acA15DF97705821c
MATICETH	Uniswap v2	Ethereum	0.30	0x819f3450dA6f110BA6Ea52195B3beaFa246062dE
MATICETH	Uniswap v3	Ethereum	0.05	0xe3baA96aD46457d9e6cDD4e32ABC11e2C124eC49
MATICETH	Uniswap v3	Ethereum	0.30	0x290A6a7460B308ee3F19023D2D00dE604bcf5B42
MATICETH	Uniswap v3	Ethereum	1.00	0x99C7550be72F05ec31c446cD536F8a29C89fdB77
MATICETH	Quickswap	Polygon	0.30	0xadbf1854e5883eb8aa7baf50705338739e558e5b
MATICETH	Quickswap v3	Polygon		0x479e1b71a702a595e19b6d5932cd5c863ab57ee0
MATICETH	Uniswap v3	Polygon	0.05	0x86f1d8390222A3691C28938eC7404A1661E618e0
MATICETH	Uniswap v3	Polygon	0.30	0x167384319B41F7094e62f7506409Eb38079AbfF8
MATICUSDT	Uniswap v3	Ethereum	0.30	0x68F73e2180024DB5B54E0E119d4F5128953F9417
MATICUSDT	Quickswap	Polygon	0.30	0x604229c960e5cacf2aaeac8be68ac07ba9df81c3
MATICUSDT	Quickswap v3	Polygon		0x5b41eedcfc8e0ae47493d4945aa1ae4fe05430ff
MATICUSDT	Uniswap v3	Polygon	0.05	0x781067Ef296E5C4A4203F81C593274824b7C185d
MATICUSDT	Uniswap v3	Polygon	0.30	0x781067Ef296E5C4A4203F81C593274824b7C185d

Continued on next page

Table A2.3 – Continued from previous page

Pair	Market	Chain	Fee (%)	Pair Address
SANDETH	Uniswap v2	Ethereum	0.30	0x3dd49f67E9d5Bc4C5E6634b3F70BfD9dc1b6BD74
SANDETH	Uniswap v3	Ethereum	1.00	0x5b97B125CF8aF96834F2D08c8f1291BD47724939
SANDETH	Uniswap v3	Ethereum	0.30	0x5859ebE6Fd3BBC6bD646b73a5DbB09a5D7B6e7B7
SNXETH	Uniswap v2	Ethereum	0.30	0x43AE24960e5534731Fc831386c07755A2dc33D47
SNXETH	Uniswap v3	Ethereum	0.30	0xEDe8dd046586d22625Ae7fF2708F879eF7bdb8CF
SNXETH	Uniswap v3	Ethereum	1.00	0xf5ce0293C24FD0990E0a5758E53f66a36cA0118f
SNXETH	Uniswap v3	Optimism	0.30	0x0392b358CE4547601BEFa962680BedE836606ae2
UNIETH	Uniswap v2	Ethereum	0.30	0xd3d2E2692501A5c9Ca623199D38826e513033a17
UNIETH	Uniswap v3	Ethereum	0.30	0x1d42064Fc4Beb5F8aAF85F4617AE8b3b5B8Bd801
UNIETH	Uniswap v3	Ethereum	0.05	0xfaA318479b7755b2dBfDD34dC306cb28B420Ad12
UNIETH	Uniswap v3	Ethereum	1.00	0x360b9726186C0F62cc719450685ce70280774Dc8
WBTCETH	Uniswap v3	Arbitrum	0.05	0x2f5e87C9312fa29aed5c179E456625D79015299c
WBTCETH	Uniswap v3	Arbitrum	0.30	0x149e36E72726e0BceA5c59d40df2c43F60f5A22D
WBTCETH	Uniswap v3	Arbitrum	1.00	0x99dFc0126ED31E0169fc32dB6B89adF9FeE9a77e
WBTCETH	Uniswap v2	Ethereum	0.30	0xBb2b8038a1640196FbE3e38816F3e67Cba72D940
WBTCETH	Uniswap v3	Ethereum	0.01	0xe6ff8b9A37B0fab776134636D9981Aa778c4e718
WBTCETH	Uniswap v3	Ethereum	0.05	0x4585FE77225b41b697C938B018E2Ac67Ac5a20c0
WBTCETH	Uniswap v3	Ethereum	0.30	0xCBCdF9626bC03E24f779434178A73a0B4bad62eD
WBTCETH	Uniswap v3	Ethereum	1.00	0x6Ab3bba2F41e7eAA262fa5A1A9b3932fA161526F
WBTCETH	Uniswap v3	Optimism	0.05	0x85C31FFA3706d1cce9d525a00f1C7D4A2911754c
WBTCETH	Uniswap v3	Optimism	0.30	0x73B14a78a0D396C521f954532d43fd5fFe385216
WBTCETH	Quickswap	Polygon	0.30	0xdc9232e2df177d7a12fdff6ecbab114e2231198d
WBTCETH	Quickswap v3	Polygon		0xac4494e30a85369e332bdb5230d6d694d4259dbc
WBTCETH	Uniswap v3	Polygon	0.05	0x50eaEDB835021E4A108B7290636d62E9765cc6d7
WBTCETH	Uniswap v3	Polygon	0.30	0xfe343675878100b344802A6763fd373fDeed07A4

Chapter 3

Can Automated Market Makers Provide Efficient Price Discovery?

3.1 Introduction

For a market to succeed, it must provide two core functions: liquidity and price discovery.[1] Automated Market Makers (AMMs) challenge the common understanding of how these functions operate in a successful market. Unlike traditional markets, where Liquidity Providers (LPs) actively manage and adjust prices based on real-time information, AMMs rely on passive liquidity provision. Prices are determined based on a mathematical function known to all market participants. Traders arriving at the market update prices by trading. This passive approach to liquidity provision prevents LPs from reacting to incoming order flow and consequently exposes them to adverse selection, a central challenge faced by all market makers (Glosten & Milgrom, 1985; Kyle, 1985). This heightened exposure to adverse selection increases the cost of providing liquidity, which would reduce liquidity and weaken price efficiency (Chordia, Roll, & Subrahmanyam, 2008). It therefore raises the question whether AMMs can support efficient price discovery when LPs have no ability to adjust their quoted prices in response to new information?

[1] See O'Hara (2003).

In traditional Limit Order Book (LOB) markets, LPs manage adverse selection risks by adjusting quotes based on all available information. In an AMM, LPs cannot update their quotes, instead they can step away from the market. By stepping away, the LP is no longer vulnerable to traders with superior information. However, the market is more illiquid, which would suggest that price discovery should be worse. This raises questions about the efficiency of these markets and whether AMMs can truly facilitate price discovery while avoiding the pitfalls of inefficiency and illiquidity.

This chapter explores the role of AMMs in the price discovery process. First, whether AMMs can provide efficient price discovery when they are the only market to trade an asset. Second, where and how price discovery occurs when AMMs compete with LOBs, and what drives each markets contribution to price discovery.

AMMs have the potential to change the future of financial markets, particularly in the trading of tokenized assets. AMMs offer a cost-effective solution to facilitate trading of tokenized assets, not only allowing more assets to become tradable but also improving gains from trade in existing tradable assets. Foley, O'Neill, and Putniņš (2023) suggests AMMs are a more efficient market design for trading high volume, low volatility assets. Similarly, liquidity in AMMs does not suffer from intraday illiquidity (Adams et al., 2023). Malinova and Park (2023) argue that using AMMs could save billions of dollars in transaction costs and improve capital allocation. Foreign exchange markets are often discussed as a market which would benefit from trading through an AMM. Led by the Bank of International Settlements Project Mariana shows the cost-saving potential of AMM-based trading systems while also highlighting the drawback of these markets though pre-funding of trades (Bank for International Settlements, 2023).

In this chapter, we compare two market structures: the LOB market (Binance), and an AMM (Uniswap and Quickswap). Through a variance decomposition based on the framework of Brogaard et al. (2022), we measure the informational content in an assets price that trade only on AMMs versus those traded on both AMMs and LOBs, assessing whether AMMs can independently facilitate efficient price discovery.

Our main finding is that AMMs can provide price discovery efficiently when they are the only market to trade an asset. There is no significant difference in the share of noise—the innovations in price that cannot be explained by information—between assets only traded on AMMs and a matched sample of assets that trade on both an AMM and a LOB.

We also find that LPs in AMMs react strategically to adverse selection risk by withdrawing liquidity when there is a high share of private information in prices. This behavior suggests that LPs in AMMs are more active and adaptive in their market-making strategies than previously assumed.

Next, we evaluate the AMMs ability to lead price discovery compared to LOBs. When AMMs compete with a LOB trading the same asset, we find that the LOB tends to lead the price discovery process. However, low-fee AMM pools show a stronger contribution to price discovery compared with their high-fee counterparts. Specifically, Uniswap v3 pools with 0.05% and 0.01% fees achieve average information leadership shares (ILS) of 13.6% and 67.5%, respectively, significantly greater than the sample average of 8.7%. Higher trading costs on LOBs, such as wider spreads, correspond with increased AMM price leadership. These findings highlight that AMMs can contribute to the price discovery process when they are more liquid than LOBs. However, AMMs cannot often match the liquidity of LOBs markets because of the increased cost of adverse selection due to passive liquidity provision, and therefore lag in the price discovery process.

Capponi, Jia, and Yu (2023) show that informed traders bid for priority inclusion in a block, paying higher gas fees. In contrast, at an aggregate level we do not find any significant relation between gas fees and the AMMs contribution to price discovery. When controlling for difference across blockchains, we find that that gas fees and the AMMs contribution to price discovery are significantly negatively related. Combined, these findings highlight that informed traders are not directly sensitive to overall gas costs, but are willing to pay higher fees for priority inclusion.

We show that the price discovery contribution of AMMs on Layer 2 blockchains such as Optimism and Polygon is significantly lower than on Ethereum.[2] However, AMMs on Arbitrum do not exhibit a statistically significant difference in their contribution relative to Ethereum. Our results show that more frequent blocks do not increase price informativeness. One plausible explanation is the opportunity cost associated with maintaining inventory on Layer 2 blockchains, which may deter informed traders and limit the incorporation of information into prices. Alternatively, consistent with Admati and Pfleiderer (1988), less frequent blocks may result in more information aggregation per block resulting in more informative prices on AMMs.

This chapters main contribution is to show that AMMs can provide a fundamental market function—price discovery. However, the benefits of AMMs being a cost effective market structure rely on all trading volume to flow through the AMM (Foley, O'Neill, & Putniņš, 2023; Malinova & Park, 2023). Without this volume the AMM tend to be illiquid and unable to compete with the LOB in its contribution to price discovery. Given sufficient volume, AMMs could enable a broader range of assets to become tradable and enhance the gains from trade for existing assets, realizing the full potential of tokenized asset markets.

3.1.1 Related literature

Existing research explores the role of price discovery between AMMs and Centralized Exchanges (CEXs). Capponi, Jia, and Yu (2023) show that informed traders pay higher gas fees to reduce execution risk and compete in Decentralized Exchanges (DEXs). Similarly, Klein et al. (2023) analyze how different order types and liquidity provision in AMMs contribute to price discovery. Alexander et al. (2023) examine the ETH, USDC trading pair and the contribution to price discovery across Uniswap, Coinbase and Bitstamp. They

[2]Layer 2 blockchains are scaling systems built on top of a Layer 1 or base blockchain, such as Ethereum, that execute transactions on a separate blockchain or completely off-chain and periodically settle their transactions in batches on the Layer 1 blockchain. The goal of Layer 2 blockchains is to reduce congestion and transaction costs while preserving the security of the Layer 1 blockchain.

find the AMMs share of price discovery increases with the liquidity improvements gained from the concentrated liquidity market design.

This chapter relates to studies on AMMs and DEXs. Aoyagi and Ito (2021) present a theoretical framework for the interaction of AMMs and LOB markets. Similarly, Barbon and Ranaldo (2022) quantify the costs of trading on DEXs. Lehar and Parlour (2025) show that AMMs and DEXs can function as stable markets. Heimbach, Schertenleib, and Wattenhofer (2022) provide an analysis on the profitability of liquidity provision. Moreover, Milionis, Moallemi, and Roughgarden (2024) and Milionis et al. (2022) quantify the un-diversifiable risk of providing liquidity in AMMs. Xu et al. (2022) provide an in-depth overview of the different AMM designs. We contribute to this growing area of literature by examining how well AMMs provide price discovery.

Our chapter also contributes to the literature on informational efficiency. Nguyen et al. (2024) measures the amount of noise in cryptocurrency markets. Comerton-Forde and Putniņš (2015) and Foley and Putniņš (2016) study the impact dark pools have on informational efficiency. Chordia, Roll, and Subrahmanyam (2008) and Rösch, Subrahmanyam, and van Dijk (2017) also use high frequency measures of informational efficiency. Boulatov and George (2013) show how hidden liquidity and its impact on informational efficiency. Informative prices are crucial because reduced profitability from acquiring information leads to less information production overall, resulting in less informative prices (Admati & Pfleiderer, 1988; Kyle, 1989).

There is a rich literature that focuses on the contribution of price discovery between different markets. Cabrera, Wang, and Yang (2009) explore the role of the futures market in price discovery. Several studies explore the price discovery dynamics between stock and options markets (Chakravarty, Gulen, & Mayhew, 2004; Hsieh, Lee, & Yuan, 2008; Muravyev, Pearson, & Paul Broussard, 2013; Patel et al., 2020). Ates and Wang (2005) examine the contribution of price discovery between open outcry and electronic markets. Similarly, Booth et al. (2002) study the information content of trades between upstairs and downstairs markets. Numerous studies also examine the role of price discovery across geographical regions (Chen & Choi, 2012; Eun & Sabherwal, 2003; Fricke & Menkhoff, 2011). Alexander, Heck, and Kaeck (2022) and Conlon et al. (2022) measure the price discovery in cryptocurrency markets. We add to this literature by highlighting the contribution of AMMs to price discovery in cryptocurrency markets.

3.2 Institutional background

AMMs have gained popularity as a method for trading cryptocurrencies in Decentralized Finance (DeFi). While there have been attempts to implement on-chain LOB markets, these efforts have largely been unsuccessful because of the high transaction costs on blockchains like Ethereum.[3] AMMs minimize the data storage costs required to operate a market on-chain. For users, trading on-chain offers the additional benefits of security and pseudonymity, as they keep custody of their funds (Aspris et al., 2021).

AMMs are not new to DeFi and historically they have been used in prediction markets (Hanson, 2003). Most AMM designs have not been liquidity sensitive, which has limited their use beyond prediction markets. Othman and Sandholm (2011) and Othman et al. (2013) outline invariants to deal with the challenge of liquidity sensitivity in AMMs.

In DeFi, AMMs facilitate trades because LPs add assets to a liquidity pool. Any user can trade against this liquidity pool on the blockchain. When a trader arrives at the AMM intending to trade, a mathematical function known as the pricing function (or invariant) determines the price for the trade. As trades occur in the liquidity pool, prices move with the demand of each asset. LPs earn a share of the trade volume as a fee to incentivize liquidity provision.

There are many AMM designs in DeFi. In this chapter, we focus on the largest AMM design popularized by Uniswap, the Constant Product Market Maker (CPMM) and

[3] Etherdelta was the first notable DEX on Ethereum to use a LOB design. On Solana, where transactions are cheaper, on-chain LOB are more feasible, with notable examples such as Phoenix and the historically significant Serum.

the extension of the CPMM which is known as concentrated liquidity. Uniswap v2 and Quickswap use the CPMM where the pricing function takes the form of

$$L = \sqrt{r_1 r_2}, \tag{3.1}$$

where r_1 and r_2 are the reserves of the two assets in the pool and L is a constant equal to their geometric mean.[4] When swapping against the CPMM, L is held constant by the change in reserves.

$$\sqrt{(r_1 + \Delta r_1)(r_2 + \Delta r_2)} = \sqrt{r_1 r_2} = L. \tag{3.2}$$

By holding the invariant L constant prices follow the demand for each asset. The quote price of asset r_1 in the liquidity pool is simply the ratio of the reserves $\frac{r_2}{r_1}$. To add and remove liquidity in the CPMM, LPs must hold the price constant by adding both assets proportional to their current weighting or ratio of assets in the pool. This ensures that any changes in liquidity do not impact the price in the pool. For more detail on the CPMM of Uniswap and Quickswap, we refer to Zhang, Chen, and Park (2018).

Uniswap v3 introduces the idea of "concentrated liquidity" within the CPMM. Concentrated liquidity more closely resembles traditional LOB markets with a liquidity surface that can vary with price. In the CPMM, liquidity can be thought of as being distributed equally between all prices $[0, \infty]$. In a concentrated liquidity market, liquidity is added to the pool between two prices $[p_a, p_b]$ where the amount of liquidity added follows the pricing function

$$(r_1 + \frac{L}{\sqrt{p_b}})(r_2 + L\sqrt{p_a}) = L^2. \tag{3.3}$$

In practice, prices p are bounded to discrete price ticks to maintain data storage costs. Different LPs can provide liquidity across different price ranges. The liquidity L can be added up for each position within one discrete tick interval. A user can trade against

[4]This pricing function is also commonly known as $xy = k$. However, this is not positive homogeneous and using the geometric mean of reserves solves this issue (Othman, 2021). Using the geometric mean is the standard used in Uniswap v3 and in the smart contract implementations in Uniswap v2 and Quickswap to mint the liquidity pool tokens for proportional wealth distribution.

L (reserves) in the current tick as they would in the CPMM. Intuitively, all the prices above (below) the current price in the pool are in the first (second) token in the pool. As new tokens are added to the pool, L is held constant, and the price is pushed away from the side where the tokens are being added. Trading in the pool can move the price into a new tick range where the L in the pool can change. To control for this, swaps are calculated iteratively, stepping through each tick range until the total swap amount is depleted. Adams et al. (2021) provide a detailed explanation of the Uniswap v3 pricing function and implementation.

In AMMs, the fees are often taken as a haircut of the assets entering or exiting the pool. The fee in the AMM emulates a spread in a LOB market. In most AMMs the fee is taken from the asset that enters the pool. Uniswap v2 and Quickswap has a 0.3% fee for all pools. Uniswap v3 pools can have 4 different fee levels: 1%, 0.3%, 0.05% or 0.01%. Each pool is unique and the same trading pair may different fee pools with their own liquidity. The fees for a liquidity pool in Uniswap v2, Quickswap and Uniswap v3 do not change overtime, however LPs can change the pool that they supply liquidity in order to change their fee return.

AMMs are state-based markets. The state that a user interacts with the market depends on where the transaction is included on the blockchain. This feature of AMMs means the ordering of transactions is important for traders in the market. There is value associated with being able to order the transactions in a block. This is often known as Maximal Extractable Value (MEV). MEV is created when the proposer of a block has an incentive to include or reorder the transactions in the block for profit. The proposer can earn more from a block than just the gas fees if they include other transactions, such as arbitrage between a CEX and DEX or sandwich attacks as in Daian et al. (2019).

Most of our sample is covered by Uniswap on the Ethereum blockchain. Ethereum has had two consensus models: proof of work and proof of stake. When Ethereum was based on a proof of work consensus model, extracting the value from the block was left up to the miner who won the block. This miner could build a block with their own transactions around other user transactions found in the mempool. Building an optimal block is costly and building the most optimal block is infeasible for most miners. However, with the shift towards proof of stake, Proposer Builder Separation (PBS) was introduced.

PBS separates the block building process into four roles: the searcher, the builder, the relay, and the proposer. PBS ensures the block with the highest economic value is produced by shifting the costly task of building the optimal block away from the proposer and to a builder. The role of a searcher is to sequence transactions into bundles and submit them to a builder. The searcher needs incentivize the builder to be included in the block. This is typically done in the form of transaction fees. Builders create a completed block from the bundles sent to them by searchers (or their own transactions). Their own transactions can be sandwich attacks on other users or more commonly arbitrage between CEXs and DEXs as discussed in Heimbach, Pahari, and Schertenleib (2024). Builders send these blocks to the relays with a bid to send to the proposer. The relays validate the block for correct cryptographic signatures and announce to proposers the block headers with the associated bid. If a proposer is listening to the relay, they can select the block they want to propose by signing the block header. Importantly, the proposer does not see the contents of the block, just the block header and the bid for that block. If the proposer signs a block, the relay propagates this block to the rest of the network as a complete block for validation.[5]

The rest of the sample is based on Layer 2 blockchains, specifically Polygon, Arbitrum, and Optimism. The transaction life cycle on Layer 2 blockchains like Polygon, Arbitrum, and Optimism differ from Ethereum and vary depending on the blockchain design. All Layer 2's periodically publish batches of transactions to Ethereum blockchain for full validation, ensuring security and finality. However, the sequencing process differs among these solutions. Similar to Ethereum, Polygon separates consensus and sequencing into separate roles. Sequencing information remains exposed to actors involved in block sequencing,

[5] John et al. (2025) provide an excellent overview of the Ethereum block building and validating process.

creating a potential for exploitation of transaction data (e.g., front-running or MEV). Optimism and Arbitrum both use a centralized off-chain sequencer to order transactions on a first-come, first-served basis, similar to most modern order book markets.

3.3 Price discovery in automated market makers

In this section, we develop several testable hypotheses related to the role of AMMs in the price discovery process. Price discovery refers to how new information becomes incorporated into asset prices. While the structure of AMMs differs from traditional markets, their behavior shares important economic features with theoretical models of informed trading. We focus on two central research questions: (i) whether AMMs can provide efficient price discovery as the only market, and (ii) whether AMMs can lead the price discovery process when multiple markets trade an asset.

Theoretical models such as Glosten and Milgrom (1985) describe markets where informed traders maximize profit by trading toward fundamental value, while market makers adjust quotes to manage adverse selection risk. Market makers set bid-ask spreads to break even in expectation, earning from noise traders while losing to informed traders. Prices adjust over time as dealers update beliefs based on the flow of trades.

AMMs share key elements of this framework: they face the risk of trading against informed traders and must balance adverse selection against revenue from noise traders. However, AMMs differ in the way prices are set. Instead of adjusting quotes dynamically, AMMs set prices deterministically, with a fixed fee acting as a spread. LPs cannot adjust spreads (fee) in response to adverse selection, they can decrease the amount of liquidity L they supply, stepping away when adverse selection is high. To break even, LPs must balance fee income from noise traders with expected losses to informed traders.

Despite these differences, the core intuition remains: informed traders reveal information through trading, and prices adjust accordingly. AMMs continuously quote prices and allow trades at those quotes, enabling informed agents to move prices toward fundamental value. Thus, even though the pricing is suboptimal in AMMs, the market structure still supports information generation. We hypothesize AMMs can provide efficient price discovery when they are the only venue to trade an asset.

Hypothesis 1. *AMMs can provide efficient price discovery when they are the only market to trade an asset.*

A key distinction between AMMs and traditional markets lies in pricing optimality. In theoretical models, competitive liquidity provision ensures market participants dynamically adjust quotes in response to order flow and adverse selection risk. In LOB markets, this mechanism moves asset prices toward their fundamental value (Brogaard et al., 2022). However, this adaptive behavior is slower in AMMs with their passive liquidity provision. The fixed pricing function in AMMs makes adverse selection more costly to manage. To break even, LPs must rely on a sufficient volume of noise trader activity to offset expected losses to informed traders. When such noise volume is low, trading costs rise, discouraging informed (noise) trader participation, leading to a withdrawal of liquidity from LPs.

This self-reinforcing cycle reduces liquidity and limits the profit informed traders can make on their information (Chordia, Roll, & Subrahmanyam, 2008). Suggesting that AMMs would be more expensive for informed traders relative to LOBs, especially when noise trader volume in the AMM is limited. As a result, informed traders may prefer LOBs where they can trade more cheaply. We therefore hypothesize that the AMMs will lag behind the LOB when they trade the same asset.

Hypothesis 2. *LOB prices will lead the price discovery process (i.e., contribute more than 50% to price discovery).*

However, trading costs in AMMs are variable. When noise trader activity increases, LPs increase liquidity, lowering the effective cost of trading for informed agents. In such cases, AMMs may become more attractive than LOBs. Theoretical work by Aoyagi and Ito (2021) suggests that AMMs with higher liquidity levels attract more informed traders, since greater liquidity allows them to extract more value from their information. Similarly, in US equities markets, Collin-Dufresne and Fos (2015) show that informed traders strategically time their trades based on liquidity conditions. Combined, these insights suggest AMMs can lead price discovery when there is enough liquidity to offset their suboptimal pricing.

Hypothesis 3. *AMMs will lead the price discovery process when they are more liquid.*

Informed traders face a strategic decision when choosing their trading venue. Trading on an AMM entails variable explicit costs, such as gas fees, which can make trading expensive. Informed traders must ensure favorable execution relative to other market participants, often by competing for priority inclusion within a block. To maximize their informational advantage, informed traders may bid for priority inclusion with a block builder, capturing the AMMs state at the start of the block. Capponi, Jia, and Yu (2023) provide evidence to show that informed traders actively compete for block inclusion, frequently bidding higher transaction fees to secure their position. However, this introduces a tradeoff between the value of acting on information and the costs associated with timely inclusion, as high gas fees may erode potential profits.

Hypothesis 4. *AMM prices are less informative when blockchain fees are higher.*

A common critique of AMMs is that AMMs (on Ethereum) lag centralized markets simply due to the slow discrete block times only being able to update prices as fast as the blockchain allows (e.g., every 12 seconds for Ethereum based DEXs). However, this limitation is not specific to AMMs but to the underlying infrastructure. AMMs operate across various environments, including Layer 2 blockchains with faster block times.

Faster block times do not necessarily lead to more informative prices. Drawing on Admati and Pfleiderer (1988), informed traders rely on noise trader volume to disguise their private information. When trades are batched into fewer, larger intervals, aggregate order flow can reveal more about the underlying asset value, enhancing price informativeness. In contrast, in environments with very frequent updates, informed trades are more easily detected, and the scope for camouflaging information diminishes.

Hypothesis 5. *All else equal, more batching (i.e., less frequent blocks) increases price informativeness.*

3.4 Data

We collect data on the largest exchanges of each respective market design across different platforms: Uniswap and Quickswap the AMMs, and Binance, the LOB. We collect data for Uniswap v2 on Ethereum and Uniswap v3 on Ethereum, Polygon, Arbitrum and Optimism blockchains. We collect data on Quickswap v2, a direct fork of Uniswap v2 on the Polygon blockchain. Uniswap and Quickswap are permissionless exchanges where any user can deploy their own token and to be traded. As of April 2023, Uniswap on Ethereum has roughly 170 thousand different asset pairs covering 157 thousand unique tokens. In contrast, Binance has roughly 1200 asset pairs, including over 350 unique assets.

Our analysis covers the period from June 2020 to March 2023. This period captures the launch and subsequent growth of Uniswap v2. Introduced in May 2020, the launch of Uniswap v2 serves as a significant milestone in the adoption of AMMs. We collect data on the AMMs directly from blockchain nodes for the respective blockchains. Binance trade data is collected from the Binance data vision webpage. Binance tick data is collected from Tardis.Dev. We use the MVIS® CryptoCompare Digital Assets 100 Index (MVDA) from CCData as a market index to calculate market returns for the variance decomposition.[6] MVDA is a market-cap weighted index that tracks the performance of the top 100 digital assets.

Many of the AMM pools that trade are inactive and/or illiquid. We focus our study on trading pairs with over 250 trades in the month, trade actively for three or more months and have over $500 thousand Total Value Locked (TVL) across all AMM pools. We exclude stablecoin to stablecoin pairs to prevent conflicting price discovery results because of their

[6]See https://bit.ly/3PC0uKo for MVDA fact sheet.

connection to a pegged asset.[7] Where possible, we consider the wrapped version of tokens which give them the ERC20 standard and allow them to be traded on AMMs.[8]

Table 3.1 presents the summary statistics for trade and liquidity metrics for the LOB and the AMM for pair month observations. We measure dollar volume, number of trades, the mean (median) trade size, and the realized volatility in each market. Following Chapter 2, we evaluate liquidity using four metrics for both the AMM and LOB: the half quoted spread (equivalent to the AMM fee), the effective spread, market depth (along with its relative variance).[9]

Table 3.1 shows that Binance has larger trading volumes and number of trades compared to AMM pools. However, average and median trade sizes are larger in AMM pools than on Binance. The average half-quoted spread on Binance is 0.21%, whereas the AMM fee is lower only in Uniswap v3 pools, where fees (spread) are set at 0.05% and 0.01%. Despite these lower fee structures, effective spreads tend to be larger on the AMM, likely due to the larger trade sizes in the AMM pools. These larger trade sizes also contribute to higher volatility observed in AMM prices. While the AMM shows greater average depth. These descriptive statistics highlight the structural differences between AMM and LOB markets.

Our analysis is conducted in two stages. First, we compare trading pairs that only trade on an AMM with those trading simultaneously on an AMM and a LOB. This allows us to test whether AMMs can serve as effective venues for price discovery when they are the sole market. Second, we focus solely on trading pairs listed on both an AMM and the LOB to assess whether AMMs can lead the price discovery process.

For the first sample, we get a list of tokens that trade on a CEX from CCData and identify those that trade only on Uniswap. This identification yields 243 trading pairs and 1150 pair-month observations that only trade on an AMM. To ensure comparability

[7] We also remove assets that no longer represent the same value, such as those from token contract migrations.

[8] i.e., Ethereum (ETH) pairs are compared to Wrapped Ethereum (WETH) pairs.

[9] The relative variance of depth is calculated by taking the standard deviation of hourly depth and normalizing it by the mean for each pair-month.

TABLE 3.1: Summary statistics of automated market maker and limit order book trading pairs

This table presents the summary statistics for the LOB (Panel A) and AMM (Panel B) trading pairs based on pair month observations for the sample period between June 2020 and March 2023. *Volume* is the value of trade volume in millions of dollars. *Trades* are the number of thousands of trades for the month. *Mean Trade* is the average trade size in dollars for the pair in the month. *Median Trade* is the median trade size in dollars for the pair in the month. *Realized Volatility* is the average of squared hourly mid-quote returns. The *Half Quoted Spread* for the AMM is equal to the liquidity pool fee. The Half Quoted Spread for Binance is the time weighted average half quoted spread for the month. *Effective Spread* is the volume weighted effective spread in AMM winsorized at 99% before taking the weighted average to control for large outliers. *Depth*$_{[1.5\%]}$ is the average value of orders in thousands of dollars between 0% and 1.5% of the mid-quote price for the pair-month. *Depth Variance*$_{[1.5\%]}$ is the standard deviation divided by the mean of the depth between 0% and 1.5% of the mid-quote price for the pair-month.

Panel A. Limit Order Book Summary Statistics					
	Mean	Standard Deviation	p10	Median	p90
Volume ($)	4,583	20,104	3	24	4,319
Trades	4,052.5	19,592.6	16.1	87.3	5,109.0
Mean Trade ($)	482	486	113	298	1,230
Median Trade ($)	147	161	30	91	372
Realized Volatility (%)	1.48	2.04	0.24	0.85	3.42
Half Quoted Spread (%)	0.21	0.17	0.04	0.18	0.40
Effective Spread (%)	0.25	0.30	0.03	0.18	0.51
Depth$_{[1.5\%]}$ ($)	1.81	6.12	0.01	0.07	3.27
Depth Variance$_{[1.5\%]}$ (%)	45.31	24.48	23.78	40.04	70.98

Panel B. Automated Market Maker Summary Statistics					
	Mean	Standard Deviation	p10	Median	p90
Volume ($)	314	1,281	2	15	785
Trades	48.1	149.0	0.5	3.1	107.6
Mean Trade ($)	12,663	31,560	867	4,118	24,477
Median Trade ($)	6,058	17,715	164	1,823	10,499
Realized Volatility (%)	21.44	261.69	0.07	0.72	5.48
Half Quoted Spread (%)	0.31	0.21	0.05	0.30	0.30
Effective Spread (%)	4.95	20.47	0.17	0.88	4.30
Depth$_{[1.5\%]}$ ($)	4.88	44.74	0.00	0.03	2.98
Depth Variance$_{[1.5\%]}$ (%)	81.66	127.54	7.45	38.96	174.93

between AMM-only and AMM LOB-market pairs, we employ a matched sample approach. Observations are matched without replacement by market, year, month, and nearest neighbor based on trading volume, with a tolerance of up to a 20% difference. Observations failing to meet this match criterion are discarded. Importantly, our results are robust even when unmatched trading pairs are included in the analysis. The final sample comprises 92 matched pairs and 522 pair-month observations. The second sample examines tokens listed concurrently on both the AMMs and Binance. Our final dataset for this stage includes 59 asset pairs with 1709 pair month observations.

3.5 Can automated market makers provide efficient price discovery?

Price discovery is the process by which information is reflected in asset prices (Lehmann, 2002). Efficient price discovery is the process by which new information is incorporated into prices in a timely and orderly manner. In LOB markets, most of the price discovery occurs through quotes from LPs (Brogaard, Hendershott, & Riordan, 2019; Klein et al., 2023). AMMs, however, rely on trades from liquidity demanders to adjust prices and incorporate new information. When AMMs trade alongside a LOB market, the bulk of the price discovery can be performed by the LOB with arbitrageurs updating the price in the AMM. However, Hypothesis 1 suggests that AMMs can provide efficient price discovery.

To assess whether AMMs can provide efficient price discovery, we first quantify the amount of information in an assets price using a variance decomposition. Following Brogaard et al. (2022) we measure various components of information: market-wide, public, and private, as well as noise in a price series. By decomposing the variance, we can identify the amount of information in an assets price, which we can then use to compare the AMM and LOB markets.[10]

The variance decomposition derives the amount of information in a price series from innovations in price. An asset's price is made up of two parts: m_t the efficient price and s_t the pricing error:

$$p_t = m_t + s_t. \tag{3.4}$$

[10]To control for differences arising from market design, we compare both AMM prices, where in one pair the LOB price serves as the implicit benchmark and primary vehicle for price discovery.

Pricing errors have short-run effects on price and do not have any permanent effect on price in the long run. m_t follows a random walk with drift μ and innovations w_t:

$$m_t = m_{t-1} + \mu + w_t. \tag{3.5}$$

The innovations w_t reflect new information entering the price and are unpredictable ($E_{t-1}[w_t] = 1$). The drift μ can be thought of as the discount rate over the next period. An assets return for log prices p is

$$r_t = p_t - p_{t-1} = \mu + w_t + \Delta s_t. \tag{3.6}$$

The innovations in a assets price w_t can be split into three sources of information:

$$w_t = \theta_{r_m}\varepsilon_{r_m,t} + \theta_x\varepsilon_{x,t} + \theta_r\varepsilon_{r,t}, \tag{3.7}$$

and thus an assets return is:

$$r_t = \mu + \theta_{r_m}\varepsilon_{r_m,t} + \theta_x\varepsilon_{x,t} + \theta_r\varepsilon_{r,t} + \Delta s_t, \tag{3.8}$$

where $\varepsilon_{r_m,t}$ and $\varepsilon_{x,t}$ are the unexpected innovations in market return and signed dollar volume respectively. $\varepsilon_{r,t}$ is the innovation in the assets price. Each variable represents different sources of information incorporated into an assets price: market wide information ($\theta_{r_m}\varepsilon_{r_m,t}$), firm-specific based on private information ($\theta_x\varepsilon_{x,t}$), and the remaining firm specific information not captured by trading on private information ($\theta_r\varepsilon_{r,t}$). Changes in the pricing error, Δs_t, can be correlated with the innovations in the efficient price, w_t. This model assumes that both the permanent (information) and transient (noise) components are driven by the same shocks (Beveridge & Nelson, 1981). This assumption explicitly allows for correlation between information and noise, consistent with theory.

We estimate these components for pair month observations using a structural VAR model based on hourly returns with 12 hours of lags to account for serial correlation and any other lagged effects:

$$\begin{aligned} r_{m,t} &= \sum_{l=1}^{12} a_{1,l} r_{m,t-l} + \sum_{l=1}^{12} a_{2,l} x_{t-l} + \sum_{l=1}^{12} a_{3,l} r_{t-l} + \epsilon_{r_m,t}, \\ x_t &= \sum_{l=1}^{12} b_{1,l} r_{m,t-l} + \sum_{l=1}^{12} b_{2,l} x_{t-l} + \sum_{l=1}^{12} b_{3,l} r_{t-l} + \epsilon_{x,t}, \\ r_t &= \sum_{l=1}^{12} c_{1,l} r_{m,t-l} + \sum_{l=1}^{12} c_{2,l} x_{t-l} + \sum_{l=1}^{12} c_{3,l} r_{t-l} + \epsilon_{r,t}, \end{aligned} \tag{3.9}$$

where $r_{m,t}$ is the market return, x_t is the signed dollar volume of trading in the given asset (positive values for net buying and negative values for net selling), and r_t is the asset return.

By explicitly modeling the contemporaneous relationships between variables, the structural VAR innovations $\{\epsilon_{r_m t}, \epsilon_{x_t}, \epsilon_{r_t}\}$ are contemporaneously uncorrelated. The VAR model identifies market-wide, private, and public information components by capturing their long-term effects on asset prices through cumulative impulse responses $\{\theta_{r_m,t}, \theta_{x,t}, \theta_{r,t}\}$. Permanent returns linked to market-return shocks estimate market-wide information, shocks to signed dollar volume reveal private information, and shocks to asset returns—controlling for market and trading variables—highlight public information. Noise emerges as the transitory return component after removing information-driven influences. The cumulative impulse responses then quantify the contributions of each information component to price variance within the structural VAR framework.

Taking the variance of the innovations in the efficient price we get $\sigma_w^2 = \theta_{r_m}^2 \sigma_{\epsilon_{r_m}}^2 + \theta_x^2 \sigma_{\epsilon_x}^2 + \theta_r^2 \sigma_{\epsilon_r}^2$. The contribution to the variation in the efficient price from each of the information components is $\theta_{r_m}^2 \sigma_{\epsilon_{r_m}}^2$ (market-wide information), $\theta_x^2 \sigma_{\epsilon_x}^2$ (private firm-specific information), and $\theta_r^2 \sigma_{\epsilon_r}^2$ (public firm-specific information). The estimated components of variance can be normalized to give shares of the variance in a price series.

$$
\begin{aligned}
MarketInfoShare &= \theta_{r_m}^2 \sigma_{\varepsilon_{r_m}}^2 / (\sigma_w^2 + \sigma_s^2), \\
PrivateInfoShare &= \theta_x^2 \sigma_{\varepsilon_x}^2 / (\sigma_w^2 + \sigma_s^2), \\
PrivateInfoShare &= \theta_r^2 \sigma_{\varepsilon_r}^2 / (\sigma_w^2 + \sigma_s^2), \\
NoiseShare &= \sigma_s^2 / (\sigma_w^2 + \sigma_s^2),
\end{aligned}
\tag{3.10}
$$

$MarketInfoShare$, $PrivateInfoShare$, and $PublicInfoShare$, are the corresponding shares of variance from those various sources of information. Meanwhile, $NoiseShare$ reflects the relative importance of pricing errors due to illiquidity, an over (under) reaction to information, or other microstructure frictions. We get the inputs for the VAR using the procedure in Appendix 3.1.

Table 3.2 provides summary statistics on the variance decomposition shares, particularly highlighting the noise component. These estimates allow us to quantify the net level of unexplained noise within asset prices. Focusing on this noise component, we assess whether AMMs alone can support effective price discovery. Both samples of trading pairs reveal substantial noise and private information shares, each exceeding 40%. This result aligns with Nguyen et al. (2024), who find that variance in cryptocurrency prices is largely driven by noise with a roughly 47% noise share.[11]

When an AMM is the only market available to trade an asset, the amount of information in its prices may be limited because of the market structure of AMMs. AMMs operate on blockchains which require traders to pay an explicit fee (usually as gas) to get inclusion in the block. In combination with AMMs needing trades to update prices, this may impact the amount of information in prices.

MEV, the profit that block producers or transaction orders can extract by controlling the ordering, inclusion, or exclusion of transactions within a block, notably sandwich attacks—where the block builder (searcher) front runs and back runs another trade to

[11] Appendix 3.2 provides summary statistics for the LOB variance shares.

TABLE 3.2: Summary statistics for shares of variance by market

This table presents the summary statistics for the LOB (Panel A) and AMM (Panel B) trading pairs based on pair month observations for the sample period between June 2020 and March 2023. $NoiseShare$, $MarketInfoShare$, $PrivateInfoShare$, and $PublicInfoShare$, are the shares of variance explained by noise, market information, private information and public information respectively.

	Mean	Standard Deviation	p10	Median	p90
Panel A. AMM and LOB Summary Statistics					
$MarketInfoShare$ (%)	7.44	14.86	0.04	2.24	17.58
$PrivateInfoShare$ (%)	42.35	25.03	4.99	45.22	73.88
$PublicInfoShare$ (%)	5.45	9.37	0.21	2.14	11.72
$NoiseShare$ (%)	44.76	25.11	17.59	39.67	88.37
Panel B. AMM Only Summary Statistics					
$MarketInfoShare$ (%)	5.64	11.67	0.05	1.23	14.77
$PrivateInfoShare$ (%)	45.02	21.87	11.47	48.85	71.57
$PublicInfoShare$ (%)	4.80	5.73	0.21	2.99	12.47
$NoiseShare$ (%)	44.53	19.82	20.07	43.34	72.32

push the price in the pool to extreme levels often with large trade sizes—may increase the amount of noise in the variance decomposition. Traders often protect themselves from sandwich attacks using tools such as slippage tolerances (Chemaya et al., 2024).[12] However, the default values set for these protections could introduce further noise into the analysis. There is a structural relationship between trade prices and the mid-quote prices before and after a trade, as trade prices are the geometric mean between these mid-quote prices (Angeris et al., 2023). This relationship impacts LPs, as noted by Milionis et al. (2022), through the concept of loss verse rebalancing, representing the undiversifiable risk of providing liquidity in an AMM. Structurally, this results in post-trade mid-quote prices that deviate from trade prices, suggesting noisier prices in AMMs compared to LOB markets. We control for this difference in noise by considering the AMM prices only and by taking the price at the end of the block.

The impact of liquidity on the levels of information in AMMs prices is ambiguous. On the one hand, liquidity reduces price impact and noise in the price series (Chordia, Roll, & Subrahmanyam, 2008). It also acts as an incentive for informed traders to participate, as a more liquid market allows them to profit more from their information. This encourages greater information acquisition, leading to increased overall information production and

[12] Slippage tolerance on DEXs is a user-specified threshold that sets the maximum price a trader is willing to accept when trading on an AMM. If the execution price of the trade exceeds this tolerance, the transaction automatically reverts, preventing execution at unfavorable prices.

more informative prices (Admati & Pfleiderer, 1988; Kyle, 1989). On the other hand, the presence of informed traders can negatively impact liquidity due to adverse selection. This can cause liquidity being withdrawn from the AMM, worsening informational efficiency. Adverse selection is a well-documented source of illiquidity in markets (Glosten & Milgrom, 1985; Kyle, 1985).

To test Hypothesis 1, if AMMs can provide efficient price discovery, we compare variance shares of pairs that only trade on the AMM with matched pairs that trade on the AMM and LOB. We regress the different variance shares on a dummy variable for if the AMM is the only market to trade the asset. This regression examines if the amount of information in the AMMs price is lower when they are the only market to trade the asset. We estimate the following regression:

$$VarShare_{i,t} = \alpha_{i,t} + \beta_1 AMMOnly_{i,t} + \beta_2 Liquidity_{i,t} + \beta Controls_{i,t} + \epsilon_{i,t}, \tag{3.11}$$

where $VarShare$ is a measure of variance shares such as the $NoiseShare$, $MarketInfoShare$, $PrivateInfoShare$, or the $PublicInfoShare$. The independent variables of focus will be the AMMOnly dummy variable. This dummy variable will show whether the AMM has more (less) information when it is the only market to trade an asset. The inherit benchmark is the LOB information content. We include controls for market liquidity due to the interconnection between adverse selection costs and illiquidity (Glosten & Milgrom, 1985; Kyle, 1985). Specifically, we consider three measures of liquidity, the effective spread, the 1.5% depth and the variance of the 1.5% depth. Each measure captures a different aspect of liquidity: width, depth, and resiliency, respectively. We also control for the realized volatility, whether the market uses concentrated liquidity and explicit (gas) costs. All continuous variables are in are in natural log form.

Table 3.3 presents regression coefficients from Equation 3.11. The AMMOnly dummy variable in Regression 1 shows that the $NoiseShare$ is significantly lower for pairs traded only on AMMs compared with those trading across both AMMs and LOB. While the reduction in noise share is marginal and not economically significant, it highlights a key finding in support of our hypothesis: AMMs can provide price discovery when they are the only market to trade an asset. There is no less information in prices compared to assets that also trade on a LOB.

TABLE 3.3: Regressions to assess the determinants variance shares

This table reports coefficient estimates from OLS regressions using pair-month observations to examine the determinants of the AMM variance shares. The dependent variables $NoiseShare$, $MarketInfoShare$, $PrivateInfoShare$, and $PublicInfoShare$, are the shares of variance explained by noise, market information, private information and public information respectively (Regressions 1-4). $AMMOnly$ is a dummy variables that equal to 1 if the pair only trades on a AMM, otherwise 0. $EffectiveSpread$ is the volume weighted effective spread in the AMM winsorized at 99% before taking the weighted average to control for large outliers. $Depth_{[1.5\%]}$ is the average value of orders in thousands of dollars between 0% and 1.5% of the mid-quote price for the pair-month on the AMM. $DepthVariance_{[1.5\%]}$ is the standard deviation divided by the mean of the depth between 0% and 1.5% of the mid-quote price for the pair-month on the AMM. $Volume$ is the value of trade volume in dollars on the AMM. $RealizedVolatility$ is the average of squared hourly mid-quote returns on Binance. Gas is the average price of gas on the blockchain in dollars in a month. $CLCPMM$ is a dummy variable for that is equal to one if the market uses concentrated liquidity from Uniswap v3. All explanatory variables are in natural log form. The sample comprises 1709 observations from 59 trading pairs during the period June 2020 to March 2023. Standard errors are clustered by pair month. t-statistics are reported in parentheses. Significance at the 10%, 5%, and 1% levels is indicated by *, **, and ***, respectively.

	(1)	(2)	(3)	(4)
Dependent variable:	$NoiseShare$	$MarketInfoShare$	$PrivateInfoShare$	$PublicInfoShare$
$AMMOnly$	−0.036**	−0.004	0.039**	0.002
	(−2.26)	(−0.42)	(2.09)	(0.31)
$EffectiveSpread$	0.092***	0.002	−0.103***	0.009
	(5.51)	(0.27)	(−6.74)	(0.95)
$Depth_{[1.5\%]}$	0.004	0.014*	−0.020**	0.002
	(0.76)	(1.84)	(−2.22)	(0.56)
$DepthVariance_{[1.5\%]}$	−0.042***	0.017	−0.003	0.027***
	(−3.76)	(1.50)	(−0.19)	(4.39)
$Volume$	−0.035***	0.014**	0.026**	−0.006
	(−4.59)	(2.14)	(2.50)	(−1.53)
$RealizedVolatility$	0.052***	−0.024***	−0.017	−0.012*
	(4.61)	(−3.52)	(−1.35)	(−1.95)
Gas	0.007	−0.006	−0.007	0.006*
	(0.72)	(−0.79)	(−0.60)	(1.73)
$CLCPMM$	0.121***	0.027	−0.123***	−0.025***
	(4.76)	(1.53)	(−4.07)	(−2.64)
Intercept	1.712***	−0.514***	−0.369*	0.171**
	(10.96)	(−3.38)	(−1.71)	(2.40)
Observations	522	522	522	522
Adjusted R^2	39.3%	20.3%	20.8%	6.4%

Significance: * $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$

The observed difference in $NoiseShare$ appears primarily driven by variations in private information, as noted by Regression 3. Token listings on CEXs are non random and could lead to biases in our results. In contrast, DEXs operate with a permissionless listing process, allowing anyone to list new asset pairs. This distinction likely affects the type of information reflected in prices—whether private or public—while having a lesser impact on overall share of noise. Importantly, any potential bias from this process would likely skew results against our findings, reinforcing the robustness of AMMs price discovery.

$PrivateInfoShare$ is negatively related to market depth and effective spread. This suggests that LPs in AMMs step away from the market when adverse selection is high, reducing the liquidity in the market. LPs strategically withdraw liquidity when they perceive high levels of private information, protecting themselves from unfavorable trades against informed participants.

Our analysis also shows that explicit transaction costs, such as gas fees, do not significantly affect the information or noise shares estimated in the variance decomposition. However, concentrated liquidity pools as used in Uniswap v3 are positively related to the $NoiseShare$, largely because of reduced levels of both private and public information.

3.6 Can automated market makers lead price discovery?

Information must enter AMM prices, regardless of whether it originates from an arbitrage opportunity with a CEX or from an informed trader trading in the AMM first. In section 3.3 we develop several testable hypothesis to related to AMM price leadership.

AMMs face informational disadvantage because of their market structure. In AMMs, prices cannot be updated without a trade, meaning that price adjustments rely solely on the actions of liquidity takers, who may possess varying levels of information. As a result, the price set by these traders can leave AMMs exposed to adverse selection, as the AMM price only becomes informed post-trade. This dynamic underscores the potential inefficiencies in AMMs, as LPs remain uninformed about the true market price until after the trade occurs, highlighting the unique challenges that AMMs face in the price discovery process.

To test our hypotheses related to AMMs price leadership, we measure each markets contribution of new information in prices by calculating the price discovery shares for

Uniswap and Binance through ILS, extending the method of Putniņš (2013) to be used for n markets following Patel et al. (2020).

Other widely used empirical measures of price discovery include the information share (IS) proposed by Hasbrouck (1995) and the component share (CS) from Gonzalo and Granger (1995). Both measures decompose price changes into permanent and temporary components. The IS metric evaluates whether a market efficiently incorporates new information while minimizing the impact of liquidity shocks (i.e., noise), whereas the CS metric assesses the sensitivity of each market to lagged transitory shocks from the other. Both measures effectively capture the relative avoidance of noise. Table 3.1 highlights the difference in noise (realized volatility) between Uniswap and Binance. To account for this, we use the ILS measure, which focuses on incorporating information into prices, independent of the noise differences between markets. The IS is widely used in literature and we also report the results following the IS in Appendix 3.5.

We estimate a reduced form Vector Error Correction Model (VECM) model for each pair month based on end of block mid-quote prices (p_t) for each market that trades the asset. We estimate the VECM with 300 lags, equating to roughly one hour based on a 12-second block time for the slowest market in our sample (Ethereum AMMs) (details on the estimation process are provided in Appendix 3.3).[13]

$$\Delta p_t = aZ_{t-1} + \sum_{i=1}^{300} b_i \Delta p_{t-i} + \epsilon_t, \tag{3.12}$$

where Δp_t is the $n \times 1$ mid-quote return vector, α is the $n \times (n-1)$ matrix of error correction coefficients, Z_{t-1} is the $n \times 1$ co-integrating vector, b_i is the $n \times n$ coefficient matrix for lag i and ϵ_t is the $n \times 1$ vector of residuals.

We calculate IS_1 to IS_n, and CS_1 to CS_n, from the VECM estimates following Hasbrouck (1995) and Gonzalo and Granger (1995) which is detailed in Appendix 3.3.

[13] Under Ethereum's previous proof-of-work consensus, block timestamps varied with the likelihood of block completion. However, since the transition to proof-of-stake, blocks are more consistently spaced at approximately 12-second intervals, only deviating when a slot is missed.

Next, we calculate market i's propensity to reflect new information (how much market i's price responds to an innovation in the efficient price) through the ratio $\beta_i = \frac{IS_i}{CS_i}$, Normalizing the information leadership propensities gives ILS:

$$ILS_n = \frac{\beta_n^2}{\sum_{i=1}^{n} \beta_i^2}. \tag{3.13}$$

The choice of sampling frequency for VECM and price discovery share estimation can significantly affect results. Sampling can be based on various clocks and frequencies, such as fixed time intervals, blockchain block times, or even trade or volume clocks. In our study, we analyze markets operating in different time regimes: one in continuous time and others in discrete time. Selecting a frequency that is too high can disproportionately favor the continuous-time (faster) market, introducing bias.

Alternative approaches, such as sampling based on trades or volume, could mitigate this issue.[14] However, significant differences between AMMs and LOBs in terms of trade frequency and trade size could also skew the results. More active markets or those experiencing larger liquidity shocks from large trades might artificially appear to contribute more to price discovery. We sample prices from both markets at the end of each block to address these challenges.[15] This approach ensures that each market has sufficient opportunity to update prices within the sampling period, providing a fair basis for comparison.

AMMs rely on trades to update prices, which suggests that they would be a slow market that merely mirrors price movements on CEXs through arbitrage when the AMM quotes a stale price. This assumption holds at an aggregate level, across all AMM pools, the average ILS is 4.3%. However, when examining individual pools with varying fee structures, the low-fee pools contribute meaningfully to price discovery. For instance, the Uniswap v3 pool with a 0.05% fee has an average ILS of 13.6%. Similarly, the 0.01% fee pool exhibits an even higher ILS at 67.5%, although this pool predominantly trades

[14] As discussed in Hasbrouck (2021)

[15] There might be a slight discrepancy between the block's on-chain timestamp and the actual time it was processed in a distributed system due to unsynchronized block times. The impact of this deviation on our results would be minimal.

FUNToken against Ethereum and other major stablecoins, which is likely to bias the results.[16] In contrast, other AMM pools have very low price discovery shares, confirming Hypothesis 2 that prices on the AMM follow the LOB.

Figure 3.1 plots the monthly average ILS for different Uniswap pools, highlighting the contribution of the Uniswap v3 0.05% and 0.01% pools compared to Uniswap v2, Uniswap v3 0.3%, and Uniswap v3 1% pools.[17] Notably, prior to the release of Uniswap v3, the Uniswap v2 pools exhibited an ILS around 10%, which significantly declined after Uniswap v3 was introduced. This figure provides partial support for the idea that low fee or low cost AMM pools having a meaningful contribution to the price discovery process.

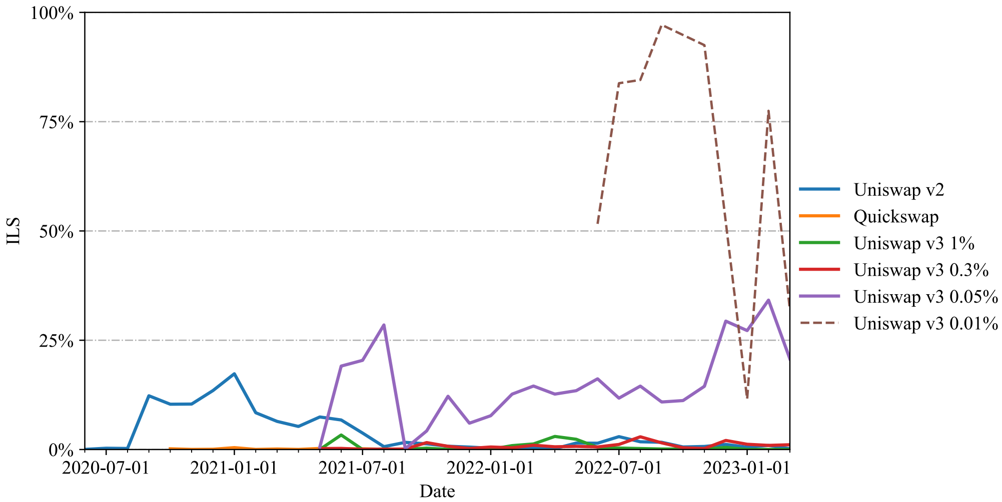

FIGURE 3.1: Automated market maker price discovery shares

This figure plots the Yan-Zivot-Putnins information leadership share (ILS) for AMM pools from 58 trading pairs during the period June 2020 to March 2023. The plotted ILS is the mean price discovery share for each different Uniswap and Quickswap market (Uniswap v2, Quickswap, Uniswap v3 1%, Uniswap v3 0.3%, Uniswap v3 0.05%, and Uniswap v3 0.01%).

Across the different blockchains in our sample the average ILS for AMMs are consistently low: Ethereum (4.8%), Arbitrum (3.8%), Optimism (1.2%), and Polygon (1.4%). These findings suggest that AMMs contribute marginally to price discovery across different blockchains.

[16] As a robustness test, we remove the 0.01% fee pools from our sample. We find our results are robust to excluding these pools.

[17] Figure A3.1 presents the same figure plotting the information share. The results are consistent with Figure 3.1

Figure 3.1 illustrates that AMMs, particularly the low-fee pools, can lead price discovery. To statistically investigate the factors driving price leadership we regress liquidity measures, trading costs, and volatility on the price discovery shares (ILS) of the AMM markets. We estimate the following regression model:

$$PDS_{i,t} = \alpha_{i,t} + Liquidity_{i,t} + Volatility_{i,t} + Gas_t + Controls_{i,t} + \epsilon_{i,t}, \tag{3.14}$$

where PDS is the AMM price discovery share (ILS). We consider three measures of liquidity, the effective spread, the 1.5% depth and the variance of the 1.5% depth. Each measure of liquidity captures a different aspect of liquidity: width, depth, and resiliency, respectively. We consider the explicit cost of inclusion on-chain by measuring the value of gas on the respective blockchain. Volatility is measured as the realized volatility (the average of squared returns) in the LOB market. In this regression, we control for trade volume, pair fixed effects and time fixed effects. All continuous variables are in natural log form.

Table 3.4 presents the coefficient estimates from Ordinary Least Squares (OLS) regressions based on Equation 3.14. To test whether the AMM leads price discovery when it is more liquid. The quoted spread (fee) in the AMM shows a significant negative relationship with both the ILS and IS, consistent with the results from Figure 3.1. These findings suggest that the fee in the AMM plays a significant role in whether the AMM contributes to price discovery. However, on its own, the effective spread does not exhibit a significant impact on AMM price discovery shares. When combined with other measures of liquidity or pair and time fixed effects, the effective spread is positively related to the AMM price discovery shares. This is likely because of the correlation between effective spread and trade size in the AMM. The AMM price discovery share is also significantly and positively related to the volume on the AMM. Overall, when using spreads to assess AMM liquidity, our findings suggest the AMM leads price discovery when it is more liquid.

We also assess AMM liquidity by considering market depth. A positive relationship between depth and the AMMs price discovery share is expected, as greater depth should

TABLE 3.4: Regressions to assess the determinants of automated market maker price discovery

This table reports coefficient estimates from OLS regressions using pair-month observations to examine the determinants of the AMM price discovery shares. The dependent variable is the Yan-Zivot-Putnins information leadership share (ILS). *QuotedSpread* is the fee in the AMM. $EffectiveSpread$ is the volume weighted effective spread in the AMM winsorized at 99% before taking the weighted average to control for large outliers. $Depth_{[1.5\%]}$ is the average value of orders in thousands of dollars between 0% and 1.5% of the mid-quote price for the pair-month on the AMM. $DepthVariance_{[1.5\%]}$ is the standard deviation divided by the mean of the depth between 0% and 1.5% of the mid-quote price for the pair-month on the AMM. $Volume$ is the value of trade volume in dollars on the AMM. $RealizedVolatility$ is the average of squared hourly mid-quote returns on Binance. *Gas* is the average price of gas on the blockchain in dollars in a month. All explanatory variables are in natural log form. The sample comprises 1709 observations from 59 trading pairs during the period June 2020 to March 2023. Standard errors are clustered by pair month. t-statistics are reported in parentheses. Significance at the 10%, 5%, and 1% levels is indicated by *, **, and ***, respectively.

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Dependent variable:	ILS	ILS	ILS	ILS	ILS	ILS	ILS	ILS
$QuotedSpread$	−0.079***			−0.078***	−0.068***			−0.069***
	(−7.20)			(−7.08)	(−6.86)			(−7.17)
$EffectiveSpread$		0.007		0.012**		0.011**		0.014***
		(1.27)		(2.46)		(2.15)		(2.79)
$Depth_{[1.5\%]}$			−0.011***	−0.007***			−0.005**	0.000
			(−4.89)	(−3.91)			(−2.43)	(0.09)
$DepthVariance_{[1.5\%]}$			−0.025***	−0.017***			−0.001	0.006
			(−5.87)	(−4.52)			(−0.26)	(1.64)
$Volume$	0.004	0.017***	0.028***	0.016***	0.007***	0.023***	0.025***	0.010***
	(1.49)	(5.22)	(7.80)	(4.36)	(2.58)	(6.66)	(6.75)	(3.05)
$RealizedVolatility$	0.015***	0.003	−0.007	0.003				
	(4.12)	(0.78)	(−1.62)	(0.88)				
Gas	0.000	0.000	0.000	0.000				
	(0.19)	(−0.04)	(0.15)	(0.04)				
Intercept	−0.346***	−0.186***	−0.419***	−0.561***	−0.434***	−0.302***	−0.307***	−0.426***
	(−5.53)	(−3.59)	(−6.12)	(−7.12)	(−5.22)	(−3.97)	(−3.91)	(−5.03)
Pair Effects	N	N	N	N	Y	Y	Y	Y
Month Effects	N	N	N	N	Y	Y	Y	Y
Observations	1,709	1,709	1,709	1,709	1,709	1,709	1,709	1,709
Adjusted R^2	15.7%	4.6%	9.4%	19.3%	37.4%	31.9%	31.9%	37.9%

Significance: * $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$

imply that informed traders are more likely to profit from their information through the AMM. However, the results show a negative relationship between depth and the AMMs price discovery share. This finding contradicts the notion that informed traders are more likely to trade in the AMM when it is liquid. One explanation is that LPs in the AMM withdraw liquidity when adverse selection costs from informed traders are high. However, this seems unlikely, as the AMM must adjust prices through trading regardless of whether the information is new or arises from arbitrage in another market. The significance and magnitude of the relationship between AMM price discovery shares and depth drops when controlling for pair and month fixed effects highlighting that the effect is driven by differences across trading pairs rather than changes within a given pair over time.

The variance of depth shows a negative relationship with the AMMs price discovery share, though this becomes insignificant when controlling for pair and month fixed effects. This suggests that the AMM price leadership is driven by persistent characteristics in some trading pairs and temporary fluctuations in depth do not significantly impact the AMMs role in the price discovery process. There is no clear relationship between realized volatility and AMM price leadership. Price deviations do not appear to significantly impact whether the AMM leads in price discovery.

We do not find a significant relationship between explicit costs (gas fees) and AMM information leadership as suggested by Hypothesis 4. This result contrasts with Capponi, Jia, and Yu (2023), who show that informed traders bid for priority inclusion in blockchains, and Barbon and Ranaldo (2022), who attribute inefficiencies in AMMs to high explicit costs. Both studies only focus on the Ethereum blockchain. One explanation for our findings is the inclusion of Layer 2 AMM pools in our analysis, where gas fees are significantly lower, potentially mitigating the impact of explicit costs on AMM informational efficiency.

Next we consider whether the liquidity in the LOB, impacts the AMM information leadership. Table 3.5 presents the coefficient estimates from OLS regressions based on Equation 3.14 using measures of LOB liquidity. We observe that AMM price leadership is positively related to the quoted spread in the LOB. After controlling for pair and month fixed effects, AMM price leadership also shows a positive relationship with the effective spread. These findings suggest the AMM leads price discovery when the LOB is less liquid, suggesting higher LOB costs drive informed traders toward the AMM.

We also consider whether the AMMs price discovery share is impacted by aspects of the CEXs in relative terms, such as the AMM being the main liquid market where price discovery would be more likely to occur. To test this, we regress on relative measures of spread, depth, and volatility following Equation 3.14. The relative measures are the log difference between the AMM and LOB measures used to examine the determinants of the price discovery shares in Table 3.4.

TABLE 3.5: Regressions to assess the determinants of automated market maker price discovery from limit order book variables

This table reports coefficient estimates from OLS regressions using pair-month observations to examine the limit order book determinants of the AMM price discovery shares. The dependent variable is the Yan-Zivot-Putnins information leadership share (ILS). $LOBQuotedSpread$ is the average half quoted spread on Binance. $LOBEffectiveSpread$ is the volume weighted effective spread on Binance. $LOBDepth_{[1.5\%]}$ is the average value of orders in thousands of dollars between 0% and 1.5% of the mid-quote price for the pair-month on Binance. $LOBDepthVariance_{[1.5\%]}$ is the standard deviation divided by the mean of the depth between 0% and 1.5% of the mid-quote price for the pair-month on Binance. $LOBVolume$ is the value of trade volume in dollars on Binance. $RealizedVolatility$ is the average of squared hourly mid-quote returns on Binance. All explanatory variables are in natural log form. The sample comprises 1709 observations from 59 trading pairs during the period June 2020 to March 2023. Standard errors are clustered by pair month. t-statistics are reported in parentheses. Significance at the 10%, 5%, and 1% levels is indicated by *, **, and ***, respectively.

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Dependent variable:	ILS	ILS	ILS	ILS	ILS	ILS	ILS	ILS
$LOBQuotedSpread$	0.011***			0.016***	0.004			0.000
	(3.73)			(3.78)	(1.31)			(−0.11)
$LOBEffectiveSpread$		−0.004		−0.022***		0.016**		0.012
		(−0.39)		(−2.89)		(2.08)		(1.49)
$LOBDepth_{[1.5\%]}$			−0.020**	−0.013**			−0.033***	−0.031***
			(−2.00)	(−2.05)			(−3.86)	(−3.54)
$LOBDepthVariance_{[1.5\%]}$			−0.005	−0.010			−0.012	−0.013
			(−0.57)	(−1.08)			(−1.27)	(−1.34)
$LOBVolume$	0.000	−0.004	0.011*	0.003	−0.019***	−0.017***	−0.006	−0.006
	(0.03)	(−1.48)	(1.72)	(0.94)	(−3.38)	(−3.14)	(−1.00)	(−0.90)
$RealizedVolatility$		0.002	−0.010					
		(0.24)	(−1.39)					
Intercept	0.117***	0.095**	−0.019	0.088***	0.368***	0.416***	0.478***	0.514***
	(6.04)	(2.16)	(−0.30)	(4.13)	(3.20)	(3.32)	(3.90)	(4.00)
Pair Effects	N	N	N	N	Y	Y	Y	Y
Month Effects	N	N	N	N	Y	Y	Y	Y
Observations	1,709	1,709	1,709	1,709	1,709	1,709	1,709	1,709
Adjusted R^2	0.6%	0.0%	0.4%	1.1%	28.2%	28.3%	28.6%	28.6%

Significance: * $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$

Table 3.6 reports coefficient estimates from OLS regressions. Consistent with our findings in Table 3.4 we find a significant negative relation between the difference in the quoted spread and the AMM price discovery share. We also find a significant positive relation between the volume ratio and the AMM price discovery share. We do not find a significant relation between the difference in the effective spread and the AMM price discovery shares.

Our findings for the ratio of quoted spread are consistent with this hypothesis. However, the findings in Table 3.6 show that when controlling for pair and date fixed effects, the difference between the depth in the AMM and the LOB is negatively related to the price discovery share in the AMM. This finding suggests that a deeper LOB or less liquid AMM is positively related to the AMMs price discovery share. We do not find any significant relation between the variance of depth and the AMM price discovery share.

TABLE 3.6: Regressions to assess the relative determinants of automated market maker price discovery

This table reports coefficient estimates from OLS regressions using pair-month observations to examine the determinants of the AMM price discovery shares. The dependent variable is the Yan-Zivot-Putnins information leadership share ILS. $QuotedSpreadRatio$ is a ratio between the fee in the AMM and the average half quoted spread for Binance. $EffectiveSpreadRatio$ is the ratio of the volume weighted effective spread between the AMM and Binance. $Depth_{[1.5\%]}Ratio$ is the ratio of the average value of orders in dollars between 0% and 1.5% of the mid-quote price for the AMM and Binance, respectively. $DepthVariance_{[1.5\%]}Ratio$ is the ratio of the standard deviation divided by the average of orders in dollars between 0% and 1.5% of the mid-quote price for the AMM and Binance, respectively. $VolumeRatio$ is the ratio of the value of trade volume in dollars between the AMM and Binance. Gas is the average price of gas on the blockchain in dollars in a month. All explanatory variables are in natural log form. The sample comprises 1709 observations from 59 trading pairs during the period June 2020 to March 2023. Standard errors are clustered by pair month. t-statistics are reported in parentheses. Significance at the 10%, 5%, and 1% levels is indicated by *, **, and ***, respectively.

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Dependent variable:	ILS	ILS	ILS	ILS	ILS	ILS	ILS	ILS
$QuotedSpreadRatio$	−0.030***			−0.037***	−0.032***			−0.036***
	(−6.84)			(−6.74)	(−6.17)			(−6.36)
$EffectiveSpreadRatio$		0.008		0.024***		0.009		0.018***
		(1.17)		(3.58)		(1.53)		(2.82)
$Depth_{[1.5\%]}Ratio$			−0.007***	−0.003*			−0.005**	−0.002
			(−3.18)	(−1.69)			(−2.13)	(−0.82)
$DepthVarience_{[1.5\%]}Ratio$			−0.010***	0.000			0.001	0.006*
			(−2.88)	(−0.11)			(0.23)	(1.68)
$VolumeRatio$	0.008***	0.016***	0.018***	0.014***	0.016***	0.025***	0.026***	0.021***
	(5.30)	(6.35)	(7.61)	(5.28)	(6.00)	(6.73)	(7.20)	(5.70)
Intercept	0.080***	0.047***	0.055***	0.041***	0.058	−0.011	0.014	0.036
	(10.76)	(3.94)	(10.10)	(3.75)	(1.03)	(−0.19)	(0.26)	(0.60)
Pair Effects	N	N	N	N	Y	Y	Y	Y
Month Effects	N	N	N	N	Y	Y	Y	Y
Observations	1,709	1,709	1,709	1,709	1,709	1,709	1,709	1,709
Adjusted R^2	10.8%	5.7%	6.9%	13.2%	35.4%	32.8%	32.9%	36.3%

Significance: * $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$

Aoyagi and Ito (2021) theorizes that informed trading in the AMM should reduce adverse selection in the LOB, reducing spreads in the LOB market. Contrary to this theory, our results suggest when information first enters the market through the AMM, spreads in the LOB are higher. It is possible, but unlikely, that market makers do not learn from the AMM price. However, when liquidity in the LOB is low (i.e., high spreads) the AMM is an attractive market for informed traders as they can profit more from their information. Informed traders appear to be sensitive to the cost structure in each market, limiting their engagement with the AMM when these costs outweigh potential profit from their information. This dynamic may explain why we do not observe a significant reduction in LOB spreads, despite informed trader presence in the AMM.

Table 3.6 suggests that the share of volume is a significant determinant of AMM information leadership. To explore this relationship further, we examine whether the share of AMM volume influences the determinants of AMM price discovery. An exchanges share of volume is expected to correlate with information leadership, consistent with findings by Barclay and Warner (1993) and Aoyagi and Ito (2021), which suggest that informed traders follow volume (noise).

We segment the sample into quintiles based on the share of volume, with quintile 1 representing the lowest volume share and quintile 5 the highest. Quintile 1 exhibits a mean Information Leadership Share (ILS) of 0.6%, while quintile 5 records a mean ILS of 13.1%, consistent with Barclay and Warner (1993). We assess whether volume share impacts the determinants of AMM information leadership following Equation 3.14 to each quintile.

Table 3.7 provides OLS coefficient estimates for each quintile sample. Quintiles 4 and 5, representing higher volume shares, exhibit the strongest explanatory power for the determinants of price discovery. This shows that volume share is a critical factor in price leadership. However, while higher volume correlates with greater explanatory power, AMM price discovery remains relatively low on average, even in high-volume settings, suggesting other structural inefficiencies may limit AMM information leadership.

Hypothesis 5 suggests that the faster block times on Layer 2 blockchains would not increase the informativeness of their price. To test this hypothesis we estimate the following regression model:

$$PDS_{i,t} = \alpha_{i,t} + Blockchain_{i,t} + Gas_t + Controls_{i,t} + \epsilon_{i,t}, \tag{3.15}$$

where PDS is the AMM price discovery share (ILS) and Blockchain is a dummy variable which is equal to 1 if the AMM pool is on the associated blockchain: Polygon, Optimism, and Arbitrum. We consider the explicit cost of inclusion on-chain by measuring the value of gas on the respective blockchain. AMMs on Ethereum are the baseline. We use the same controls from Equation 3.14. All continuous variables are in natural log form.

Table 3.8 presents the coefficient estimates from OLS regressions based on Equation 3.15. The results reveal notable differences in AMM price leadership across Layer 2 blockchains. AMMs on Polygon exhibit significantly lower price discovery shares, despite

TABLE 3.7: Regressions to assess the determinants of automated market maker price discovery by share of volume

This table reports coefficient estimates from OLS regressions using pair-month observations to examine the determinants of the AMM price discovery shares by share of volume. The dependent variable is the Yan-Zivot-Putnins information leadership share ILS. Columns 1 through 5 represent a subsample of which is cut at each quintile, (Q1) through (Q5) respectively. (Q1) is the group with the lowest share of volume while (Q5) is the group with the highest share of volume. $QuotedSpread$ is the fee in the AMM. $EffectiveSpread$ is the volume weighted effective spread in the AMM winsorized at 99% before taking the weighted average to control for large outliers. $Depth_{[1.5\%]}$ is the average value of orders in thousands of dollars between 0% and 1.5% of the mid-quote price for the pair-month on the AMM. $DepthVariance_{[1.5\%]}$ is the standard deviation divided by the mean of the depth between 0% and 1.5% of the mid-quote price for the pair-month on the AMM. $Volume$ is the value of trade volume in dollars on the AMM. $RealizedVolatility$ is the average of squared hourly mid-quote returns on Binance. Gas is the average price of gas on the blockchain in dollars in a month. All explanatory variables are in natural log form. The sample comprises 1709 observations from 59 trading pairs during the period June 2020 to March 2023. Standard errors are clustered by pair month. t-statistics are reported in parentheses. Significance at the 10%, 5%, and 1% levels is indicated by *, **, and ***, respectively.

	(1)	(2)	(3)	(4)	(5)
Dependent variable: ILS	(*Q*1)	(*Q*2)	(*Q*3)	(*Q*4)	(*Q*5)
$QuotedSpread$	−0.020	−0.001*	−0.052**	−0.126***	−0.120***
	(−1.37)	(−1.90)	(−2.28)	(−4.68)	(−5.46)
$EffectiveSpread$	0.004	0.000	−0.007***	0.012	0.041***
	(0.72)	(−1.25)	(−2.60)	(1.21)	(2.85)
$Depth_{[1.5\%]}$	0.000	0.000	0.003***	−0.003	−0.009**
	(0.25)	(0.41)	(2.71)	(−1.13)	(−2.34)
$DepthVariance_{[1.5\%]}$	0.005	0.000	−0.009**	−0.010	−0.069***
	(1.00)	(−0.03)	(−2.49)	(−0.99)	(−4.34)
$Volume$	−0.001	0.001	−0.011**	0.019**	0.045***
	(−0.55)	(1.14)	(−2.31)	(2.06)	(3.65)
$RealizedVolatility$	0.000	0.003*	0.007**	0.014	0.022
	(0.07)	(1.76)	(2.22)	(1.45)	(1.17)
Gas	0.000	0.000	0.000	0.002	−0.001
	(0.44)	(−1.07)	(0.38)	(1.30)	(−0.23)
Intercept	−0.065	0.014	−0.111	−0.775***	−1.004***
	(−1.48)	(0.94)	(−1.22)	(−3.44)	(−3.83)
Observations	342	342	341	341	343
Adjusted R^2	4.6%	2.6%	17.4%	30.0%	30.0%

Significance: * $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$

the chain's lower explicit costs compared to Ethereum. In contrast, AMMs on Optimism and Arbitrum, with centralized transaction sequencing, are not significantly different to AMMs on Ethereum. The use of centralized sequencers on these chains might introduce potential fairness concerns. However, their incentive to maintain a first-come, first-served transaction ordering may help with informational efficiency. Overall, we do not find support for our hypothesis that more batching increases price informativeness.

One reason for this may be that informed traders face opportunity costs when holding inventory on a CEX or a specific blockchain, as capital tied up in one venue may limit their ability to act on arbitrage or impound information in prices elsewhere. These costs play a critical role in venue selection, as traders must weigh the benefits of holding assets on a CEX, where information is aggregated and arbitraged at the end of each block, against the flexibility of holding assets on a blockchain or Layer 2.

TABLE 3.8: Regressions to assess blockchain effects on automated market maker price discovery

This table reports coefficient estimates from OLS regressions using pair-month observations to examine the determinants of the AMM price discovery shares by blockchain. The dependent variable is the Yan-Zivot-Putnins information leadership share ILS. *Polygon*, *Optimism*, and *Arbitrum* are dummy variables equal to 1 if the AMM pool trades on that exchange. *Gas* is the average price of gas on the blockchain in dollars in a month. We control for the *QuotedSpread*, $EffectiveSpread$, $Depth_{[1.5\%]}$, $DepthVariance_{[1.5\%]}$, $Volume$, and *RealizedVolatility*. All explanatory variables are in natural log form. The sample comprises 1709 observations from 59 trading pairs during the period June 2020 to March 2023. Standard errors are clustered by pair month. t-statistics are reported in parentheses. Significance at the 10%, 5%, and 1% levels is indicated by *, **, and ***, respectively.

	(1)	(2)	(3)	(4)
Dependent variable:	ILS	ILS	ILS	ILS
Polygon	−0.070***			−0.083***
	(−5.21)			(−5.13)
Optimism		−0.005		−0.059***
		(−0.42)		(−3.09)
Arbitrum			−0.029	−0.022
			(−1.38)	(−1.08)
Gas	−0.003***	0.000	0.001	−0.003***
	(−3.66)	(−0.04)	(1.26)	(−2.93)
Intercept	−0.603***	−0.561***	−0.558***	−0.610***
	(−7.32)	(−7.12)	(−7.14)	(−7.36)
Controls	Y	Y	Y	Y
Observations	1,709	1,709	1,709	1,709
Adjusted R^2	20.2%	19.3%	19.4%	20.4%

Significance: * $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$

When controlling for blockchain differences, we find that gas costs are correlated with AMM price leadership. This result aligns with studies suggesting that such findings may be unique to the Ethereum blockchain and not generalizable to AMMs across other blockchains. Notably, this contrasts with Capponi, Jia, and Yu (2023), who argue that informed traders bid for priority inclusion in blockchains, and Barbon and Ranaldo (2022), who attribute AMM inefficiencies to high explicit costs. These results emphasize the importance of considering blockchain-specific factors when analyzing the role of explicit costs on AMM market quality.

3.7 Conclusion

This chapter investigates whether AMMs can provide efficient price discovery. Through a variance decomposition, we find that, for assets only traded on AMMs, there is no significant change in the level of information within prices compared with assets traded on both AMMs and LOBs. This finding highlights that AMMs can provide efficient price discovery. Our results show that LPs in AMMs respond to heightened risks of adverse selection by withdrawing their liquidity.

We also find that AMM prices often lag the price in a LOB market. However, in low-fee, high-liquidity pools where trading costs are low, AMMs can lead price discovery. This emphasizes the role of trading costs in influencing AMM price leadership.

AMMs promise to be a viable alternative for efficient trading, particularly in foreign exchange markets (Foley, O'Neill, & Putniņš, 2023; Malinova & Park, 2023). Given sufficient volume to compensate LPs for any adverse selection costs, AMMs can deliver cost savings and improved gains from trade, making them a promising market structure for tokenized assets.

Appendix 3.1 Estimation of variance decomposition

To estimate the level of information in a price series, we follow the variance decomposition method of Brogaard et al. (2022). Two key inputs are needed to construct the components of variance which we can get by estimating a reduced form structural Vector Autoregression (VAR) model: the variance of the innovations in each variable, $\sigma^2_{\epsilon_{r_m}}$, $\sigma^2_{\epsilon_x}$, $\sigma^2_{\epsilon_r}$ and the long-run cumulative responses to these shocks θ_{r_m}, θ_x, θ_r.

$$\begin{aligned} r_{m,t} &= a_0^* + \sum_{l=1}^{12} a_{1,l}^* r_{m,t-l} + \sum_{l=1}^{12} a_{2,l}^* x_{t-l} + \sum_{l=1}^{12} a_{3,l}^* r_{t-l} + e_{r_m,t} \\ x_t &= b_0^* + \sum_{l=1}^{12} b_{1,l}^* r_{m,t-l} + \sum_{l=1}^{12} b_{2,l}^* x_{t-l} + \sum_{l=1}^{12} b_{3,l}^* r_{t-l} + e_{x,t} \\ r_t &= c_0^* + \sum_{l=1}^{12} c_{1,l}^* r_{m,t-l} + \sum_{l=1}^{12} c_{2,l}^* x_{t-l} + \sum_{l=1}^{12} c_{3,l}^* r_{t-l} + e_{r,t} \end{aligned} \tag{A3.1}$$

The reduced form residuals can be written as linear models of the structural-model residuals:

$$\begin{aligned} e_{r_m,t} &= \epsilon_{r_m,t} \\ e_{x,t} &= \epsilon_{x,t} + b_{1,0}\epsilon_{r_m,t} = b_{1,0}e_{r_m,t} + \epsilon_{x,t} \\ e_{r,t} &= \epsilon_{r,t} + (c_{1,0} + c_{2,0}b_{1,0})\epsilon_{r_m,t} + c_{2,0}\epsilon_{x,t} = c_{1,0}e_{r_m,t} + c_{2,0}e_{x,t} + \epsilon_{r,t} \end{aligned} \tag{A3.2}$$

Although the structural-model residuals are uncorrelated contemporaneously by design, Equation A3.2 shows that the reduced-form residuals exhibit contemporaneous correlation. This correlation allows for inference about the structural-model residuals. Specifically, we estimate $b_{1,0}$ by regressing the reduced-form innovation $e_{x,t}$ on $e_{r_m,t}$ following the second equation in Equation A3.2. Similarly, we estimate $c_{1,0}$ and $c_{2,0}$ by regressing the reduced-form innovation $e_{r,t}$ on $e_{r_m,t}$ and $e_{x,t}$ as indicated in the third equation in Equation A3.2. Using the estimated parameters $b_{1,0}$, $c_{1,0}$ and $c_{2,0}$, along with the estimated variances of the reduced-form residuals $\sigma^2_{e_{r_m}}$, $\sigma^2_{e_x}$, and $\sigma^2_{e_r}$ we derive estimates for the variances of the structural model shocks by rearranging the variance expression of Equation A3.2.

$$\begin{aligned}
\sigma_{\epsilon_{rm}} &= \sigma_{e_{rm}} \\
\sigma_{\epsilon_x} &= \sigma_{e_x} - b_{1,0}^2\sigma_{e_{rm}}^2 \\
\sigma_{\epsilon_r} &= \sigma_{e_r} - (c_{1,0}^2 + 2c_{1,0}c_{2,0}b_{1,0})\sigma_{e_{rm}}^2 - c_{2,0}^2\sigma_{e_x}^2
\end{aligned} \tag{A3.3}$$

To estimate the long-run cumulative impulse response functions of the structural model, we compute the equivalent reduced-form shocks and feed them through the reduced-form model. Specifically:

1. A structural shock to market returns $[\epsilon_{r_m,t}, \epsilon_{x,t}\epsilon_{r,t}]' = [1, 0, 0]'$ has a reduced-form equivalent $[e_{r_m,t}, e_{x,t}e_{r,t}]' = [1, b_{1,0}, (c_{1,0} + c_{2,0}b_{1,0})]'$
2. A structural shock to market returns $[\epsilon_{r_m,t}, \epsilon_{x,t}\epsilon_{r,t}]' = [0, 1, 0]'$ has a reduced-form equivalent $[e_{r_m,t}, e_{x,t}e_{r,t}]' = [0, 1, c_{2,0}]'$
3. A structural shock to market returns $[\epsilon_{r_m,t}, \epsilon_{x,t}\epsilon_{r,t}]' = [0, 0, 1]'$ has a reduced-form equivalent $[e_{r_m,t}, e_{x,t}e_{r,t}]' = [0, 0, 1]'$

The cumulative return response to each of these shocks, evaluated at $t = 36$ (the point where the responses stabilize), provides estimates for θ_{r_m}, θ_x, and θ_r respectively.

The structural VAR innovations $\{\epsilon_{r_m,t}, \epsilon_{x,t}, \epsilon_{r,t}\}$ are the shocks that contain information and noise. The information in each shock is the estimated long-run effect on the price, which is the permanent component in the Beveridge and Nelson (1981) decomposition. Therefore, the market-wide information at time t is $\theta_{r_m}\epsilon_{r_m,t}$ where $\epsilon_{r_m,t}$ is the innovation in the structural VAR and θ_{r_m} is the long-run effect of a unit shock, inferred from the cumulative impulse response function of returns to a unit shock. Similarly, the private firm-specific information revealed through trading and the public firm-specific information revealed through other sources are estimated as $\theta_x\epsilon_{x,t}$ and $\theta_r\epsilon_{r,t}$, where $\epsilon_{x,t}$ and $\epsilon_{r,t}$ are the innovations in the structural VAR and θ_x and θ_r are the long-run cumulative impulse responses of returns to a unit shock.

The sum of the estimated information components gives the innovation in the efficient price, $w_t = \theta_{r_m}\varepsilon_{r_m,t} + \theta_x\varepsilon_{x,t} + \theta_r\varepsilon_{r,t}$. The change in the pricing error (noise return), is the realized return that is not attributable to information or discount rate (drift): $s_t = r_t - \mu - w_t = r_t - a0 - \theta_{r_m}\varepsilon_{r_m,t} - \theta_x\varepsilon_{x,t} - \theta_r\varepsilon_{r,t}$, and the noise variance, σ_s^2, is estimated as the variance of the time series of Δs_t

Taking the variance of the innovations in the efficient price we get $\sigma_w^2 = \theta_{r_m}^2\sigma_{\epsilon_{r_m}}^2 + \theta_x^2\sigma_{\epsilon_x}^2 + \theta_r^2\sigma_{\epsilon_r}^2$. Recall, the errors in the structural model are contemporaneously uncorrelated by construction and therefore the covariance terms are all zero. The contribution to the variation in the efficient price from each of the information components is $\theta_{r_m}^2\sigma_{\epsilon_{r_m}}^2$ (market-wide information), $\theta_x^2\sigma_{\epsilon_x}^2$ (private firm-specific information), and $\theta_r^2\sigma_{\epsilon_r}^2$ (public firm-specific information). The estimated components of variance are therefore

$$\begin{aligned} MarketInfo &= \theta_{r_m}^2\sigma_{\varepsilon_{r_m}}^2 \\ PrivateInfo &= \theta_x^2\sigma_{\varepsilon_x}^2 \\ PrivateInfo &= \theta_r^2\sigma_{\varepsilon_r}^2 \\ Noise &= \sigma_s^2 \end{aligned} \tag{A3.4}$$

Normalizing these variance components to sum to 100% gives variance shares:

$$\begin{aligned} MarketInfoShare &= \theta_{r_m}^2\sigma_{\varepsilon_{r_m}}^2/(\sigma_w^2 + \sigma_s^2) \\ PrivateInfoShare &= \theta_x^2\sigma_{\varepsilon_x}^2/(\sigma_w^2 + \sigma_s^2) \\ PrivateInfoShare &= \theta_r^2\sigma_{\varepsilon_r}^2/(\sigma_w^2 + \sigma_s^2) \\ NoiseShare &= \sigma_s^2/(\sigma_w^2 + \sigma_s^2) \end{aligned} \tag{A3.5}$$

$MarketInfo$, $PrivateInfo$, $PublicInfo$, and $Noise$ are the variance contributions of market-wide information, trading on private firm-specific information, firm-specific information other than that revealed through trading, and noise, respectively. $MarktInfoShare$, $PrivateInfoShare$, and $PublicInfoShare$, are the corresponding shares of variance from those various sources of stock-price movements. Meanwhile, $NoiseShare$ reflects the relative importance of pricing errors due to over- or under-reaction to information, illiquidity, price pressures, or other microstructure frictions.

Appendix 3.2 Variance decomposition summary statistics

TABLE A3.1: Summary statistics for shares of variance

This table presents the summary statistics for the variance decomposition from a sample of pairs traded on the LOB based on pair month observations for the period between June 2020 and March 2023. $NoiseShare$, $MarketInfoShare$, $PrivateInfoShare$, and $PublicInfoShare$, are the shares of variance explained by noise, market information, private information and public information respectively.

	Mean	Standard Deviation	p10	Median	p90
$MarketInfoShare$	11.80	19.35	0.12	4.40	26.12
$PrivateInfoShare$	27.14	16.04	5.98	25.75	48.28
$PublicInfoShare$	42.62	18.27	19.26	42.59	65.97
$NoiseInfoShare$	18.44	11.93	7.71	14.73	37.08

Appendix 3.3 Estimation of price discovery shares

We estimate price discovery shares following the method outlined in Patel et al. (2020). For each pair-month, we estimate a reduced form VECM of the log price series ($p_{1,t}$ to $p_{n,t}$) with 300 lags (prices are sampled based on the Ethereum block time where trading is continuous in the AMM and the LOB).

$$\Delta p_t = aZ_{t-1} + \sum_{i=1}^{300} b_i \Delta p_{t-i} + \epsilon_t \tag{A3.6}$$

where Δp_t is the $n \times 1$ mid-quote return vector, α is the $n \times (n-1)$ matrix of error correction coefficients, Z_{t-1} is the $n \times 1$ co-integrating vector, b_i is the $n \times n$ coefficient matrix for lag i and ϵ_t is the $n \times 1$ vector of residuals.

From the reduced form VECM estimates in Equation A3.6 we derive the corresponding infinite lag Vector Moving Average (VMA) representation in structural form assuming recursive contemporaneous causality running from the first through to the last price series.

$$\begin{aligned}
\Delta p_{1,t} &= \sum_{l=0}^{\infty} A_{1,l}\varepsilon_{1,t-1} + \sum_{l=0}^{\infty} A_{2,l}\varepsilon_{2,t-1} + \cdots + \sum_{l=0}^{\infty} A_{n,l}\varepsilon_{n,t-1} \\
\Delta p_{2,t} &= \sum_{l=0}^{\infty} B_{1,l}\varepsilon_{1,t-1} + \sum_{l=0}^{\infty} B_{2,l}\varepsilon_{2,t-1} + \cdots + \sum_{l=0}^{\infty} B_{n,l}\varepsilon_{n,t-1} \\
&\vdots \qquad\qquad\qquad \vdots \qquad\qquad\qquad \vdots \\
\Delta p_{n,t} &= \sum_{l=0}^{\infty} N_{1,l}\varepsilon_{1,t-1} + \sum_{l=0}^{\infty} N_{2,l}\varepsilon_{2,t-1} + \cdots + \sum_{l=0}^{\infty} N_{n,l}\varepsilon_{n,t-1}
\end{aligned} \tag{A3.7}$$

We obtain the structural VMA coefficients by computing the orthogonalized impulse response functions and the (contemporaneously uncorrelated) structural VMA errors ($\varepsilon_{1,t}$ to $\varepsilon_{n,t}$) by mapping their relation to the reduced form errors.

The permanent price impacts of shocks to the n price series are obtained from the structural VMA. For example, a unit shock to the first price ($\varepsilon_{1,t} = 1$) has a permanent effect on all of the prices equal to $\theta_{\varepsilon 1} = \sum_{l=0}^{\infty} A_{1,l}$. The permanent price impact is the same for all prices $\theta_{\epsilon 1} = \sum_{l=0}^{\infty} A_{1,l} = \sum_{l=0}^{\infty} B_{1,l} = \sum_{l=0}^{\infty} N_{1,l}$ because permanent impacts are innovations in the efficient value and the efficient value is common to all prices as they refer to the same underlying asset (Hasbrouck, 1995).

Following Hasbrouck (1995) the temporary-permanent decomposition, innovations in the permanent component (the efficient price, m_t) are given by:

$$\Delta m_t = \theta_{\varepsilon 1}\varepsilon_{1,t} + \theta_{\varepsilon 2}\varepsilon_{2,t} + \cdots + \theta_{\varepsilon n}\varepsilon_{n,t} \tag{A3.8}$$

The variance of the innovations in the efficient price is therefore:

$$\begin{aligned} Var(\Delta m_t) &= Var(\theta_{\varepsilon 1}\varepsilon_{1,t} + \theta_{\varepsilon 2}\varepsilon_{2,t} + \cdots + \theta_{\varepsilon n}\varepsilon_{n,t}) \\ &= \theta_{\varepsilon 1}^2 Var(\varepsilon_{1,t}) + \theta_{\varepsilon 2}^2 Var(\varepsilon_{2,t}) + \cdots + \theta_{\varepsilon n}^2 Var(\varepsilon_{n,t}) \end{aligned}$$

because the structural VMA errors are uncorrelated by construction (contemporaneous correlation in the reduced form errors is absorbed in the recursive contemporaneous effects that are part of the structural VMA). Hasbrouck (1995) information shares (IS) are obtained as each price's contribution to the variance of the efficient price innovations

$$IS_n = \frac{\theta_{\varepsilon n}^2 Var(\varepsilon_{n,t})}{Var(\Delta m_t)} \tag{A3.9}$$

The Gonzalo and Granger (1995) component shares (CS) are obtained by normalizing the permanent price impacts of each price series in the reduced form model. The permanent price impacts of the reduced form model are based on the simple impulse response functions from the reduced form VECM instead of orthogonalized impulse response functions. We ensure all permanent price impacts are non-negative so that the CS take the range $[0, 1]$.

$$CS_n = \frac{\theta_{\epsilon n}}{\sum_{i=1}^{n} \theta_{\epsilon i}} \tag{A3.10}$$

Finally, we calculate the information leadership share (ILS). In the two-price case, market's propensity to reflect new information (how much market i's price responds to an innovation in the efficient price) can be obtained from the ratio $\beta_i = \frac{IS_i}{CS_i}$, which when normalized gives the information leadership share

$$ILS_n = \frac{\beta_n^2}{\sum_{i=1}^{n} \beta_i^2} \tag{A3.11}$$

The IS estimated in Hasbrouck (1995) are not unique and depend on the ordering of the prices in the estimation procedure. This is due to the recursive contemporaneous causality assumed to run from the first through to the last price series (this assumption is implicit in procedures that use Cholesky factorization of the covariance matrix of reduced form errors, and explicit in the structural VMA that we present above). We apply the standard approach used in the literature and estimate IS (and subsequently ILS) for each price under all possible orderings of the prices. This gives a range of IS values for each price, with the minimum and maximum values providing the lower and upper bounds on IS. We take the median IS and ILS from the range of IS and ILS values to get a single estimate for each price.

Appendix 3.4 Correlation tables

TABLE A3.2: Correlation of variance decomposition variables

This table presents the correlation matrix of variables used in regressions on the variance decomposition results. $NoiseShare$, $MarketInfoShare$, $PrivateInfoShare$, and $PublicInfoShare$, are the shares of variance explained by noise, market information, private information and public information respectively. $EffectiveSpread$ is the volume weighted effective spread in the AMM winsorized at 99% before taking the weighted average to control for large outliers. $Depth_{[1.5\%]}$ is the average value of orders in thousands of dollars between 0% and 1.5% of the mid-quote price for the pair-month on the AMM. $DepthVariance_{[1.5\%]}$ is the standard deviation divided by the mean of the depth between 0% and 1.5% of the mid-quote price for the pair-month on the AMM. $RealizedVolatility$ is the average of squared hourly mid-quote returns on Binance. $Volume$ is the value of trade volume in dollars on the AMM. Gas is the average price of gas on the blockchain in dollars in a month. $AMMOnly$ is a dummy variables that equal to 1 if the pair only trades on a AMM, otherwise 0. $CLCPMM$ is a dummy variable for that is equal to one if the market uses concentrated liquidity from Uniswap v3.

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
$NoiseShare$ (1)	1.00											
$MarketInfoShare$ (2)	-0.16	1.00										
$PrivateInfoShare$ (3)	-0.78	-0.40	1.00									
$PublicInfoShare$ (4)	-0.25	-0.03	-0.07	1.00								
$EffectiveSpread$ (5)	0.28	-0.02	-0.28	0.08	1.00							
$Depth_{[1.5\%]}$ (6)	-0.10	0.22	-0.03	0.00	-0.03	1.00						
$DepthVariance_{[1.5\%]}$ (7)	-0.05	0.30	-0.22	0.28	0.20	0.15	1.00					
$RealizedVolatility$ (8)	0.19	-0.05	-0.16	0.02	0.62	-0.02	0.17	1.00				
$Volume$ (9)	-0.19	0.27	0.02	0.01	-0.05	0.55	0.12	-0.01	1.00			
Gas (10)	0.12	-0.07	-0.08	0.01	0.04	0.09	-0.07	-0.02	0.17	1.00		
AMM Only (11)	-0.00	-0.07	0.06	-0.04	-0.11	-0.01	-0.12	-0.03	-0.00	-0.00	1.00	
ConcLiquidity (12)	-0.06	0.26	-0.10	0.05	-0.06	0.19	0.48	-0.05	0.08	-0.12	0.00	1.00

TABLE A3.3: Correlation of price discovery shares variables

This table presents the correlation matrix of variables used in regressions on the price discovery shares results. ILS is the Yan-Zivot-Putnins information leadership share. IS is the the Hasbrouck information share. $QuotedSpread$ is the fee in the AMM. $EffectiveSpread$ is the volume weighted effective spread in the AMM winsorized at 99% before taking the weighted average to control for large outliers. $Depth_{[1.5\%]}$ is the average value of orders in thousands of dollars between 0% and 1.5% of the mid-quote price for the pair-month on the AMM. $DepthVariance_{[1.5\%]}$ is the standard deviation divided by the mean of the depth between 0% and 1.5% of the mid-quote price for the pair-month on the AMM. $RealizedVolatility$ is the average of squared hourly mid-quote returns on Binance. $Volume$ is the value of trade volume in dollars on the AMM. Gas is the average price of gas on the blockchain in dollars in a month. $LOBQuotedSpread$ is the average half quoted spread on Binance. $LOBEffectiveSpread$ is the volume weighted effective spread on Binance. $LOBDepth_{[1.5\%]}$ is the average value of orders in thousands of dollars between 0% and 1.5% of the mid-quote price for the pair-month on Binance. $LOBDepthVariance_{[1.5\%]}$ is the standard deviation divided by the mean of the depth between 0% and 1.5% of the mid-quote price for the pair-month on Binance. $LOBVolume$ is the value of trade volume in dollars on Binance.

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)	(13)	(14)
ILS (1)	1.00													
IS (2)	0.96	1.00												
$QuotedSpread$ (3)	-0.24	-0.27	1.00											
$EffectiveSpread$ (4)	-0.01	-0.04	0.04	1.00										
$Depth_{[1.5\%]}$ (5)	0.08	0.06	-0.09	-0.06	1.00									
$DepthVariance_{[1.5\%]}$ (6)	-0.03	-0.08	0.14	0.01	0.06	1.00								
$RealizedVolatility$ (7)	0.00	0.06	0.01	-0.05	-0.12	-0.06	1.00							
$Volume$ (8)	0.11	0.15	-0.20	-0.06	0.09	0.02	-0.05	1.00						
Gas (9)	-0.07	-0.08	-0.00	-0.06	0.01	-0.02	0.21	0.07	1.00					
$LOBQuotedSpread$ (10)	0.06	0.13	-0.02	-0.06	-0.15	-0.06	0.57	-0.08	0.10	1.00				
$LOBEffectiveSpread$ (11)	0.03	0.10	-0.02	-0.03	-0.10	-0.11	0.23	-0.13	-0.08	0.37	1.00			
$LOBDepth_{[1.5\%]}$ (12)	-0.07	-0.09	-0.03	-0.02	0.12	0.10	-0.12	0.11	0.03	-0.22	-0.22	1.00		
$LOBDepthVariance_{[1.5\%]}$ (13)	0.03	0.08	-0.01	-0.04	0.05	0.02	0.17	-0.02	0.02	0.31	0.15	-0.12	1.00	
$LOBVolume$ (14)	-0.06	-0.07	-0.00	-0.01	0.08	0.11	-0.04	0.08	0.04	-0.13	-0.16	0.87	-0.07	1.00

Appendix 3.5 Key findings based on information shares

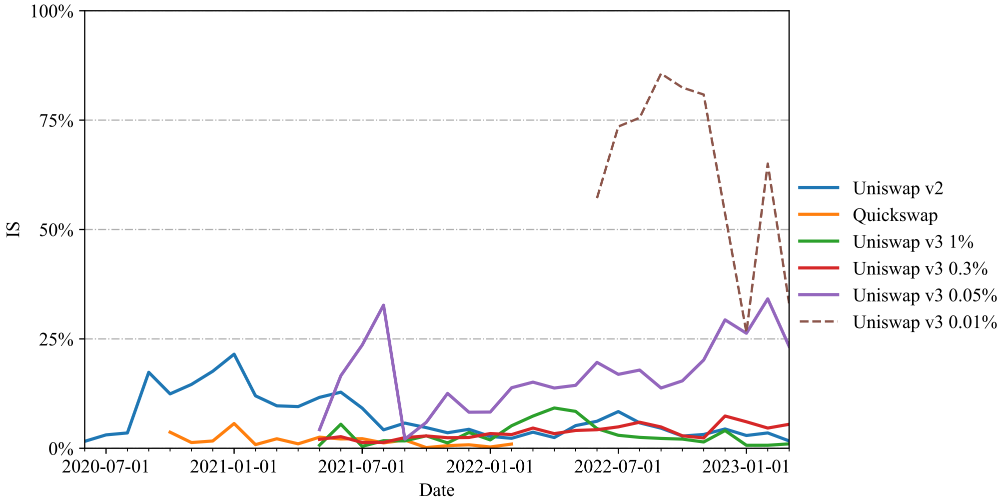

FIGURE A3.1: Automated market maker information shares

This figure plots the Hasbrouck information share (IS) for AMM pools from 58 trading pairs during the period June 2020 to March 2023. The plotted ILS is the mean price discovery share for each different Uniswap and Quickswap market (Uniswap v2, Quickswap, Uniswap v3 1%, Uniswap v3 0.3%, Uniswap v3 0.05%, and Uniswap v3 0.01%).

TABLE A3.4: Regressions to assess the determinants of automated market maker information shares

This table reports coefficient estimates from OLS regressions using pair-month observations to examine the determinants of the AMM price discovery shares. The dependent variable is the the Hasbrouck information share (IS). $QuotedSpread$ is the fee in the AMM. $EffectiveSpread$ is the volume weighted effective spread in the AMM winsorized at 99% before taking the weighted average to control for large outliers. $Depth_{[1.5\%]}$ is the average value of orders in thousands of dollars between 0% and 1.5% of the mid-quote price for the pair-month on the AMM. $DepthVariance_{[1.5\%]}$ is the standard deviation divided by the mean of the depth between 0% and 1.5% of the mid-quote price for the pair-month on the AMM. $Volume$ is the value of trade volume in dollars on the AMM. $RealizedVolatility$ is the average of squared hourly mid-quote returns on Binance. Gas is the average price of gas on the blockchain in dollars in a month. All explanatory variables are in natural log form. The sample comprises 1709 observations from 59 trading pairs during the period June 2020 to March 2023. Standard errors are clustered by pair month. t-statistics are reported in parentheses. Significance at the 10%, 5%, and 1% levels is indicated by *, **, and ***, respectively.

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Dependent variable:	IS	IS	IS	IS	IS	IS	IS	IS
$QuotedSpread$	−0.074***			−0.069***	−0.058***			−0.058***
	(−8.10)			(−7.53)	(−7.11)			(−7.19)
$EffectiveSpread$		0.000		0.003		0.003		0.005
		(−0.06)		(0.75)		(0.72)		(1.14)
$Depth_{[1.5\%]}$			−0.009***	−0.006***			−0.005**	−0.001
			(−4.94)	(−4.17)			(−2.55)	(−0.38)
$DepthVariance_{[1.5\%]}$			−0.028***	−0.022***			−0.002	0.002
			(−7.44)	(−6.69)			(−0.82)	(0.61)
$Volume$	0.004**	0.015***	0.027***	0.015***	0.014***	0.026***	0.030***	0.016***
	(2.00)	(5.28)	(9.12)	(4.77)	(6.30)	(8.81)	(9.41)	(5.42)
$RealizedVolatility$	0.024***	0.015***	0.003	0.013***				
	(7.14)	(3.87)	(0.72)	(4.02)				
Gas	0.000	0.000	0.000	0.000				
	(0.32)	(0.25)	(0.42)	(0.41)				
Intercept	−0.212***	−0.038	−0.292***	−0.408***	−0.456***	−0.340***	−0.354***	−0.454***
	(−3.96)	(−0.82)	(−5.02)	(−6.16)	(−6.41)	(−5.19)	(−5.22)	(−6.30)
Pair Effects	N	N	N	N	Y	Y	Y	Y
Month Effects	N	N	N	N	Y	Y	Y	Y
Observations	1,709	1,709	1,709	1,709	1,709	1,709	1,709	1,709
Adjusted R^2	18.1%	5.8%	11.8%	21.9%	42.3%	37.1%	37.5%	42.3%

Significance: * $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$

TABLE A3.5: Regressions to assess the determinants of automated market maker information shares from limit order book variables

This table reports coefficient estimates from OLS regressions using pair-month observations to examine the limit order book determinants of the AMM price discovery shares. The dependent variable is the the Hasbrouck information share (IS). $LOBQuotedSpread$ is the average half quoted spread on Binance. $LOBEffectiveSpread$ is the volume weighted effective spread on Binance. $LOBDepth_{[1.5\%]}$ is the average value of orders in thousands of dollars between 0% and 1.5% of the mid-quote price for the pair-month on Binance. $LOBDepthVariance_{[1.5\%]}$ is the standard deviation divided by the mean of the depth between 0% and 1.5% of the mid-quote price for the pair-month on Binance. $LOBVolume$ is the value of trade volume in dollars on Binance. $RealizedVolatility$ is the average of squared hourly mid-quote returns on Binance. All explanatory variables are in natural log form. The sample comprises 1709 observations from 59 trading pairs during the period June 2020 to March 2023. Standard errors are clustered by pair month. t-statistics are reported in parentheses. Significance at the 10%, 5%, and 1% levels is indicated by *, **, and ***, respectively.

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Dependent variable:	IS	IS	IS	IS	IS	IS	IS	IS
$LOBQuotedSpread$	0.015***			0.013***	0.008***			0.002
	(5.72)			(3.57)	(2.59)			(0.63)
$LOBEffectiveSpread$		0.011		−0.006		0.026***		0.020***
		(1.29)		(−0.93)		(3.43)		(2.71)
$LOBDepth_{[1.5\%]}$			−0.023***	−0.016***			−0.035***	−0.030***
			(−2.68)	(−2.85)			(−4.55)	(−3.82)
$LOBDepthVariance_{[1.5\%]}$			0.008	0.000			−0.007	−0.009
			(0.92)	(−0.03)			(−0.90)	(−1.09)
$LOBVolume$	−0.002	−0.002	0.011**	0.007**	−0.016***	−0.013***	−0.003	−0.002
	(−1.22)	(−0.77)	(1.99)	(2.19)	(−3.10)	(−2.62)	(−0.47)	(−0.28)
$RealizedVolatility$		0.004	−0.004					
		(0.63)	(−0.65)					
Intercept	0.205***	0.215***	0.117**	0.182***	0.385***	0.456***	0.491***	0.563***
	(10.94)	(5.46)	(2.09)	(9.30)	(3.75)	(4.11)	(4.50)	(4.98)
Pair Effects	N	N	N	N	Y	Y	Y	Y
Month Effects	N	N	N	N	Y	Y	Y	Y
Observations	1,709	1,709	1,709	1,709	1,709	1,709	1,709	1,709
Adjusted R^2	2.3%	1.4%	2.0%	2.6%	30.1%	30.4%	30.6%	30.8%

Significance: * $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$

TABLE A3.6: Regressions to assess the relative determinants of automated market maker information shares

This table reports coefficient estimates from OLS regressions using pair-month observations to examine the determinants of the AMM price discovery shares. The dependent variable is the the Hasbrouck information share (IS). $QuotedSpreadRatio$ is a ratio between the fee in the AMM and the average half quoted spread for Binance. $EffectiveSpreadRatio$ is the ratio of the volume weighted effective spread between the AMM and Binance. $Depth_{[1.5\%]}Ratio$ is the ratio of the average value of orders in dollars between 0% and 1.5% of the mid-quote price for the AMM and Binance, respectively. $DepthVariance_{[1.5\%]}Ratio$ is the ratio of the standard deviation divided by the average of orders in dollars between 0% and 1.5% of the mid-quote price for the AMM and Binance, respectively. $VolumeRatio$ is the ratio of the value of trade volume in dollars between the AMM and Binance. Gas is the average price of gas on the blockchain in dollars in a month. All explanatory variables are in natural log form. The sample comprises 1709 observations from 59 trading pairs during the period June 2020 to March 2023. Standard errors are clustered by pair month. t-statistics are reported in parentheses. Significance at the 10%, 5%, and 1% levels is indicated by *, **, and ***, respectively.

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Dependent variable:	IS	IS	IS	IS	IS	IS	IS	IS
$QuotedSpreadRatio$	−0.032***			−0.033***	−0.031***			−0.033***
	(−8.25)			(−7.04)	(−7.13)			(−6.96)
$EffectiveSpreadRatio$		−0.003		0.012**		0.002		0.010*
		(−0.56)		(2.11)		(0.43)		(1.69)
$Depth_{[1.5\%]}Ratio$			−0.006***	−0.003*			−0.004**	−0.001
			(−3.24)	(−1.78)			(−2.22)	(−0.86)
$DepthVarience_{[1.5\%]}Ratio$			−0.016***	−0.007**			−0.001	0.002
			(−5.38)	(−2.37)			(−0.57)	(0.70)
$VolumeRatio$	0.010***	0.016***	0.020***	0.014***	0.020***	0.027***	0.030***	0.023***
	(7.59)	(7.24)	(9.82)	(6.13)	(9.03)	(8.60)	(9.63)	(7.30)
Intercept	0.113***	0.098***	0.087***	0.091***	0.101**	0.048	0.057	0.088*
	(18.11)	(9.29)	(18.61)	(9.62)	(2.08)	(0.97)	(1.20)	(1.75)
Pair Effects	N	N	N	N	Y	Y	Y	Y
Month Effects	N	N	N	N	Y	Y	Y	Y
Observations	1,709	1,709	1,709	1,709	1,709	1,709	1,709	1,709
Adjusted R^2	17.2%	9.7%	12.2%	18.2%	41.5%	38.2%	38.5%	41.7%

Significance: * $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$

TABLE A3.7: Regressions to assess the determinants of automated market maker information shares by share of volume

This table reports coefficient estimates from OLS regressions using pair-month observations to examine the determinants of the AMM price discovery shares by share of volume. The dependent variable is the the Hasbrouck information share (IS). Columns 1 through 5 represent a subsample of which is cut at each quintile, (Q1) through (Q5) respectively. (Q1) is the group with the lowest share of volume while (Q5) is the group with the highest share of volume. $QuotedSpread$ is the fee in the AMM. $EffectiveSpread$ is the volume weighted effective spread in the AMM winsorized at 99% before taking the weighted average to control for large outliers. $Depth_{[1.5\%]}$ is the average value of orders in thousands of dollars between 0% and 1.5% of the mid-quote price for the pair-month on the AMM. $DepthVariance_{[1.5\%]}$ is the standard deviation divided by the mean of the depth between 0% and 1.5% of the mid-quote price for the pair-month on the AMM. $Volume$ is the value of trade volume in dollars on the AMM. $RealizedVolatility$ is the average of squared hourly mid-quote returns on Binance. Gas is the average price of gas on the blockchain in dollars in a month. All explanatory variables are in natural log form. The sample comprises 1709 observations from 59 trading pairs during the period June 2020 to March 2023. Standard errors are clustered by pair month. t-statistics are reported in parentheses. Significance at the 10%, 5%, and 1% levels is indicated by *, **, and ***, respectively

	(1)	(2)	(3)	(4)	(5)
Dependent variable: IS	($Q1$)	($Q2$)	($Q3$)	($Q4$)	($Q5$)
$QuotedSpread$	−0.025**	−0.007**	−0.047***	−0.105***	−0.099***
	(−2.05)	(−2.14)	(−2.77)	(−4.67)	(−5.48)
$EffectiveSpread$	0.001	−0.006***	−0.012***	0.001	0.025*
	(0.15)	(−4.53)	(−4.30)	(0.10)	(1.96)
$Depth_{[1.5\%]}$	−0.001	0.001	0.005***	−0.001	−0.009***
	(−0.46)	(1.36)	(3.24)	(−0.35)	(−2.90)
$DepthVariance_{[1.5\%]}$	0.002	−0.004***	−0.016***	−0.017*	−0.065***
	(0.41)	(−2.60)	(−4.52)	(−1.89)	(−4.97)
$Volume$	0.002	0.001	−0.012***	0.007	0.033***
	(0.73)	(0.86)	(−2.95)	(0.93)	(3.22)
$RealizedVolatility$	0.006**	0.013***	0.020***	0.033***	0.032**
	(2.14)	(4.41)	(4.48)	(3.78)	(2.06)
Gas	0.001	0.000	0.000	0.003**	0.000
	(0.87)	(−0.34)	(0.13)	(2.06)	(0.02)
Intercept	−0.083*	0.047	0.050	−0.317*	−0.572***
	(−1.77)	(1.64)	(0.61)	(−1.67)	(−2.67)
Observations	342	342	341	341	343
Adjusted R^2	12.4%	12.3%	22.0%	31.1%	30.5%

Significance: * $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$

TABLE A3.8: Regressions to assess blockchain effects on automated market maker information shares

This table reports coefficient estimates from OLS regressions using pair-month observations to examine the determinants of the AMM price discovery shares by blockchain. The dependent variable is the the Hasbrouck information share (IS). *Polygon*, *Optimism*, and *Arbitrum* are dummy variables equal to 1 if the AMM pool trades on that exchange. *Gas* is the average price of gas on the blockchain in dollars in a month. We control for the $QuotedSpread$, $EffectiveSpread$, $Depth_{[1.5\%]}$, $DepthVariance_{[1.5\%]}$, $Volume$, and $RealizedVolatility$. All explanatory variables are in natural log form. The sample comprises 1709 observations from 59 trading pairs during the period June 2020 to March 2023. Standard errors are clustered by pair month. t-statistics are reported in parentheses. Significance at the 10%, 5%, and 1% levels is indicated by *, **, and ***, respectively.

	(1)	(2)	(3)	(4)
Dependent variable:	IS	IS	IS	IS
Polygon	−0.079***			−0.095***
	(−6.29)			(−6.27)
Optimism		−0.010		−0.071***
		(−0.78)		(−3.71)
Arbitrum			−0.040**	−0.032*
			(−2.25)	(−1.89)
Gas	−0.003***	0.000	0.002**	−0.003***
	(−3.91)	(0.27)	(2.03)	(−2.92)
Intercept	−0.455***	−0.408***	−0.403***	−0.463***
	(−6.57)	(−6.16)	(−6.16)	(−6.65)
Controls	Y	Y	Y	Y
Observations	1,709	1,709	1,709	1,709
Adjusted R^2	23.4%	21.8%	22.1%	23.8%

Significance: * $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$

Chapter 4

Misconduct in Decentralized Exchanges

4.1 Introduction

Cryptocurrencies have long been associated with illicit activity, a perception reinforced by reports linking nearly $180 billion in total value to illicit addresses between 2020 and 2024 (Chainalysis, 2025). The pseudonymous nature of blockchains and the lack of regulation in cryptocurrency markets have made digital assets a preferred medium for darknet transactions and other unlawful activities (Foley, Karlsen, & Putniņš, 2019). Recently, we have seen new markets emerge on blockchains, Decentralized Exchanges (DEXs), which inherit the properties of blockchains presenting unique challenges for oversight and regulation of market abuse.

DEXs possess several features that create an environment conducive to market misconduct. Unlike Centralized Exchanges (CEXs), they typically operate without Know-Your-Customer (KYC) requirements, allowing participants to trade pseudonymously. These markets also function without intermediaries or direct regulatory oversight, and their permissionless nature lowers entry barriers for bad actors. Despite these vulnerabilities, the transparency of public blockchains provides the ability to directly measure and assess misconduct in these markets.

There is growing concern among regulators over misconduct in DEXs (and cryptocurrency markets more broadly), yet enforcing compliance remains a significant challenge.[1] Despite this, the actual extent and impact of market misconduct on DEXs remains poorly understood. This chapter aims to bridge this gap by systematically quantifying the prevalence of market misconduct in DEXs.

We develop an index that aggregates multiple indicators of market misconduct on DEXs, focusing on rug pulls, wash trading, sandwich attacks, and money laundering. This index can be interpreted as the proportion of volume related to market misconduct, creating a measure of market integrity. We construct the index using transaction-level data from major DEXs—Uniswap, Pancakeswap, Quickswap, and Sushiswap—on the Ethereum, Binance Smart Chain (BSC), and Polygon blockchains.

We identify $345 billion in trading volume on DEXs that is linked to market misconduct, on average 12% of all DEX trading volume. While our estimates are not directly comparable to those reported by Chainalysis—who identify $180 billion in illicit activity across all of crypto—the magnitude of misconduct we uncover is nearly twice as large.[2] This discrepancy underscores both the pervasiveness and scale of market abuse in DEXs.

A sandwich attack is where a block builder exploits a traders transaction by both front and back running their trade on a DEX. We directly identify sandwich attacks by tracing specific patterns of trading activity: an attacker's buy (sell) is immediately followed by a victim's buy (sell), and then the attacker's offsetting sell (buy), all within a single block. Using this method, we detect $192 billion in trading volume associated with attacker trades. The vast majority of these sandwich attacks occur on Ethereum DEXs.

[1] See https://bit.ly/4d4Q1nf for European Union report on embedded supervision of Decentralized Finance (DeFi).

[2] Chainalysis focuses on on-chain transaction volumes associated with hacks, ransomware, darknet markets, money laundering, and other non-market forms of misconduct. In contrast, we study market misconduct on DEXs, for which there is very little overlap with the activity captured in those figures. While we also consider money laundering, our analysis is limited to the share of DEX trading volume originating from wallets linked to similar categories of misconduct.

Wash trading is where a trader buys and sells the same asset to create the illusion of volume. To identify wash trading, we analyze patterns of circular trading at a user level, focusing on repeated, non-economically motivated transactions between a single wallet. We detect \$73 billion in wash trading volume, predominantly in constant product markets. Activity varies across chains and tokens: on Ethereum, it is concentrated in smaller, less reputable tokens, while on Polygon it is concentrated in larger token pairs such as wrapped Polygon and wrapped Ethereum. Wash trading is also highly concentrated at the wallet level, with individual addresses generating billions in volume through either large paired trades or numerous small swaps per transaction. As a result, reported trading volume on DEXs substantially overstates genuine economic activity, and current regulatory efforts, including recent Department Of Justice (DOJ) actions, capture only a small fraction of this behavior.[3]

A rug pull is a scam where a bad actor launches a project and token, creates a market to trade it, attracts investors, and then abruptly withdraws funds and liquidity, leaving others holding worthless tokens. We identify rug pulls using a set of heuristics that capture common patterns: a price run-up followed by a reversal, profits accrued by the pool creator, and a halt in trading activity in the pool. Applying these criteria, we detect \$56 billion in total volume associated with DEX pools we identify as rug pulls. Rug pulls are most common on Uniswap v2 on Ethereum and Pancakeswap v2 on BSC, where we observe over 371 thousand individual rug pulls across the two platforms and approximately \$40 billion in trading volume. While our methodology captures substantial rug pull activity, our estimates are conservative relative to prior work and likely understate the overall prevalence and volume of such misconduct on DEXs (Cernera et al., 2023).

Money laundering involves concealing the origins of illegally obtained funds. We estimate an upper bound of money laundering on DEXs by tracking trading activity from wallet addresses that received funds from the Tornado Cash mixer, as well as wallet addresses linked to known exploits using data from the RektDB database. Addresses we

[3]See https://bit.ly/4jJV674 for DOJ press release on wash trading.

label as money launderers account for $35 billion in trading volume. Notably, about 70% of this volume flows through Uniswap v2 on Ethereum.

We find that both market design and blockchain-level factors—such as transaction costs and block times—meaningfully influence the prevalence of misconduct in DEXs. Rug pulls are sensitive to blockchain transaction costs (gas), with less rug pulls when fees are high. Moreover, Polygon DEXs consistently exhibit lower misconduct across all integrity measures. Market design also plays a role in the market integrity with constant product DEXs experiencing a higher proportion of wash trading, rug pulls, and money laundering while having a smaller proportion of sandwich attacks when compared to concentrated liquidity markets.

Despite the scale of misconduct, market integrity issues do not appear to deter other traders from trading on DEXs. By measuring the response of DEX volumes to changes in the market integrity indexes we show that traders can selectively engage with pools, avoiding those that show signs of manipulation, such as rug pulls or wash trading. For Liquidity Providers (LPs), they earn fees based on volume, regardless of if it is from a bad actor. With 12% of volume related to misconduct, 12% of the fees earnt by LPs comes from misconduct. Our evidence suggests that most other market participants are insensitive to misconduct. As a result, misconduct on DEXs does not appear to disrupt overall market function or deter trading activity. Regardless, the levels of misconduct in DEXs highlight the need for strong regulatory frameworks to appropriately police misconduct.

4.1.1 Related literature

Market integrity is a key part of market quality and directly shapes the confidence of market participants. Brogaard, Li, and Yang (2022) show that manipulative practices negatively affect market quality, reducing liquidity and price efficiency. Ensuring fairness in markets is a significant regulatory challenge. Foley, Karlsen, and Putniņš (2022) develop a global index to track market integrity, highlighting the role of regulatory interventions in mitigating misconduct. Traditional markets experience a wide range of misconduct and market manipulation, including spoofing, quote stuffing, and pump-and-dump schemes (Putniņš, 2020). These forms of misconduct are not limited to traditional markets. Bolz et al. (2024) and Dhawan and Putniņš (2023) quantify the impact of pump-and-dumps in cryptocurrency markets.

This chapter closely relates to research examining wash trading in cryptocurrency markets. In CEXs, Cong et al. (2023b) document widespread manipulation, showing abnormal distributions in trade sizes and transactions, with wash trading common on unregulated exchanges. Deng and Zhou (2024) propose methods to detect wash trading using market liquidity, while Sila et al. (2024) highlight how volatility, exchange flows, and public attention influence manipulative activity, often reflected in trade size roundness and deviations from Benford's Law. Existing studies typically infer wash trading statistically from trade data providing, at best, an estimation of wash trading.

In contrast, DEXs allow direct insight into trading activity, as every on-chain transaction incurs real costs and provides a traceable history of trading activity. Moreover, features of the market design on Automated Market Makers (AMMs) can be exploited to identify wash trading more clearly than in order book markets: circular trades in AMMs generate fees for the trader, whereas in order books similar patterns could reflect market-making activity rather than manipulation or misconduct. Gan et al. (2024) apply circular trading and clustering methods to Uniswap, identifying \$27.5 billion in wash trading across roughly 30 thousand pools. We extend this approach across multiple DEXs, identifying \$73 billion in wash trading volume from 11 exchanges on three different blockchains. Victor and Weintraud (2021) measure wash trading on order book–based DEXs using wallet clustering. In this chapter we avoid clustering wallets to focus on a lower bound estimate of wash trading, broad clustering algorithms may misclassify large networks of potentially unrelated traders.

In addition, this chapter also relates to research on Maximal Extractable Value (MEV), which relates to the profit a block builder can gain from reordering the transactions within a block. Daian et al. (2019) introduce MEV and identify sandwich attacks as a common form trading in DEXs with similarities to high-frequency trading, where traders profit by front- and back-running transactions. Qin, Zhou, and Gervais (2022) provide empirical evidence of the prevalence of sandwich attacks, liquidations, and arbitrage across AMM pools, reporting total profits of $174.3 million from sandwich attacks alone. Bagourd and Francois (2023) extend MEV analysis to Layer 2 blockchains, showing its relevance beyond Ethereum, particularly for MEV related arbitrage and liquidation opportunities.

Research on MEV has also explored ways to mitigate its impact on DEX traders. Heimbach and Wattenhofer (2022) propose reducing sandwich attack risks by adjusting trader slippage tolerances, while Heimbach, Pahari, and Schertenleib (2024) examine non-malicious forms of MEV, such as arbitrage, and their effects on gas fees. Despite these insights, the full extent of sandwich attacks across multiple blockchains and exchanges remains unknown. This chapter addresses that gap by quantifying the volume of trading activity associated with sandwich attacks across a range of DEXs on Etheruem, Polygon, and BSC.

Moreover, this chapter also relates to research on rug pulls, a form of exit scam in which a developer withdraws liquidity from a token pool on a DEX , collapsing its price and profiting at the expense of other participants. Aliyev, Allahverdiyeva, and Putniņš (2023) analyze Ethereum-based tokens, distinguishing between soft and hard rug pulls and identifying patterns in price, activity, and profitability that signal potential scams. Mazorra, Adan, and Daza (2022) demonstrate that features observable before a rug pull can predict scam tokens on Uniswap with high accuracy. Expanding beyond Ethereum, Cernera et al. (2023) document widespread rug pulls on both BSC and Ethereum, noting that while many (38.1% on Ethereum; 60% on BSC) do not yield profit, losses are minimal and a single successful trade can offset multiple failures.

DEXs can also serve as a tool for obfuscating on-chain transactions, contributing to money laundering and other illicit activity. A broad range of cryptocurrency-based crimes—including phishing, hacks, exploits, blackmail, ponzi schemes pig butchering, and romance scams—are documented, with studies highlighting their mechanisms and economic impact (Cong et al., 2025; Griffin & Mei, 2024; Hornuf et al., 2025). Theft from centralized exchanges is also substantial, with over $10 billion stolen, often involving loss of customer funds (Nguyen & Putniņš, 2023). Darknet markets have historically relied on cryptocurrencies to facilitate transactions, Foley, Karlsen, and Putniņš (2019) estimate that nearly half of all Bitcoin transactions are related to darknet markets. Ransomware attacks illustrate how a small number of actors can drive significant fund movements and affect blockchain dynamics such as transaction fees (Paquet-Clouston, Haslhofer, & Dupont, 2019; Sokolov, 2021). Studies on mixers and obfuscation techniques show how funds can be effectively laundered on public blockchains (Nadler & Schär, 2023; See, 2023). We build on this literature by quantifying the extent that DEXs are used by bad actors to obfuscate transactions.

Finally, this chapter relates to research on AMMs and DEXs. Barbon and Ranaldo (2022) examine AMM market quality, highlighting the influence of both transaction (gas) and market (swap) fees on market quality. Lehar and Parlour (2025) compare AMMs with limit order books, providing insight into liquidity provision and its implications for market efficiency. Aspris et al. (2021) emphasizes the security benefits of DEXs, noting that self-custody reduces counterparty risk. We add to this literature by quantifying the amount of misconduct—rug pulls, wash trading, sandwich attacks, and money laundering—in DEXs, highlighting challenges for regulating these currently unregulated markets.

4.2 Institutional background

Cryptocurrencies and blockchains pose unique oversight challenges for financial misconduct because of their pseudonymous nature, open access, and unclear regulatory boundaries. Unlike the traditional financial system, which has well-defined legal and regulatory structures, cryptocurrency markets lack standardized oversight, creating risks for participants.

DEXs, which facilitate peer-to-peer trading on these blockchains without centralized intermediaries, inherit these features. This raises concerns about market integrity in DEXs and the potential for misconduct with the challenges of overseeing these new markets.

4.2.1 Regulatory concerns

Globally, we are seeing growing concern with DEXs and market misconduct. Regulators view DEXs as high-risk venues for market misconduct because of their novel market structure and lack of centralized oversight. The European Securities and Markets Authority (ESMA) has highlighted concerns about manipulation in DeFi, particularly the emergence of MEV strategies such as sandwich attacks, which exploit blockchain transactions to front-run trades (European Securities and Markets Authority, 2023b). Additionally, ESMA categorizes different smart contracts used in DeFi, emphasizing that their automated and self-executing nature complicates enforcement and accountability (European Securities and Markets Authority, 2023a).

In the United States (U.S.), regulators have increasingly pursued enforcement actions against misconduct in cryptocurrency markets. The U.S. DOJ has charged entities with engaging in wash trading on cryptocurrency platforms, artificially inflating trading volumes to mislead investors, with some of this activity limited to tokens only listed on DEXs.[4] Regulatory action has extended to MEV-related exploits, the DOJ indictment of traders who manipulated the Ethereum blockchain to extract $25 million in illicit profits.[5] Broader concerns about illicit activities in DeFi are clear in the U.S. Treasury's sanctions against Tornado Cash, a privacy-focused protocol accused of facilitating money laundering for illicit actors.[6] These actions highlight the increasing regulatory scrutiny of DEXs and DeFi protocols, particularly in areas where traditional market manipulation is difficult to enforce.

[4]See https://bit.ly/4jJV674 for DOJ press release on wash trading.
[5]See https://bit.ly/3EVE071 for DOJ press release on MEV.
[6]See https://bit.ly/3Z6KkiP for U.S. treasury press release on Tornado Cash sanctions.

Regulatory oversight of DEXs remains a challenge because of the absence of centralized control. Unlike traditional exchanges, where compliance responsibilities fall on a central entity, DEXs distribute key functions across smart contracts, front-end interfaces, and network validators. This fragmentation complicates enforcement, as it is unclear which actors—protocol developers, transaction facilitators, or LPs—should bear responsibility for ensuring market integrity.

4.2.2 Decentralized exchanges and automated market makers

DEXs differ from traditional exchanges, as they leverage blockchain technology to execute trades through smart contracts rather than centralized order books. Most DEXs facilitate trades through an AMM. Using AMMs to facilitate trading reduces the data costs required to store a distributed ledger of a Limit Order Book (LOB) on a blockchain.

In AMMs, LPs contribute capital to pools (smart contracts), and traders execute transactions against these pools at prices determined by a predefined mathematical function. Unlike traditional market makers, LPs in AMMs do not actively set prices, AMM prices are set based on liquidity demander trading flow, with the LPs earning a small fee for every trade they facilitate.

The Constant Product Market Maker (CPMM), pioneered by Uniswap, is the first AMM design that gained significant popularity in DeFi. The liquidity pool maintains the product of the reserves of two assets constant, expressed as $L = \sqrt{r_1 r_2}$. In the CPMM, the quote price is simply the ratio of reserves $P = r_2/r_1$. In a CPMM, the amount received in a trade depends on the change in square root price and the invariant (L), with trades adjusting the reserves while holding the invariant L constant. The swap fee (typically 0.3%) acts like a spread and is taken from the asset that enters the pool. To have no impact on price, LPs add or remove funds in the pool based on the current price.

The Concentrated Liquidity Constant Product Market Maker (CL-CPMM) extends the CPMM by letting LPs allocate liquidity across specific price ranges rather than across all prices. Introduced in Uniswap V3, this refinement increases capital efficiency by focusing liquidity where trades are more likely to occur. Within each defined range $[P_a, P_b]$, the reserves satisfy a modified invariant, which leverages the real reserves in the pool. The CL-CPMM calculates trades just like the CPMM. This design leads to deeper markets near the active trading price while preserving the core features of an AMM. For a more detailed explanation of the CPMM and CL-CPMM, we refer to Chapter 2.

4.2.3 Market misconduct in decentralized exchanges

Market misconduct and manipulation take various forms, including but not limited to price manipulation, order-based strategies, benchmark manipulation, and abuse of market power. Putniņš (2020) provides a comprehensive overview of the different market manipulation strategies bad actors could employ. While market manipulation is mostly a legal term that encompasses various behaviors, the broader concept of market misconduct extends further to include any activity that undermines market integrity (i.e., money laundering).

The structural differences between AMMs and centralized markets have introduced new forms of market misconduct that were not possible in traditional markets. A key example of this is a sandwich attack. A sandwich attack is where a block builder exploits a traders transaction by both front and back running a trade are unique to blockchain based markets. To execute a sandwich attack, you need to order the transaction on the blockchain and add your own transactions, which is often reserved for block builders.

The design of AMMs also limits traditional forms of market manipulation. Manipulative practices such as spoofing, layering, and pinging rely on discretionary order placement and execution. AMMs, in contrast, eliminate these strategies as LPs let the AMM determine price and remove the ability for LPs to place traditional orders.

This chapter focuses on four key forms of market misconduct in DEXs: wash trading, money laundering, rug pulls, and sandwich attacks. Wash trading is a market manipulation strategy from traditional markets. In cryptocurrencies, wash trading is widespread across CEXs (Cong et al., 2023b). We add to this by examining the prevalence of wash trading in DEXs. Money laundering continues to attract significant media and regulatory attention. We examine the role of DEXs in the laundering processes. We also examine sandwich attacks identified by ESMA as misconduct unique to DEXs (European Securities and Markets Authority, 2023b). Last, we investigate rug pulls, a type of exit scam enabled by the low barriers to launching tokens and markets on DEXs.

4.3 Data and measures

We collect tick-level DEX data from CoinDesk Data and blockchain RPC nodes for the period June 2020 to December 2024. From CoinDesk Data, we get data on Uniswap v2 (Ethereum, BSC), Pancakeswap v2 and v3 (BSC), and Sushiswap v2 and v3 (Ethereum). We collect Uniswap v3 data on Ethereum from an Ethereum RPC node and collect Polygon-based data for Quickswap v2 and v3, as well as Uniswap v3, from a Polygon RPC node.

All collected data is denominated in token value. To convert prices to USD, we use CoinDesk Data's reference rates for common base tokens (e.g., WETH, USDC, USDT, DAI, WBTC). The reference rate is calculated as the 24-hour volume-weighted average price across multiple cryptocurrency exchanges.

Table 4.1 provides summary statistics for the 11 exchanges in our sample. Uniswap v3 on Ethereum is the largest by volume, averaging $1.1 billion daily and capturing nearly half of the market share by total volume. At a blockchain level, Ethereum leads the total volume share with 77.2% of volume traded, followed by BSC and Polygon with 18.1% and 4.7% respectively. CPMM exchanges list the most trading pairs, with Pancakeswap v2 on BSC offering 1.3 million and Uniswap v2 on Ethereum offering 395 thousand trading pairs.

Table 4.1: Summary statistics for exchanges

This table presents descriptive statistics for each DEX for the period June 2020 to December 2024. *Daily Volume* is the average daily volume in millions of dollars. *Daily Total Value Locked (TVL)* is the average daily total value locked in millions of dollars. *Pairs* is the total number of trading pairs on the DEX. *Volume share* is the percentage share of total volume for the exchange.

Exchange	Chain	Daily Volume ($M)	Daily TVL ($M)	Pairs	Volume Share (%)
Uniswap v2	Ethereum	359.2	2,438.8	391,308	21.03
Uniswap v2	BSC	0.3	0.5	1,049	0.00
Uniswap v3	Ethereum	1,103.7	2,602.9	29,184	49.72
Uniswap v3	Polygon	61.3	104.8	25,710	2.35
Pancakeswap v2	BSC	287.6	2,767.5	1,294,265	13.50
Pancakeswap v3	BSC	177.3	309.1	39,203	3.85
Sushiswap	Ethereum	131.6	1,308.5	4,138	7.26
Sushiswap	BSC	0.2	1.8	1,632	0.01
Sushiswap v3	Ethereum	0.4	2.0	288	0.01
Quickswap	Polygon	24.6	278.0	57,303	1.33
Quickswap v3	Polygon	33.8	38.1	4,535	0.95

4.3.1 Sandwich attacks

For a trade on a DEX to be published on-chain, it must first enter the mempool. The mempool is a list of all pending transactions from which block builders select transactions to include in the next block. All block builders can see the mempool and have the decision on whether to include any transaction it in the next block. Ideally, the block builder picks the transactions where they maximize their earnings from transaction (gas) fees.

However, because transactions in the mempool are visible to block builders, a malicious block builder can exploit the information of a DEX trade by including their own transaction to front run the trade, driving up the token price, and then back-running to restore their inventory. This activity, commonly known as a sandwich attack and lets block builders profit beyond their block's transaction (gas) fees. Since this entire process occurs within a single block, the block builder faces no risk. This practice is a form of MEV where block builders profit from their ability to include and reorder transactions in the block (Daian et al., 2019).[7] Sandwich attacks force traders to pay higher prices while also having greater price impacts because of the design of AMMs. While techniques to mitigate sandwich attacks exist, they remain a concern in DEXs (Heimbach & Wattenhofer, 2022).

[7]MEV is the profit that block producers or transaction orders can extract by controlling the ordering, inclusion, or exclusion of transactions within a block.

To measure sandwich attacks, we follow Qin, Zhou, and Gervais (2022) by identifying sequential trades within a pool in a single block. If a victim submits a buy order, a block builder can front-run by buying first, then back-run by selling after the victim's trade. We track patterns where an attacker's buy (sell) is immediately followed by the victim's buy (sell), and then the attacker's sell (buy), ensuring both attacker transactions originate from the same address. To ensure we capture the transactions that are only sandwich attacks and not by chance or coincidence, we only classify the transaction as a sandwich attack if the attacker trades with a similar size. For example, we exclude sandwich attacks where the front running trade is not within 1% of the volume of the back running trade.[8]

Table 4.2 presents summary statistics on sandwich attacks across the DEXs in our sample. We identify \$191.8 billion in trading volume associated with attacker trades and \$51.5 billion related to victim trades. Consistent with overall volume shares, most sandwich attacks occur on Uniswap v3 on Ethereum and Uniswap v2 on Ethereum, suggesting sandwich attacks are not specific to any exchange. At a blockchain level, Polygon and BSC have the least amount of sandwich attack volume, with \$0.6 and \$6.6 billion, respectively.

TABLE 4.2: Summary statistics for sandwich attacks

This table presents summary statistics on sandwich attacks for each DEX for the period June 2020 to December 2024. The *Mean Victim Trade* is the average trade size of the victim of the sandwich attack in dollars. *Total Victim Volume* is the total volume of the victim of the sandwich attack in millions of dollars. The *Mean Attacker Trade* is the average volume from the two attacker trades in the sandwich attack in dollars. *Total Attacker Volume* is the total volume from the two attacker trades in the sandwich attack in millions of dollars.

Exchange	Chain	Mean Victim Trade (\$)	Total Victim Volume(\$M)	Mean Attacker Trade (\$)	Total Attacker Volume(\$M)
Uniswap v2	Ethereum	5,078	18,717.0	13,407	49,590.0
Uniswap v2	BSC	373	0.4	1,425	1.6
Uniswap v3	Ethereum	41,224	24,605.1	213,506	127,435.6
Uniswap v3	Polygon	829	178.5	884	190.7
Pancakeswap v2	BSC	1,086	1,254.5	1,837	2,126.2
Pancakeswap v3	BSC	1,578	889.3	8,053	4,544.9
Sushiswap	Ethereum	32,501	5,361.4	45,548	7,517.9
Sushiswap	BSC	486	1.1	863	2.0
Sushiswap v3	Ethereum	5,465	4.8	56,778	50.0
Quickswap	Polygon	706	407.7	564	327.6
Quickswap v3	Polygon	778	31.5	1,052	42.6

[8]The criteria for similar volumes ensures that the attacker is selling approximately the same amount of tokens they initially purchased, confirming their intent to exploit the price movement rather than to hold the asset for potential investment purposes.

In our dataset, the address *jaredfromsubway.eth* is the most active MEV bot, completing approximately 1.5 million sandwich attacks and generating over $10 billion in trading volume on Ethereum DEXs.[9] In contrast, the largest sandwich attack bot on BSC records only $200 million in volume, while on Polygon the top bot accounts for only $30 million, indicating that sandwich attack activity is significantly smaller on these blockchains.[10] *jaredfromsubway.eth* is one of the more notorious MEV bots with early reports estimating the bot profits around $2 million a month while paying up to $10 million a month in gas fees.[11]

4.3.2 Wash trading

Wash trading is a form of market manipulation where a trader buys and sells the same asset to create the illusion of volume on a market. It suggests the market is more liquid and active than it is, ultimately hindering decision making for market participants.

In cryptocurrency, markets wash trading on CEXs is a huge concern, with exchanges creating fake volumes to end up high on the leader boards of cryptocurrency data providers to attract more customers. Cong et al. (2023b) examines and highlights the prevalence of wash trading in centralized cryptocurrency exchanges. Notably, CEXs can report trades without any underlying transfer of assets between users, allowing volumes to be artificially inflated. In contrast, DEX transactions are executed on-chain and leave a verifiable record of asset movement. As a result, identification of wash trading on DEXs reflects genuine economic activity rather than fabricated volume, providing a more accurate measure of real wash trading.

We leverage the on-chain transparency of DEXs to directly measure wash trading. Exploiting features of AMM design, we identify uneconomical trading patterns whose sole purpose is to inflate volume. In traditional order-book markets, circular trading—activity

[9] The address for the *jaredfromsubway.eth* MEV bot is 0xae2Fc483527B8EF99EB5D9B44875F005ba1FaE13

[10] The address for the most prominent BSC sandwich attacker is 0x215Ac1a98A7AD09439C2c3251b24918aA3d27aF2. The address for the most prominent Polygon sandwich attacker is 0x553db61A162483b3768b335CD61A42F8C782205c.

[11] See Eigenphi blog post on *jaredfromsubway.eth* here https://bit.ly/3ZnH2Ys

that nets to zero—is typically consistent with market making, where traders earn the spread while managing inventory risk. In AMMs, the pool acts as the counterparty, and circular trading is costly with the trader needing to pay the spread (swap fees). As a result, circular trades in AMMs most likely reflect wash trading rather than liquidity provision, providing a clear, direct identification strategy unavailable in order-book markets.

Following Gan et al. (2024), we identify wash trading in DEXs by looking at user level circular trading. We differ in the fact we do not cluster wallet addresses, instead we treat each wallet address as a single unique entity likely underestimating the true amount of wash trading. For each hour, we calculate the total amount of circular trading from a unique user based on the ratio of the $NetVolume/TotalVolume$. We classify volume as wash trading if the ratio is below a 5% threshold. Notably, we exclude any volume from sandwich attacks prior to the calculation of the wash trading volume, as these trades would appear similar to wash trading.

Table 4.3 presents summary statistics on wash trading across the sampled DEXs. Wash trading volumes are highly skewed, with medians substantially lower than means across all exchanges. We identify $72.7 billion in wash trading volume. Most wash trading concentrates on Uniswap v2 on Ethereum and Pancakeswap v2 on BSC, accounting for $40.6 billion and $11.2 billion of volume, respectively.

TABLE 4.3: Summary statistics for wash trading

This table presents summary statistics on wash trading for each DEX for the period June 2020 to December 2024. The *Mean Wash Trade* is the average wash trade volume in an hour in dollars. The *Median Wash Trade* is the median wash trade volume in an hour in dollars. *Total Value Wash Traded* is the total wash trade volume in millions of dollars.

Exchange	Chain	Mean Wash Trade ($)	Median Wash Trade ($)	Total Value Wash Traded ($M)
Uniswap v2	Ethereum	68,115	956	40,610.2
Uniswap v2	BSC	29,131	17	19.6
Uniswap v3	Ethereum	68,599	5,008	10,595.1
Uniswap v3	Polygon	1,468	101	526.2
Pancakeswap v2	BSC	9,223	109	11,188.9
Pancakeswap v3	BSC	3,451	336	943.6
Sushiswap	Ethereum	127,428	6,537	8,430.9
Sushiswap	BSC	3,279	24	34.3
Sushiswap v3	Ethereum	24,781	2,028	10.5
Quickswap	Polygon	540	40	166.6
Quickswap v3	Polygon	2,261	377	223.0

Notably, our identification of wash trading captures the pools and tokens outlined in the recent U.S. DOJ enforcement action on wash trading in cryptocurrency markets.[12] Across the four tokens publicly announced in the press release, we identify over $88 thousand in wash trading on DEXs. This finding helps validate our identification method while also highlighting how wide spread wash trading is across DEXs with $72.7 billion identified across our sample and the U.S. DOJ case only scratching the surface of the issue.

Wash trading activity varies substantially across chains and tokens. On Ethereum, the largest wash trading volumes are concentrated in smaller, less reputable tokens. For example, Tendies Token accounts for approximately $7.5 billion in wash trading volume across roughly 1,200 trades, while another token, Sickburn Inu records $5.2 billion across more than 1,100 trades. In contrast, on Polygon, wash trading occurs most frequently in larger token pairs, particularly WPOL traded against WETH, USDT0, and USDC.e, each of which has over 70,000 wash trades.

Wash trading activity is also highly concentrated at the wallet level. On Ethereum, a single address generates over $10 billion in wash trading volume across 2,420 trading pairs. This address typically executes paired buy and sell transactions of substantial size, often two swaps of approximately $150,000 within a single transaction, resulting in $300,000 of wash trading volume per event.[13] On BSC, the largest wash trader generates $2.3 billion in volume but follows a different strategy, maximizing the number of swaps per transaction rather than trade size.[14] By executing nearly 100 swaps within a single transaction, this address inflates reported volume with relatively small capital outlays. These individual examples highlight the significant heterogeneity in wash trading strategies on DEXs.

Reported trading volume on DEXs substantially overstates genuine economic activity. A large share of observed volume is generated by wash trading, which, while executed through verifiable on-chain transactions, does not reflect meaningful activity. Consequently, the prevalence of wash trading obscures the true scale of trading on DEXs and

[12] See https://bit.ly/4jJV674 for the DOJ press release on wash trading.

[13] The wallet address for the largest Ethereum wash trader is 0xF5dea4b05305636Dc37442F4BBEEA73dF2D7788D

[14] The wallet address for the largest BSC wash trader is 0x3fFD66b51d8687792E9965a8688458a3A6264bE2

creates challenges for assessing the market quality and benefits of DEXs and AMMs.

4.3.3 Rug pulls

Rug pulls are a prevalent form of market misconduct in DEXs. In a typical rug pull, a developer creates and promotes a new token, establishes a trading market—often via a DEX—and attracts liquidity from unsuspecting participants. Once sufficient capital has been drawn into the pool, the developer sells their holdings, intentionally collapsing the token's price and extracting value at the expense of other traders.

We classify a pool as a rug pull based on the following heuristics outlined by Aliyev, Allahverdiyeva, and Putniņš (2023). First, the pool must have a price pattern which follows a run-up and reversal—where the end price is lower than the pool's start price, after a price appreciation. Next, the pool has to be inactive. A pool is classified as inactive if the pool has no trades for the last 10% of its life. Next, the pool creator needs to profit from their activity in the pool.[15] The profit heuristic ensures that projects that would be considered a failure because the token creator also loses money do not get misclassified as a rug pull. Finally, we extend the heuristics outlined by Aliyev, Allahverdiyeva, and Putniņš (2023) to exclude any rug pull pools that have over $1 million in reserves after the rug pull. This extension ensures that any large pools with the same heuristics are not misclassified as rug pulls.

Table 4.4 presents summary statistics on rug pulls. We identify $56.4 billion in total volume from rug pull pairs. Price run ups and reversals are much smaller on concentrated liquidity DEXs like Uniswap v3, resulting lower losses for participants. Rug pulls are frequent on Uniswap v2 on Ethereum and Pancakeswap v2 on BSC with over 371 thousand rug pulls between the two DEXs and $18.2 and $22.2 billion in volume, respectively.

[15] We measure pool creator profits using the counterparty token—typically ETH or BNB—by tracking whether the creator's trades or liquidity provision within the pool result in a net increase in ETH (BNB) following the price crash. This measure is based exclusively on realized gains from trading activity in the pool, rather than changes in wallet balances across points in time.

Notably, we do not observe any rug pulls on Sushiswap v3 on Ethereum, although this exchange only trades 315 pairs.

Compared to Cernera et al. (2023), who report exit scam patterns in 65.6% of BSC pools and 81.1% of Ethereum pools, our analysis identifies substantially lower shares—22.0% of BSC pools and 20.3% of Ethereum pools. This suggests that our estimates take a conservative approach limting our classification to rug pulls and not exit scams more broadly. As a result, we are likely understate the true prevalence and volume of rug pulls across DEXs.

TABLE 4.4: Summary statistics for rug pulls

This table presents summary statistics on wash trading for each DEX for the period June 2020 to December 2024. *Price Run-up* is the median pair price run-up for the starting pool price and max pair price in percentage terms. *Price Reversal* is the median price reversal from the max price in the pair. *Rug Pulls* is the total number of rug pulls on the exchange. *Total Rug Pull Volume* is the total volume in the rug pull pools in millions of dollars.

Exchange	Chain	Price Run-up (%)	Price Reversal (%)	Rug Pulls	Total Rug Pull Volume ($M)
Uniswap v2	Ethereum	237.99	−99.92	75,779	18,165.9
Uniswap v2	BSC	38.07	−98.74	57	2.2
Uniswap v3	Ethereum	59.28	−70.49	10,426	15,420.4
Uniswap v3	Polygon	50.76	−64.56	4,012	146.5
Pancakeswap v2	BSC	261.10	−100.00	293,711	22,284.8
Pancakeswap v3	BSC	654.60	−98.57	253	75.6
Sushiswap	Ethereum	957.97	−99.98	340	298.1
Sushiswap	BSC	266.11	−99.79	52	0.1
Sushiswap v3	Ethereum			0	
Quickswap	Polygon	1,524.98	−100.00	39	4.6
Quickswap v3	Polygon	34.25	−76.47	378	2.3

4.3.4 Money laundering

Money laundering involves concealing the origins of illegally obtained funds. To measure its extent in DEXs, we compile a list of blockchain wallet addresses linked to mixers or hacks. Specifically, we track all wallets receiving funds from the Tornado Cash mixer on Ethereum, BSC, and Polygon, as well as wallets associated with known hacks, scams, and ransomware using data from the RektDB database. We estimate the amount of money laundering activity by calculating the trading volume from these addresses on each exchange.

Our measure likely represents an upper bound of money laundering, as not all mixer-related funds originate from illicit activities. However, mixers obscure fund origins and are frequently used by bad actors to hide the source of illicit funds, making some portion of this volume related to money laundering.

Table 4.5 presents summary statistics on money laundering across the sampled DEXs. Most illicit addresses are active on Uniswap v2 on Ethereum and Pancakeswap v2 on BSC.[16] In total, we identify $35 billion in volume from addresses we label money launderers, with $25 billion coming only from Uniswap v2 on Ethereum. We do not find any money laundering activity on either Quickswap exchange.

TABLE 4.5: Summary statistics for money laundering

This table presents summary statistics on money laundering for each DEX for the period June 2020 to December 2024. The *Mean Launder Trade* is the average trade volume from a launderer in dollars. *Total Launder Volume* is the total trade volume from launderers in millions of dollars. *Money Launderers* in the total number of unique money launderers on the DEX.

Exchange	Chain	Mean Launder Trade ($)	Total Launder Volume($M)	Money Launderers
Uniswap v2	Ethereum	25,596	25,262.4	15,472
Uniswap v2	BSC	85	0.1	20
Uniswap v3	Ethereum	77,150	671.3	51
Uniswap v3	Polygon	267,268	0.5	1
Pancakeswap v2	BSC	2,215	1,715.1	14,466
Pancakeswap v3	BSC	5,220	599.3	3,616
Sushiswap	Ethereum	28,746	7,372.9	8,119
Sushiswap	BSC	187	0.9	352
Sushiswap v3	Ethereum	2,882	2.3	103
Quickswap	Polygon			0
Quickswap v3	Polygon			0

We identify substantial trading volume on DEXs associated with addresses linked to money laundering activity. Addresses listed in the RektDB database account for $11.2 billion in trading volume, while addresses associated with Tornado Cash generate $34.3 billion. $9.9 billion in volume originates from addresses that appear in both datasets (88.5% of the volume from the RektDB database). As a result, a large share of the identified volume is driven by Tornado Cash related wallets, which likely represents an upper bound on laundering-related activity, given that mixer usage does not necessarily imply illicit intent.

[16] Appendix 4.2 presents summary statistics by address source, distinguishing between addresses identified in the RektDB database, Tornado Cash, and those appearing in both.

4.4 Market integrity indexes

From each measure of market misconduct—sandwich attack, wash trading, money laundering, and rug pulls—we develop indexes to measure the market integrity. Our index can be interpreted as the proportion of DEX activity that is related to market misconduct. For each measure of misconduct, we normalize the contribution of volume from the form of misconduct by the total exchange volume.

We develop an index for sandwich attacks based on the trading volume that can be attributed to the attacker. Specifically, the trading volume of the trades that front run and back run in a sandwich attack. For each month, we normalize the attacker volume on the total exchange volume to get an index on sandwich attack (SAI):

$$SAI_{i,t} = \frac{SandwichAttackerVolume_{i,t}}{TotalVolume_{i,t}}, \tag{4.1}$$

where SAI is the sandwich attack index for exchange chain i for month t.

Similarly, from the volume associated with wash trading on an exchange, we develop a wash trading index (WTI). For each exchange, we normalize the total wash trading volume by the total trading volume:

$$WTI_{i,t} = \frac{WashTradingVolume_{i,t}}{TotalVolume_{i,t}}, \tag{4.2}$$

where WTI is the wash trading index for exchange chain i for month t.

Our index to track rug pulls (RPI) is based on the total trading volume in pools we classify as rug pulls. This volume includes the volume associated with the initiator of the pool but also the other traders who purchased the token. For each exchange, we normalize the total rug pull volume by the total trading volume:

$$RPI_{i,t} = \frac{RugPullVolume_{i,t}}{TotalVolume_{i,t}}, \tag{4.3}$$

where RPI is the rug pull index for exchange chain i for month t.

Finally, our index on money laundering (MLI) tracks the swap volume of wallets that may be related to money laundering—wallets associated with known hacks, scams, and ransomware along with wallets that receive funds from the Tornado Cash mixer. We normalize the total volume associated with these wallets by the total volume on the exchange to get an index of money laundering:

$$MLI_{i,t} = \frac{MoneyLaunderingVolume_{i,t}}{TotalVolume_{i,t}}, \tag{4.4}$$

where MLI is the money laundering index for exchange chain i for month t.

We combine these indexes into a market integrity index (MII) to track the total amount of misconduct in DEXs. The MII gives a measure of market integrity that can be interpreted as the percentage of volume on DEXs that is related to market misconduct. We calculate the MII by taking the sum of the four indexes in each measure of market misconduct:

$$MII_{i,t} = SAI_{i,t} + WTI_{i,t} + RPI_{i,t} + MLI_{i,t}, \tag{4.5}$$

where MII is the market integrity index for exchange chain i for month t. SAI is the sandwich attack index, WTI is the wash trading index, RPI is the rug pull index and MLI is the money laundering index.

In the market integrity index MII, we do not double count volume across the different indexes to maintain a proportion of volume on that exchange, which takes a value between 0 and 1. For example, a rug pull pool may also have wash trading in order to attract other market participants. In these instances, we only consider the volume once to avoid double counting. Overall, we identify $345 billion in trading volume associated with misconduct in DEXs.

Table 4.6 presents average index values across DEXs for wash trading, sandwich attacks, rug pulls, money laundering, and overall market integrity. Wash trading is the most widespread form of misconduct, with the proportion of volume linked to wash trading over

5% for multiple exchanges, reaching as high as 33% on Uniswap v2 on BSC. Sandwich attacks account for nearly 10% of trading volume on Uniswap v2 and v3 on Ethereum. Rug pulls are concentrated in CPMM-based markets on Ethereum and BSC, comprising 4% of volume on Uniswap v2 (Ethereum and BSC) and 8% on Pancakeswap v2 on BSC. In contrast, money laundering and rug pull volumes are typically below 1% on most DEXs and never exceed 4%. Misconduct is concentrated on certain exchanges, 25% of volume on Uniswap v2 on Ethereum and 10% of volume on Uniswap v3 on Ethereum is related to market misconduct. On average, 12% of total DEX volume is associated with misconduct, implying that roughly one in every eight dollars traded on DEXs is related to misconduct.

Table 4.6: Summary statistics for the market integrity indexes

This table presents average index value for each DEX for the period June 2020 to December 2024. The table presents the average for the wash trading index (WTI), the sandwich attack index (SAI), the rug-pull index (RPI), the money laundering index (MLI), and the market integrity index (MII). Each index is a proportion of the volume associated with the form of misconduct and the total volume on that exchange. Values are reported in percentages.

Exchange	Chain	WTI	SAI	RPI	MLI	MII
Uniswap v2	Ethereum	8.59	9.18	4.37	3.86	24.28
Uniswap v2	BSC	33.33	2.12	4.16	0.32	37.52
Uniswap v3	Ethereum	0.81	9.07	1.16	0.03	10.98
Uniswap v3	Polygon	0.84	0.24	0.25	0.00	1.32
Pancakeswap v2	BSC	4.74	1.03	8.58	0.53	13.95
Pancakeswap v3	BSC	0.85	3.21	0.06	0.82	4.92
Sushiswap	Ethereum	10.75	3.90	0.17	3.30	17.17
Sushiswap	BSC	16.50	1.08	0.04	0.18	17.79
Sushiswap v3	Ethereum	4.59	14.75	0.00	2.87	22.03
Quickswap	Polygon	1.09	2.08	0.03	0.00	3.20
Quickswap v3	Polygon	0.78	0.12	0.02	0.00	0.92

Figure 4.1 illustrates the average index values across the DEXs in our sample. Wash trading is the primary driver of the overall market integrity index through time, accounting for an average of 6% of DEX volume. On average, the market integrity index suggests that 12% of total volume across all DEXs stems from market misconduct. Misconduct levels remain relatively stable over time, though wash trading shows an upward trend beginning in 2021 and continuing through to the end of our sample.[17]

[17] Appendix 4.3 illustrates the total dollar amount of market misconduct in each month. Notably the dollar amount of misconduct appears flat, with no significant trend to suggest it is growing.

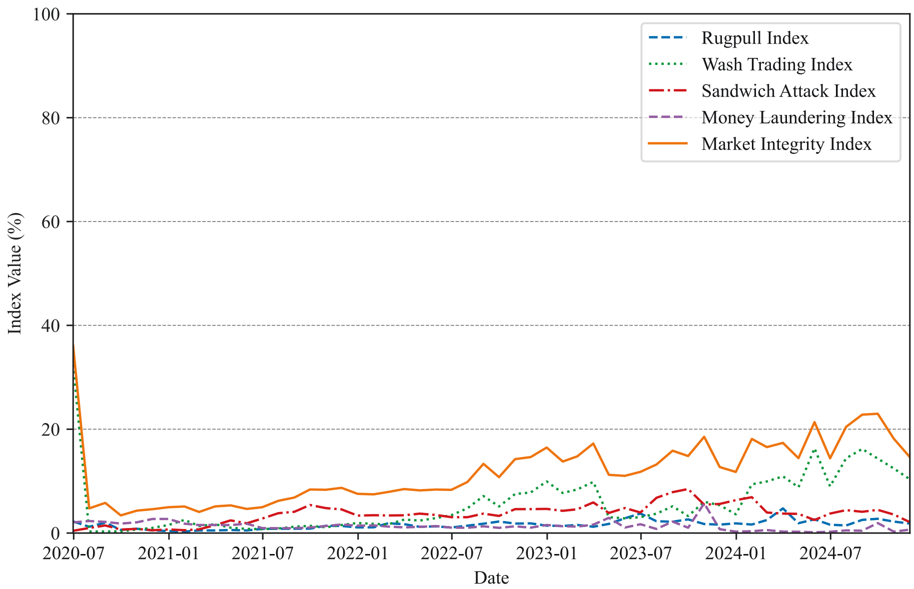

FIGURE 4.1: Market integrity indexes

This figure plots the average index values across the different DEXs for the sample period from June 2020 to December 2024. The solid orange line represents the market integrity index. The dotted green line represents the wash trading index. The dash-dotted red line represents the sandwich attack index. The dashed purple line represents the money laundering index.

4.5 Empirical analysis

4.5.1 What explains market misconduct in decentralized exchanges?

Given the scale of market misconduct observed, a natural question arises: are certain DEXs more attractive to bad actors than others? Specifically, we examine whether misconduct concentrates in high-activity venues—suggesting opportunistic behavior that follows liquidity—or whether it is more prevalent in smaller exchanges where oversight and monitoring may be weaker.

To test whether specific exchange characteristics attract misconduct, we estimate the following regression model:

$$Integrity_{i,t} = \alpha + Volume_{i,t} + TVL_{i,t} + Pairs_{i,t} + CPMM_i + \epsilon_{i,t}, \tag{4.6}$$

where $Integrity_{i,t}$ is a measure of market integrity on exchange i in month t, based on the indexes for wash trading, sandwich attacks, rug pulls, money laundering, or the total market integrity index. $Volume_{i,t}$ is the monthly trading volume, $TVL_{i,t}$ is the TVL on the exchange, and $Pairs_{i,t}$ is the number of token pairs listed. $CPMM_i$ is a dummy variable equal to 1 if the exchange operates using a CPMM design, or otherwise 0.

Table 4.7 reports Ordinary Least Squares (OLS) estimates from Equation 4.6, where we regress market integrity indexes on exchange-level characteristics. Wash trading is positively related with trading volume and negatively related with TVL. In contrast, sandwich attacks are positively related with the TVL on an exchange and negatively related with number of trading pairs. Overall, different types of misconduct exhibit distinct relationships with exchange characteristics, highlighting the structural factors that shape integrity risks on DEXs.

Table 4.7: Regressions to assess the determinants of market integrity

This table reports coefficient estimates from OLS regressions using exchange-month observations to examine the determinants of the different market integrity indexes. The dependent variables are the different market integrity indexes. The dependent variable in Regression 1 is WTI the wash trading index. The dependent variable in Regression 2 is SAI the sandwich attack index. The dependent variable in Regression 3 is RPI the rug pull index. The dependent variable in Regression 4 is MLI the money laundering index. The dependent variable in Regression 5 is MII the market integrity index. $Volume$ is the dollar volume for the exchange in a month. TVL is the total value locked for the exchange in a month. $Pairs$ is the total number of trading pairs at the end of the month. $CPMM$ is a dummy variable equal to 1 if the exchange trades based on a constant product market maker, otherwise 0. The sample comprises 398 observations from 11 exchanges for the period June 2020 to December 2024. Standard errors are clustered by month. All explanatory variables are in natural log form. t-statistics are reported in parentheses. Significance at the 10%, 5%, and 1% levels is indicated by *, **, and ***, respectively.

	(1)	(2)	(3)	(4)	(5)
Dependent variable:	WTI	SAI	RPI	MLI	MII
$Volume$	0.027***	−0.005	0.003	0.002	0.025***
	(3.88)	(−1.64)	(1.15)	(0.45)	(2.80)
$Pairs$	0.002	−0.004**	0.010***	−0.003***	0.005
	(0.75)	(−2.05)	(13.20)	(−3.38)	(1.27)
TVL	−0.042***	0.010***	−0.005*	0.001	−0.034***
	(−5.04)	(3.65)	(−1.76)	(0.27)	(−3.53)
$CPMM$	0.121***	−0.020**	0.014**	0.017**	0.122***
	(5.97)	(−2.37)	(2.50)	(2.16)	(5.05)
Intercept	0.210***	−0.009	−0.045***	−0.027	0.141**
	(3.33)	(−0.34)	(−6.77)	(−1.29)	(2.09)
Observations	398	398	398	398	398
Adjusted R^2	19.5%	4.3%	35.8%	4.6%	9.9%

Significance: * $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$

Unsurprisingly, Regression 3 shows that the rug pull index is positively related to the number of listed trading pairs. Bad actors need to create tokens and list them on DEXs to rug pull. Not every rug pull is a success, and these bad actors often create thousands of new tokens and markets to attract new investors.

The most significant finding is in the different market designs. The CPMM dummy is significantly associated with several misconduct measures. Specifically, CPMM DEXs have higher levels of wash trading and rug pulls, but experience fewer sandwich attacks. Overall, the total market integrity index is significantly higher for CPMM DEXs, suggesting that market design has an impact on misconduct in DEXs. These findings highlight how that the underlying architecture of a DEX both blockchain and market design influences the prevalence and type of misconduct.

Next, we investigate whether blockchain-level factors contribute to differences in market misconduct across DEXs. Variations in transaction ordering and block construction across different blockchains may influence how sandwich attacks—and other forms of misconduct—can be executed. This motivates whether specific blockchain features impact the level of market misconduct observed on DEXs.

To test whether there are any blockchain specific factors that influence DEX market integrity, we estimate the following regression:

$$Integrity_{i,t} = \alpha + ChainDummy_i + Volume_{i,t} + TVL_{i,t} + Pairs_{i,t} + \epsilon_{i,t}, \tag{4.7}$$

where $Integrity_{i,t}$ is a measure of market integrity on exchange i in month t, based on the indexes for wash trading, sandwich attacks, rug pulls, money laundering, or the total market integrity index. $ChainDummy_i$ is a set of dummy variables indicating whether an exchange operates on BSC or Polygon, with Ethereum serving as the baseline because of its large share of trading volume. In the model, we control for exchange-level characteristics including trading volume, TVL, and the number of trading pairs.[18]

[18]The CPMM dummy is excluded to preserve statistical power; with the small sample, including fixed-effect dummies such as chain indicators would absorb variation without improving explanatory power.

Table 4.8 presents estimates from an OLS regression from Equation 4.7. Blockchain factors significantly impact the level of market misconduct on DEXs. Polygon-based exchanges show substantially lower misconduct across all integrity measures compared to Ethereum and BSC, as shown by the consistently negative and statistically significant Polygon dummy variable in all regressions. Both BSC and Polygon have fewer sandwich attacks and lower instances of wash trading and money laundering compared to Ethereum. We find no significant difference in the share of rug pull volume between BSC and Ethereum, indicating that DEXs on both blockchains exhibit similar exposure to this form of misconduct. Notably, the total market integrity index is approximately 25% lower on Polygon and 15% lower on BSC compared to Ethereum-based DEXs. Overall, these findings suggest that both market design and blockchain based factors have a meaningful impact on the prevalence of misconduct in DEXs.

TABLE 4.8: Regressions to assess blockchain determinants of market integrity

This table reports coefficient estimates from OLS regressions using exchange-month observations to examine the determinants of the different market integrity indexes. The dependent variables are the different market integrity indexes. The dependent variable in Regression 1 is WTI the wash trading index. The dependent variable in Regression 2 is SAI the sandwich attack index. The dependent variable in Regression 3 is RPI the rug pull index. The dependent variable in Regression 4 is MLI the money laundering index. The dependent variable in Regression 5 is MII the market integrity index. BSC is a dummy variable equal to 1 if the exchange trades on BSC, otherwise 0. $Polygon$ is a dummy variable equal to 1 if the exchange trades on polygon, otherwise 0. $Volume$ is the dollar volume for the exchange in a month. TVL is the total value locked for the exchange in a month. $Pairs$ is the total number of trading pairs at the end of the month. The sample comprises 398 observations from 11 exchanges for the period June 2020 to December 2024. Standard errors are clustered by month. All explanatory variables are in natural log form. t-statistics are reported in parentheses. Significance at the 10%, 5%, and 1% levels is indicated by *, **, and ***, respectively.

	(1)	(2)	(3)	(4)	(5)
Dependent variable:	WTI	SAI	RPI	MLI	MII
BSC	−0.039**	−0.091***	0.000	−0.028***	−0.149***
	(−2.12)	(−9.31)	(0.11)	(−3.29)	(−5.56)
$Polygon$	−0.117***	−0.091***	−0.026***	−0.030***	−0.250***
	(−9.59)	(−10.89)	(−9.21)	(−3.98)	(−18.60)
$Volume$	0.002	0.007***	0.001	−0.001	0.008*
	(0.38)	(4.61)	(1.15)	(−0.60)	(1.99)
$Pairs$	0.017***	0.008***	0.012***	0.002	0.036***
	(8.73)	(6.07)	(11.29)	(0.85)	(13.11)
TVL	−0.029***	−0.015***	−0.005***	−0.001	−0.047***
	(−4.58)	(−5.06)	(−2.88)	(−0.18)	(−6.90)
Intercept	0.465***	0.157***	−0.014*	0.049**	0.638***
	(7.62)	(3.99)	(−1.78)	(2.32)	(11.64)
Observations	398	398	398	398	398
Adjusted R^2	20.6%	34.0%	42.8%	11.8%	35.5%

Significance: * $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$

4.5.2 Does misconduct impact market activity?

While we identify significant levels of market misconduct, a central question is whether other market participants care about misconduct? Specifically, do legitimate users avoid exchanges with higher levels of misconduct, such as frequent rug pulls or sandwich attacks? Or do they continue transacting despite integrity concerns? This distinction is critical for understanding whether misconduct has a real effect or whether its impacts are largely localized in smaller pools and trading pairs.

We can test whether other market participants care about misconduct by estimating the following regression:

$$\begin{aligned} \Delta Activity_{i,t} = \alpha + \Delta Integrity_{i,t} + \Delta Integrity_{i,t-1} + Controls_{i,t} \\ + ExchangeEffects_i + \epsilon_{i,t}, \end{aligned} \tag{4.8}$$

where $\Delta Activity_{i,t}$ represents the change in exchange activity, measured either by trading volume or TVL, and $\Delta Integrity_{i,t}$ are the changes in the different market integrity indexes. We include a one-period lag in the integrity measure to account for any delayed response by market participants to observed misconduct. We also control for exchange fixed effects and exchange characteristics by including volume, TVL, and the number of trading pairs in the regressions. This regression model allows us to assess whether changes in market integrity deter activity on DEXs.

Table 4.9 reports estimates from the panel OLS regression described in Equation 4.8, with exchange activity measured by changes in trading volume not related to misconduct. Examining the volume not related to misconduct ensures that any contemporaneous relation between the volume based indexes and the measure of volume remain unrelated.

We find limited evidence to suggest that misconduct on DEXs deters broader trading activity. In the contemporaneous period, only money laundering shows a significant positive relationship with changes in non-misconduct volume, suggesting that launderers target high-activity exchanges to obscure flows. In the lagged period, only changes in the

Table 4.9: Regressions to assess if market integrity impacts volume

This table reports coefficient estimates from panel OLS regressions using exchange-month observations to examine whether misconduct impacts volume on DEXs The dependent variable is the change in the natural log of dollar volume excluding the volume related to misconduct ($\Delta Volume_c$). ΔWTI is the change in the wash trading index WTI. ΔSAI is the change in the sandwich attack index SAI. ΔRPI is the change in the rug pull index RPI. ΔMLI is the change in the money laundering index MLI. ΔMII is the change in the market integrity index MII. The sample comprises 376 observations from 11 exchanges for the period June 2020 to December 2024. Standard errors are clustered by exchange month. t-statistics are reported in parentheses. Significance at the 10%, 5%, and 1% levels is indicated by *, **, and ***, respectively.

	(1)	(2)	(3)	(4)	(5)
Dependent variable:	$\Delta Volume_c$	$\Delta Volume_c$	$\Delta Volume_c$	$\Delta Volume_c$	$\Delta Volume_c$
ΔWTI	−0.061 (−0.12)				
ΔWTI_{t-1}	0.004 (0.01)				
ΔSAI		0.268 (0.22)			
ΔSAI_{t-1}		0.642 (0.89)			
ΔRPI			1.383 (1.35)		
ΔRPI_{t-1}			2.190** (2.34)		
ΔMLI				1.765*** (6.69)	
ΔMLI_{t-1}				−1.024 (−1.40)	
ΔMII					0.349 (0.70)
ΔMII_{t-1}					0.026 (0.08)
Controls	Y	Y	Y	Y	Y
Exchange Effects	Y	Y	Y	Y	Y
Observations	376	376	376	376	376
R^2	20.0%	20.2%	20.7%	22.6%	20.3%

Significance: * $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$

rug pull index show a significant positive relation with future volume, suggesting that past increases in rug pulls lead to future exchange volume implying that increases in misconduct are related to increases in activity, rather than driving users away.

We do not find any significant evidence that misconduct reduces legitimate trading volume on DEXs. Traders on AMMs appear indifferent to misconduct, likely because of structural and informational features of these markets. The counterparty to each trade is a smart contract, not a potentially exploitative intermediary. Transparent on-chain data enables traders to screen pools for signs of manipulation and avoid those that exhibit misconduct. As a result, users can select pools and tokens that match their risk tolerance, limiting their exposure to market abuse. Overall, these findings highlight that misconduct in DEXs does not disrupt broader market activity.

Next, we examine whether market misconduct affects activity based on changes to TVL by estimating the same panel OLS regression from Equation 4.8. Changes in TVL serve as a measure of LP activity. A decline in TVL reflects a withdrawal of liquidity and may signal that market participants expect lower future trading activity and thus lower fees.

Table 4.10 presents results from the panel OLS regression specified in Equation 4.8, where exchange activity is measured by changes in TVL. Across specifications, the market integrity indexes show no significant relation with TVL, except for a weakly significant (10%) negative contemporaneous relationship with the money laundering index.

TABLE 4.10: Regressions to assess if market integrity impacts total value locked

This table reports coefficient estimates from panel OLS regressions using exchange-month observations to examine whether misconduct impacts TVL on DEXs The dependent variable is the change in the natural log of the total value locked (ΔTVL). ΔWTI is the change in the wash trading index WTI. ΔSAI is the change in the sandwich attack index SAI. ΔRPI is the change in the rug pull index RPI. ΔMLI is the change in the money laundering index MLI. ΔMII is the change in the market integrity index MII. The sample comprises 376 observations from 11 exchanges for the period June 2020 to December 2024. Standard errors are clustered by exchange month. t-statistics are reported in parentheses. Significance at the 10%, 5%, and 1% levels is indicated by *, **, and ***, respectively.

	(1)	(2)	(3)	(4)	(5)
Dependent variable:	ΔTVL	ΔTVL	ΔTVL	ΔTVL	ΔTVL
ΔWTI	−0.137				
	(−1.30)				
ΔWTI_{t-1}	−0.133				
	(−0.42)				
ΔSAI		0.001			
		(0.00)			
ΔSAI_{t-1}		−0.240			
		(−1.26)			
ΔRPI			0.169		
			(0.93)		
ΔRPI_{t-1}			−0.732		
			(−1.28)		
ΔMLI				−0.361*	
				(−1.86)	
ΔMLI_{t-1}				−1.113	
				(−1.08)	
ΔMII					−0.135
					(−1.64)
ΔMII_{t-1}					−0.290
					(−1.08)
Controls	Y	Y	Y	Y	Y
Exchange Effects	Y	Y	Y	Y	Y
Observations	376	376	376	376	376
R^2	15.8%	15.7%	16.0%	16.9%	16.2%

Significance: * $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$

For LPs, misconduct generates volume, leading to higher fees. Given that misconduct accounts for roughly 12% of total DEX volume, this implies that 12% of LP fee revenue stems from market misconduct. Introducing regulation to remove misconduct would therefore reduce LP revenues by approximately 12%, helping explain their indifference to integrity issues.

4.5.3 Are bad actors sensitive to blockchain fees?

DEXs operate on blockchain networks, where transaction costs, commonly referred to as gas fees, can significantly impact trading behavior. If these fees are sufficiently high, they could diminish the incentives for bad actors to engage in market misconduct, as the costs of participating in such activities may outweigh any potential gains. To examine whether blockchain transaction fees influence market integrity on DEXs, we estimate the following regression model:

$$Integrity_{i,t} = \alpha + Gas_{i,t} + Volume_{i,t} + TVL_{i,t} + Pairs_{i,t} + ChainEffects_i + \epsilon_{i,t}, \quad (4.9)$$

where $Integrity_{i,t}$ is a measure of market integrity on exchange i in month t, based on the indexes for wash trading, sandwich attacks, rug pulls, money laundering, or the total market integrity index. Gas is the average price of gas in dollars on the blockchain for the month. We control for volume, TVL, the number of trading pairs, and chain fixed effects to capture the variation within each blockchain.

Table 4.11 presents estimates from OLS regressions based on Equation 4.9. The findings suggest that only rug pulls are significantly sensitive to transaction costs (gas fees). The creation of tokens and liquidity pools required for rug pulls can be highly gas-intensive, so when gas fees are high, the transaction costs for executing a rug pull increase, potentially deterring bad actors (Cernera et al., 2023).

In contrast, other forms of market manipulation, such as sandwich attacks, are less sensitive to gas fees. Sandwich attacks guarantee a profit, even after accounting for transaction costs. As a result, changes in gas prices do not appear to impact the frequency or profitability of sandwich attacks.

TABLE 4.11: Regressions to assess the impact of blockchain fees on market integrity

This table reports coefficient estimates from OLS regressions using exchange-month observations to examine the impact of blockchain fees on the different market integrity indexes. The dependent variables are the different market integrity indexes. The dependent variable in Regression 1 is WTI the wash trading index. The dependent variable in Regression 2 is SAI the sandwich attack index. The dependent variable in Regression 3 is RPI the rug pull index. The dependent variable in Regression 4 is MLI the money laundering index. The dependent variable in Regression 5 is MII the market integrity index. Gas is the average gas price for a transaction to be included on the blockchain. $Volume$ is the dollar volume for the exchange in a month. TVL is the total value locked for the exchange in a month. $Pairs$ is the total number of trading pairs at the end of the month. The sample comprises 398 observations from 11 exchanges for the period June 2020 to December 2024. Standard errors are clustered by month. All explanatory variables are in natural log form. t-statistics are reported in parentheses. Significance at the 10%, 5%, and 1% levels is indicated by *, **, and ***, respectively.

	(1)	(2)	(3)	(4)	(5)
Dependent variable:	WTI	SAI	RPI	MLI	MII
Gas	0.001	0.001	−0.002*	0.000	0.000
	(0.36)	(0.52)	(−2.00)	(−0.27)	(0.03)
$Volume$	0.001	0.007***	0.002	−0.001	0.008*
	(0.35)	(4.58)	(1.29)	(−0.56)	(1.97)
$Pairs$	0.017***	0.008***	0.012***	0.002	0.036***
	(8.38)	(6.17)	(11.19)	(0.86)	(12.88)
TVL	−0.029***	−0.015***	−0.005***	−0.001	−0.047***
	(−4.54)	(−5.07)	(−2.77)	(−0.18)	(−6.85)
Intercept	0.447***	0.076**	−0.044**	0.016	0.491***
	(3.92)	(2.10)	(−2.43)	(0.59)	(4.26)
Chain Effects	Y	Y	Y	Y	Y
Observations	398	398	398	398	398
Adjusted R^2	20.4%	33.9%	42.9%	11.5%	35.3%

Significance: * $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$

4.6 Conclusion

There is growing regulatory concern over how—or even whether—to regulate DEXs. The main challenge lies in their novel market structure, which lacks centralized oversight. This chapter contributes to that policy debate by quantifying the scale of market misconduct in DEXs.

We identify over $345 billion in volume related to market misconduct across 11 DEXs, accounting for approximately 12% of total DEX trading volume. Using detailed transaction-level data, we construct market integrity indexes for the proportion of volume related to misconduct based on: wash trading ($73 billion, 5.6%), sandwich attacks ($192 billion, 3.6%), rug pulls ($56 billion, 1.6%), and money laundering ($35 billion, 1.3%).

Our findings highlight the need for targeted regulatory frameworks for DEXs, given the scale of misconduct observed. However, while regulation traditionally aims to promote fairness in markets, our evidence shows that other participants are insensitive to market integrity issues. Traders can selectively engage with pools, only trading tokens where they accept any risks. Meanwhile, LPs continue to earn fees on all trading volume, including volume driven by misconduct. As a result, misconduct on DEXs primarily harms directly affected traders (such as sandwich attack victims) and does not disrupt broader market function.

Appendix 4.1 Correlation table

TABLE A4.1: Correlation of market integrity variables

This table presents the correlation matrix of variables used in regressions on the market integrity results. WTI is the wash trading index. SAI is the sandwich attack index. RPI is the rug pull index. MLI is the money laundering index. MII is the market integrity index. *Pairs* is the total number of trading pairs at the end of the month. TVL is the total value locked for the exchange in a month. $Volume$ is the dollar volume for the exchange in a month. *Gas* is the average gas price for a transaction to be included on the blockchain. $CPMM$ is a dummy variable equal to 1 if the exchange trades based on a constant product market maker, otherwise 0. *Ethereum* is a dummy variable equal to 1 if the exchange trades on Ethereum, otherwise 0. *Polygon* is a dummy variable equal to 1 if the exchange trades on polygon, otherwise 0. BSC is a dummy variable equal to 1 if the exchange trades on BSC, otherwise 0.

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
WTI (1)	1.00											
SAI (2)	-0.01	1.00										
RPI (3)	0.13	0.05	1.00									
MLI (4)	0.09	0.07	0.09	1.00								
MII (5)	0.86	0.41	0.34	0.32	1.00							
Pairs (6)	0.01	-0.09	0.79	-0.01	0.15	1.00						
TVL (7)	-0.17	0.18	0.20	0.09	-0.01	0.24	1.00					
$Volume$ (8)	-0.17	0.25	0.03	0.01	-0.03	-0.02	0.69	1.00				
$CPMM$ (9)	0.30	-0.13	0.27	0.15	0.28	0.29	0.14	-0.25	1.00			
Ethereum (10)	0.02	0.55	-0.01	0.34	0.29	-0.21	0.38	0.43	-0.02	1.00		
Polygon (11)	-0.26	-0.32	-0.28	-0.22	-0.45	-0.22	-0.39	-0.31	-0.25	-0.53	1.00	
BSC (12)	0.23	-0.28	0.29	-0.15	0.12	0.44	-0.02	-0.16	0.26	-0.56	-0.41	1.00

Appendix 4.2 Summary statistics for money laundering by address source

Table A4.2: Summary statistics for money laundering split by source of address

This table presents summary statistics on money laundering for each DEX for the period June 2020 to December 2024. Panel A presents the summary statistics for address that are from the RektDB database. Panel B presents the summary statistics for address that are from interactions with Tornado Cash. Panel C presents the summary statistics for address that are from both Tornado Cash and RektDB database. The *Mean Launder Trade* is the average trade volume from a launderer in dollars. *Total Launder Volume* is the total trade volume from launderers in millions of dollars. *Money Launderers* in the total number of unique money launderers on the DEX.

Panel A. RektDB database

Exchange	Chain	Mean Launder Trade ($)	Total Launder Volume($M)	Money Launderers
Uniswap v2	Ethereum	535,861	10,233.3	159
Uniswap v2	BSC	2,912	0.0	1
Uniswap v3	Ethereum	76,347	655.4	47
Uniswap v3	Polygon	267,268	0.5	1
Pancakeswap v2	BSC	73,191	128.3	116
Pancakeswap v3	BSC	54,340	32.3	26
Sushiswap	Ethereum	32,577	203.2	83
Sushiswap	BSC	327	0.0	4
Sushiswap v3	Ethereum	1,842	0.1	5
Quickswap	Polygon			0
Quickswap v3	Polygon			0

Panel B. Tornado Cash

Exchange	Chain	Mean Launder Trade ($)	Total Launder Volume($M)	Money Launderers
Uniswap v2	Ethereum	25,525	24,818.7	15,365
Uniswap v2	BSC	85	0.1	20
Uniswap v3	Ethereum	136,051	15.9	4
Uniswap v3	Polygon			0
Pancakeswap v2	BSC	2,110	1,632.1	14,409
Pancakeswap v3	BSC	5,124	587.2	3,605
Sushiswap	Ethereum	28,898	7,278.9	8,065
Sushiswap	BSC	187	0.9	351
Sushiswap v3	Ethereum	2,934	2.2	100
Quickswap	Polygon			0
Quickswap v3	Polygon			0

Panel C. Tornado Cash and RektDB database

Exchange	Chain	Mean Launder Trade ($)	Total Launder Volume($M)	Money Launderers
Uniswap v2	Ethereum	2,189,577	9,789.6	52
Uniswap v2	BSC	2,912	0.0	1
Uniswap v3	Ethereum			0
Uniswap v3	Polygon			0
Pancakeswap v2	BSC	48,934	45.3	59
Pancakeswap v3	BSC	51,215	20.2	15
Sushiswap	Ethereum	66,510	109.3	29
Sushiswap	BSC	466	0.0	3
Sushiswap v3	Ethereum	1,043	0.0	2
Quickswap	Polygon			0
Quickswap v3	Polygon			0

Appendix 4.3 Monthly dollar amount of market misconduct

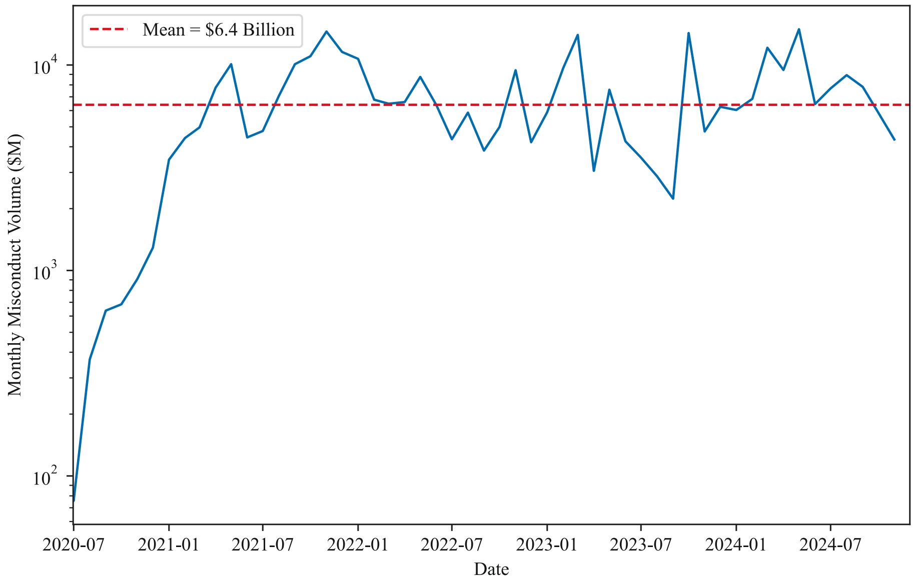

FIGURE A4.1: Monthly dollar value of market misconduct in decentralized exchanges

This figure plots the total monthly dollar volume of misconduct across the different DEXs for the sample period from June 2020 to December 2024. The dashed red line represents mean value of $6.4 Billion in monthly misconduct volume. The y-axis represents the amount of misconduct monthly based on a log scale.

Chapter 5

Insider Trading in Cryptocurrency Markets

5.1 Introduction

This chapter sheds light on insider trading in cryptocurrency markets and the impact of increased regulatory oversight. We analyze cryptocurrency listing announcements and find evidence consistent with individuals using private information to buy coins prior to listing announcements to profit from the price surge that typically follows such announcements. We leverage blockchain data to identify specific wallets (individuals) involved in such trading activity, providing more direct evidence of insider trading than is possible in stock markets.

Insider trading occurs when market participants trade using confidential information not known by the general public that originates within an organization. Insider trading is illegal in most jurisdictions, but until recently, insider trading prosecutions were mainly confined to stock and options markets.[1] However, on July 21, 2022, the U.S. Department

[1]For example, in the United States (U.S.) insider trading is subject to a maximum of 20 years imprisonment and a fine of $5 million (Zornow, Strauber, & Merzel, 2022).

Of Justice (DOJ) charged a former Coinbase employee (Coinbase is the largest U.S. cryptocurrency exchange) his brother, and a friend, with wire fraud in connection with cryptocurrency insider trading that occurred in April 2022, following a Securities and Exchange Commission (SEC) investigation.[2] The trio were accused of using confidential information about coins that were scheduled to be added for listing on the Coinbase exchange. On September 12, 2022, one of the accused plead guilty.[3] The trio were charged with wire fraud, circumventing the need to determine whether the assets traded were securities. Balaji (2024) provides a legal overview on the case and the use of wire fraud laws to police insider trading in cryptocurrency markets.

To our knowledge, this is the first academic study that specifically examines, and provides empirical evidence, of insider trading on token listing announcements in cryptocurrency markets. Leveraging blockchain data, we systematically identify instances of insider trading, including many instances not in the SEC prosecution, estimate the overall prevalence and profits of insider traders, and analyze how insiders respond to increased regulatory oversight. The prosecution provides a validation that the type of activity identified in our analysis (which was largely conducted before the prosecution was made public) is regarded by law enforcement agencies as violating securities law.

In contrast to what is possible in stock markets using public data, we provide direct trade- and account-level evidence using blockchain data. We systematically identify the accounts (wallets) that consistently trade ahead of listing announcements. We rely on wallet clustering techniques to identify 1,210 wallets that, based on their systematic trading around listing announcements, we classify as insider trading wallets. Of these wallets, we identify 49 that belong to what we refer to as "*blatant*" insider traders and 1,161 to "*semi-concealed*" insider traders. Blatant insider traders do not conceal their trading and openly trade on multiple listing announcements from one wallet, leaving an obvious pattern in blockchain data. In contrast, semi-concealed wallets spread their trading across

[2]There is a dual avenue of enforcement and prosecution available to U.S. regulators. The SEC handles civil enforcement of securities law, pursuing penalties such as fines and disgorgement. The DOJ, on the other hand, handles criminal prosecutions, which can result in imprisonment and criminal fines.

[3]Being the first case of insider trading in cryptocurrency markets, it has attracted coverage from major media outlets, including Bloomberg, the Wall Street Journal, the New York Times, Reuters, and CNBC.

numerous wallets to make their insider trading less obvious. Only via additional analysis involving wallet clustering techniques, we are able to link semi-concealed wallets through transactions with each other (e.g., profits from one wallet are passed to a new wallet to fund the next insider trade, with funds later returning to the initial wallet).

We estimate that insider trading occurs ahead of 28-48% of cryptocurrency listings—likely closer to the lower bound of 28% based on the trading of blatant insiders—and that insiders earned at least $30.3 million in trading profits during our sample. Approximately 80% of profits are made by the semi-concealed wallets. We also estimate that insiders in cryptocurrency markets make an average profit of $1.16 million ($83,071 median), which is similar in magnitude to the profits of insiders in U.S. SEC prosecuted cases in traditional markets (average of $1.27 million, median of $95,109 (Kacperczyk & Pagnotta, 2024)). This finding highlights that insiders in cryptocurrency markets can achieve comparable gains to those in traditional financial markets. Our findings support the concerns of Benedetti and Nikbakht (2021) that the value creation from cross-listing cryptocurrencies creates risks to market functions, specifically insider trading.

We also examine cryptocurrency market data and find abnormal return run-ups prior to official exchange listing announcements. This pattern mirrors what has been found in prosecuted cases of insider trading in traditional stock markets (Patel & Putniņš, 2020). We also investigate abnormal trading volumes in the weeks leading up to the listing announcements to assess whether traders might be anticipating listings based on such patterns. However, we find no evidence of abnormal trading volumes, which helps rule out alternative explanations involving early trading signals. Overall, these findings reinforce our blockchain evidence of insider trading occurring in cryptocurrency markets.

Lastly, we exploit the first ever cryptocurrency insider trading prosecution by the U.S. DOJ to investigate the response of insider traders to heightened levels of public and regulatory scrutiny. In contrast with studies of insider trading in stock markets where insiders tend to respond to prosecution risk by decreasing the level of insider trading, we find no substantial decrease in the volume nor frequency of cryptocurrency insider trading activity as regulatory scrutiny increases. Instead, we find that cryptocurrency insider traders increasingly use concealment methods when regulatory scrutiny rises, pushing the activity underground. These results are based on time-series observations and therefore some caution is required in making causal interpretations. To minimize the potential confounding effects we control for a number of factors and limit the time frame to a short window either side of the key shocks. To the best of our knowledge, this is the first study providing evidence on how insiders use concealment in response to increasing levels of regulatory oversight.

An important implication is that cryptocurrency markets cannot be policed as effectively, or by using the same approaches, as stock markets. The pseudo-anonymity of cryptocurrency participants, and an unclear regulatory jurisdiction, draws criminals to cryptocurrencies (e.g., Foley, Karlsen, and Putniņš (2019)), and it also explains the relative lack of sensitivity of cryptocurrency insider traders to increased levels of regulatory scrutiny. The concealment techniques we observe in the data, such as fragmenting trading activity across multiple wallets, rely on pseudo-anonymity. Ultimately, our findings raise the question of whether it is desirable to have regulations prohibit insider trading in cryptocurrency markets if such regulations are difficult to enforce and have a limited deterrence effect. When is no rule better than an ineffective rule?

5.1.1 Related literature

Existing research has explored the market reaction to cryptocurrency listings, with studies by Ante (2019) and Benedetti and Nikbakht (2021) suggesting evidence of informed trading around listing announcements. Additionally, Feng, Wang, and Zhang (2018) demonstrate that Bitcoin markets exhibit informed trading during significant information events. However, to the best of our knowledge, this chapter is the first to directly examine insider trading specifically related to coin listings. Moreover, we provide direct blockchain evidence of trading activity around these announcements.

This chapter is related to prior studies that find evidence of different types of misconduct in cryptocurrency markets. For example, Sokolov (2021) highlights the prevalence of ransomware attacks, showing they impact blockchain fees and congestion. Dhawan and Putniņš (2023), Hamrick et al. (2021), Kamps and Kleinberg (2018), and Li, Shin, and Wang (2021) show that cryptocurrency markets are susceptible to pump-and-dump manipulation. Gandal et al. (2018) and Griffin and Shams (2020) find evidence of fraudulent transactions that inflate Bitcoin prices. Cong et al. (2025) illustrate how ransomware groups that take payments in cryptocurrencies operate. Foley, Karlsen, and Putniņš (2019) quantify the use of cryptocurrencies in illegal online markets (e.g., to purchase drugs, weapons, and illegal pornography) and Cong et al. (2023a) find that cryptocurrencies are used for tax evasion. Cong et al. (2023b) quantify wash trading on unregulated cryptocurrency exchanges. We contribute to this growing literature by examining insider trading, another well-known form of market misconduct that until now has not received attention in the cryptocurrency literature.

Insider trading has been analyzed in traditional financial markets (Acharya & Johnson, 2010; Ahern, 2020; Akey, Grégoire, & Martineau, 2022; Augustin, Brenner, & Subrahmanyam, 2019; Kacperczyk & Pagnotta, 2019; Kallunki et al., 2018; Patel & Putniņš, 2020). Prior studies show that insiders tend to trade off profits against risks of prosecution (Kacperczyk & Pagnotta, 2024) and that inside information flows through strong social ties (Ahern, 2017). Most studies of illegal insider trading rely on prosecuted insider trading cases, which are subject to bias from non-random, incomplete detection and prosecution. The transparent and public nature of blockchains allows us to provide evidence of insider trading that is more direct and not affected by the prosecution bias.

This chapter also contributes to the literature on insider trading regulation. DeMarzo, Fishman, and Hagerty (1998) highlight the benefits of regulating insider trading. Del Guercio, Odders-White, and Ready (2017) find that aggressive enforcement leads to a reduction in illegal activity. Cohen, Malloy, and Pomorski (2012) show that opportunistic insiders are more likely to reduce trading following increased regulatory intensity. Deuskar, Khatri, and Sunder (2025) find that as regulatory scrutiny increases, insiders are more likely to trade in peer stocks. We contribute to this literature by showing how the widely documented deterrent effect of regulation and surveillance is weak, at best, in markets where there are effective means of concealment, such as in cryptocurrency markets. These findings suggest that cryptocurrency markets may need an alternate approach.

Lastly, this chapter provides empirical support to the literature that analyzes the legal challenges surrounding cryptocurrency insider trading (Anderson, 2019; Balaji, 2024; Verstein, 2019). Often, prosecution of illegal insider trading is difficult and requires empirical evidence to support the case. We provide a novel method for identifying potential instances of insider trading, which can help support these cases.

5.2 Data

We hand collect 293 Coinbase listing announcements from September 25, 2018 to January 1, 2024. September 25, 2018, is the date when Coinbase updated their listing process to rapidly list new assets.[4] Listings announced between September 25, 2018, and January 18, 2022 are posted on the Coinbase Medium blog. Sometimes, Coinbase updates the announcements after the original posting, making the exact hour of the updated posting ambiguous. In such cases, we use the Wayback Machine (an internet archive site) to validate the listing date and hour.[5]

Throughout our sample period, the Coinbase listing process has had two main changes. First, on January 18, 2022, Coinbase updated its listing process to instead announce listings on the *CoinbaseAssets* twitter page. Second, on April 11, 2022, to increase transparency about their listing process, Coinbase introduced a road-map of assets to be listed in the next quarter. In the first road-map post, they disclosed a list of 47 tokens that were to be assessed for listing. This did not guarantee that Coinbase would list these tokens.

[4]See https://bit.ly/2xQL7ai for the update to the Coinbase listing process.

[5]There are 30 listings in the sample where the hour of the announcement is still ambiguous. In these 30 listings, we assume the hour of the listing to be 16:00 UTC on the listing date. This is the earliest hour of known listings in the sample period. Our results are robust to excluding these listings.

Often it would take weeks or months for a token from the road-map announcement to be listed, if at all. Closer to the official listing date, there would typically be a follow-up announcement confirming the token's listing. This procedural step became particularly relevant when the release of the road-map sparked an internal investigation into possible insider trading. This turn of events unfolded after a Twitter user flagged an Ethereum wallet that appeared to have purchased several tokens listed in the road-map, all before this information were publicly disclosed. The internal investigation was followed by a U.S. DOJ investigation, which led to an employee, his brother, and a friend pleading guilty to insider trading on token listing announcements.[6]

Table 5.1, Panel A, outlines the sample reduction process for the listing announcements. Starting with the full set of 293 tokens announced to be listed on Coinbase until January 1, 2024, we remove 10 stablecoin and 10 wrapped or staked token listings. For these 20 listings, we do not expect there to be any price impact from the listing, as stablecoins and wrapped/staked tokens would not be profitable for an insider to trade around the listing announcements. We remove 3 tokens that are migrations, forks, or updates from tokens that were previously listed. We split the remaining 270 listings into the 4 tokens that were listed before September 25, 2018, and the 266 tokens that were listed during our sample period. From these 266 tokens, 23 are tokens that were first listed on Coinbase and therefore removed.[7] The final sample is of 252 token listing announcements, with 203 (80.5%) of these being ERC20 tokens. This is important because to identify insider traders' wallets, we focus our analysis on data from the Ethereum blockchain, where the ERC20 token standard is used.[8] The median time between the Coinbase listing announcement and the most recent listing on another CEX is 41 days (average of 114 days).

[6]The Coinbase listing process has subsequently undergone changes, such as providing guidelines on the listing prioritization, as well as in the process standards. For more information, see: https://bit.ly/4c0qkDg.

[7]Having the coin listed on another exchange prior to the Coinbase listing is a necessary condition for insiders to profit. Insiders need to be able to buy the coin somewhere before the Coinbase listing announcement so they can benefit from the price surge following the announcement.

[8]While there are other blockchain networks beyond Ethereum, such as the Solana blockchain, where similar insider trading might occur, our study focuses on the Ethereum blockchain due its dominance in issuing the tokens that subsequently get traded on Centralized Exchanges (CEXs) and the availability of detailed transaction data. Table A5.1 provides a difference in means between the ERC20 tokens that are listed and the sample mean.

Similarly, the average time between the token being tradable on any exchange and the Coinbase listing announcement is 687 days.

TABLE 5.1: Summary statistics of listing announcements

This table presents descriptive statistics of our sample of Coinbase listings. Panel A presents the breakdown of the sample composition. Panel B reports the summary statistics of token listings in the sample period. *Market Cap* is the token's market capitalization on the listing day reported in $ millions. *Trade Volume* is the sum of volume in the 7 days before the listing announcement, reported in $ millions. *Previous Listings* is the number of listings in other cryptocurrency exchanges the token had prior to being listed on Coinbase. *ERC20 Transfers* is the number of token transfers in the 7 days leading up to the listing announcement. *ERC20 Transfer Addresses* is the number of unique addresses that transfer the token in the 7 days leading up to the listing announcement. *ERC20 Transaction Volume* is the dollar value reported in $ millions of on-chain transfers in the seven days leading up to the listing announcement. The *CAR* variable denotes cumulative abnormal returns computed over different time ranges—as indicated by the different subscripts—in hours from the time of the listing announcement. All continuous variables are winsorized at 1% and 99%.

Panel A. Breakdown of listing announcements	
Number of listings	293
Listed before September 25, 2018	4
Stablecoins	10
Wrapped/Staked Tokens	10
Forks/Migrations	3
First announced by Coinbase	23
Final sample	252
ERC20 announcements	203

Panel B. Summary Statistics	Mean	Standard Deviation	p10	Median	p90
Market Cap ($ Millions)	679	1,928	10	123	1,349
Trade Volume ($ Millions)	1,085	4,340	0	22	2,071
Previous Listings	7	5	1	5	15
ERC20 Transfers	4,929	12,447	76	989	11,443
ERC20 Transfer Addresses	2,177	7,261	49	397	4,339
ERC20 Transfer Volume ($ Millions)	206	794	0	13	260
$CAR_{[-300,+100]}$	0.25	0.62	−0.20	0.07	0.95
$CAR_{[-168,-1]}$	0.14	0.31	−0.10	0.06	0.44
$CAR_{[-72,-1]}$	0.10	0.21	−0.08	0.06	0.37
$CAR_{[-24,-1]}$	0.07	0.16	−0.06	0.04	0.21

We use Ethereum blockchain data to identify wallets that are likely to belong to insider traders. To conduct this analysis, we collect ERC20 transfer data for the week leading up and following the listing announcement directly from an Ethereum node.[9] Similarly, we collect address-level Ethereum and token transaction data from Etherscan for Ethereum, ERC20, ERC721, and ERC1155 token transactions.

We also use cryptocurrency market data to assess the impact of insider trading on cryptocurrency markets. To this end, we gather hourly price and volume data for both

[9] An Ethereum node is a computer that runs specific software needed to connect to the Ethereum network of nodes. Together, they verify blocks and transactions.

CEXs and Decentralized Exchanges (DEXs). We get CEX data from CryptoCompare and DEX data from The Graph. We use Uniswap for the DEXs, as it is the market leader in traded volume throughout the sample period.[10] Where available, we use the Crypto Coin Comparison Aggregated Index (CCCAGG) for token price and volume data.[11] If the CCCAGG is not available for a token, we take the price from the cryptocurrency exchange with the largest trading volume prior to the listing announcement. We use the MVIS® CryptoCompare Digital Assets 100 Index (MVDA) as a market index to calculate abnormal returns.[12] MVDA is a market-cap weighted index that tracks the performance of the top 100 digital assets. Hourly returns are winsorized at the 0.5% and 99.5% levels. All variables that are computed per listing are winsorized at 1% and 99% to minimize the influence of outliers. Table 5.1, Panel B, presents the summary statistics for the token listings of our final sample.

5.3 Direct evidence of insider trading

This chapter focuses on insider trading that occurs on cryptocurrency listing announcements. As argued in Benedetti and Nikbakht (2021), cross-listing an asset increases its value. The surge in value can be attributed to greater information production, broadening of the investor base, or a stricter regulatory environment. Consequently, cross-listing also creates scope for insider trading that exploits advance knowledge of listing decisions. Insiders at cryptocurrency exchanges can purchase the token leading up to the announcement and sell it after the listing has increased the value of the token. In the case brought by the U.S. DOJ on Coinbase, it was insider trading on information on listing decisions that was prosecuted. However, insider trading can also occur around other information events specific to a cryptocurrency. In this chapter we focus on the insider trading around listing announcements.

[10] See https://bit.ly/3BhXyOv for market share of DEXs.
[11] See https://bit.ly/3PWftym for CCCAGG methodology.
[12] See https://bit.ly/3PC0uKo for the MVDA fact sheet.

The public and highly transparent nature of blockchains allows for a more direct analysis of the trading preceding listing announcements than what would be possible with public data from stock exchanges. Using blockchain data, we examine whether there are systematic patterns of trading, where specific individuals consistently trade ahead of listing announcements. To profit on a listing announcement, an insider must purchase the token prior to the listing announcement and sell the token after there has been an appreciation from the value creation of the listing. We investigate this pattern of behavior through blockchain transaction data to identify any wallets that could belong to insider traders.

We classify wallets that systematically trade around listing announcements into two distinct groups: *blatant* and *semi-concealed* insider traders. The group of blatant insider traders take no measures to conceal their trading activity. Semi-concealed insider traders employ basic concealment techniques by using multiple wallets to obfuscate their trading activity.

5.3.1 Identifying and classifying insider trading

To identify wallets involved in insider trading, we start by distinguishing a group of wallets that, based on their trading activities, are suspected of participating in insider trading. For insiders to profit from a listing announcement, they must buy the token in the lead-up of its announcement and sell it shortly after once the token has increased in value. In this vein, our group of *suspect wallets* is formed by identifying wallets that buy the token within seven days before its listing announcement, and then dispose of it within the seven days after the announcement.

From the suspect wallet group we then remove Etherscan-labeled wallets (addresses) of CEXs and other smart contracts/Decentralized Finance (DeFi) protocols. Due to how some CEXs pool client assets, a buy before the announcement followed by a sell after the announcement does not necessarily come from the same trader. This is a conservative approach that will likely lead us to underestimate the amount of actual insider trading. We also remove addresses with over 250 unique token transactions or over 10,000 ERC20 transfers. Addresses with many token interactions and ERC20 transactions are likely a result of some form of automated bot rather than an insider trader. These thresholds are chosen to eliminate a small number (≈1%) of highly active addresses that are unlikely to be insider traders based on their trading activity. These accounts are more likely to be bots like those engaging in Maximal Extractable Value (MEV).

Many of the wallets belonging to the suspect wallet group only trade around a single listing announcement with an economically insignificant size. This trading behavior is not systematic enough to allow us to classify these wallets as belonging to an insider trader. Therefore, to confidently classify a suspect wallet as an insider trading wallet we impose restrictions on frequency and volume of trading prior to the listing announcement. We classify a wallet as "*blatant*" insider trader if, using the same wallet, it consistently trades on three or more announcements with over $25,000 in value per announcement. We classify a wallet as "*semi-concealed*" insider trader if it is part of a cluster of related wallets that consistently trades on three or more announcements with a value of over $25,000 per announcement, but that uses different wallet addresses to obscure the pattern of insider trading. Ethereum wallets are inexpensive to create and an individual can have unlimited wallets with minimal costs. An individual can therefore easily conceal their behavior by using several connected wallets. Within this context, Lundblad, Yang, and Zhang (2022) provide evidence suggesting that insiders are more likely to trade with multiple accounts.

The choice of a cutoff of at least three listing announcements with a minimum dollar value of $25,000 may seem somewhat arbitrary.[13] Therefore, we next analyze the sensitivity of our wallet clustering method to distinct specification choices. Figure 5.1 plots the number of tokens traded by insider wallets for different listing announcement frequencies and various dollar amounts. As shown in the figure, the largest variation in the sensitivity analysis is for the blatant wallets, which compared to the semi-concealed group, are more sensitive to the choice of cutoff used on the number of listing announcements.

[13] A cutoff of more than one listing announcement is chosen to guarantee a systematic pattern of behavior. When it comes to the minimum traded value, very large values would be too restrictive and would significantly reduce the number of wallets. Similarly, small trading quantities would introduce too much noise and would dramatically increase the number of wallets under consideration.

For the semi-concealed wallet group, the listing announcement cutoff is less meaningful, with most of the variation for this group coming from the restriction on the traded dollar amount. When testing for various cutoffs and dollar restrictions, we find that choosing a cutoff of at least three listing announcements, and a dollar amount of minimum $25,000 is a good compromise to conservatively identify insider traders.

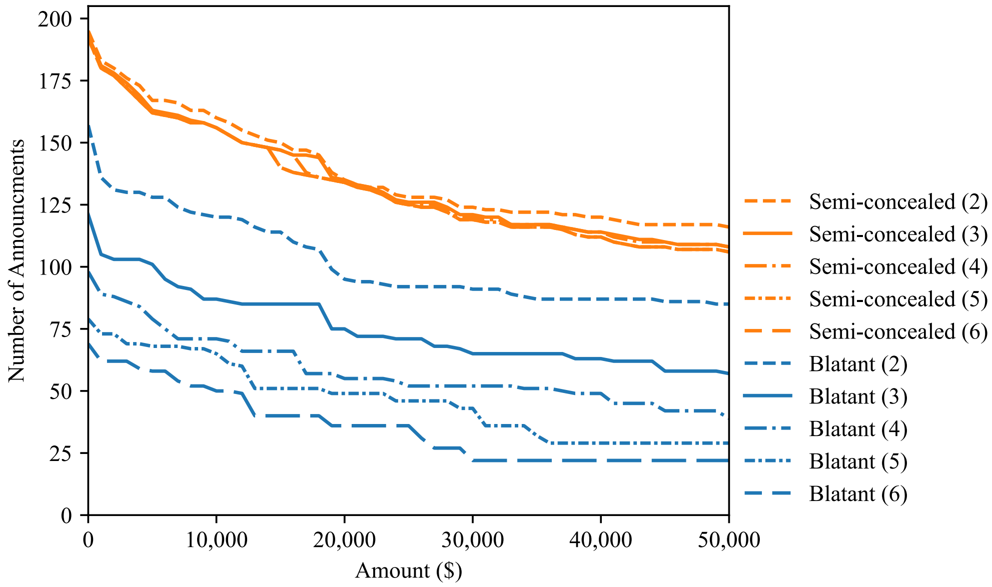

FIGURE 5.1: Wallet clustering sensitivity analysis

This figure presents the sensitivity of the wallet clustering method to different specification choices. The y-axis represents the number of announcements traded by the insider trading groups (blatant and semi-concealed). The x-axis represents different cutoffs of value traded in dollars. The blatant (semi-concealed) insider trading group is represented by the blue (orange) lines, with the variation in line style highlighting the different number of announcements that a wallet (cluster of wallets) needs to trade to be defined as either blatant or semi-concealed insider trading. The number of announcements is shown in the legend surrounded by brackets.

We cluster wallets based on the strongly connected components of a directed graph. Using a directed graph $G(V, E)$, where the nodes V refer to the unique Ethereum wallet addresses that are suspect of insider trading, and E are the interactions between these wallets, we define interactions between wallets to be any Ethereum, ERC20, and/or Non-fungible Tokens (NFT) transfer/transaction to and from the suspect wallets. We include both NFT standards ERC721 and ERC1155. To obtain the clusters of addresses, we find the strongly connected components following Nuutila and Soisalon-Soininen (1994). Prior to clustering, we remove the suspect wallets that do not meet the economic significance cutoff. We do not include nodes in the graph G where the node is part of the list of Etherscan-labeled addresses. The connections from the labeled addresses would not be meaningful for identifying insider trading because they are extremely unlikely to be from an insider trader but from a CEX or DeFi protocol.

A strongly connected component in a graph is one where each node is reachable from every other node in that component. Figure 5.2 provides and example of strongly connected components in a directed graph. The strongly connected components are surrounded by the black dotted line. The example has two strongly connected components 4, 5, 6, 7 and 8, 9, 10 since there is a path from each node to every other node in the same strongly connected component. In contrast, the weakly connected components are surrounded by the orange dotted line. The example has two weakly connected components 1, 2, 3 and 4, 5, 6, 7, 8, 9, 10 since there is a path from each node ignoring the direction of the of the edges.

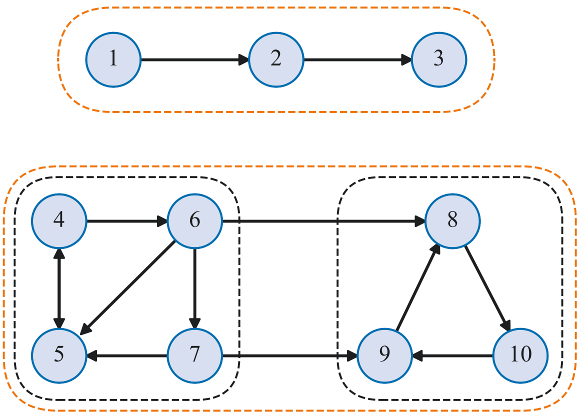

FIGURE 5.2: Example of the strongly and weakly connected components of a directed graph

This figure provides examples of a directed graph, with the strongly and weakly connected components highlighted. The strongly connected components are surrounded by the black dotted line. The directed graph has two strongly connected components {4, 5, 6, 7} and {8, 9, 10} since there is a path from each node to every other node in the same strongly connected component. The weakly connected components are surrounded by the orange dotted line. The directed graph has two weakly connected components {1, 2, 3} and {4, 5, 6, 7, 8, 9, 10} since there is a path from each node, ignoring the direction of the of the edges.

We use the strongly connected components to identify the flows to and from our suspect wallets in order to capture the connection between various suspect wallets. Figure 5.3

provides an example of a cluster of wallets that make a strongly connected component. The blue nodes are the identified suspect wallets, while the orange nodes are other related wallets. The arrows represent the direction of flow between the wallets. This figure highlights the relationship and connectivity among various suspect wallets, with suspect wallets only being two interaction hops from another suspect wallet. By limiting the graph to have edges that are interactions only to and from our suspect wallets, we are capturing the relationship between all suspect addresses that are within two interaction hops of each other.

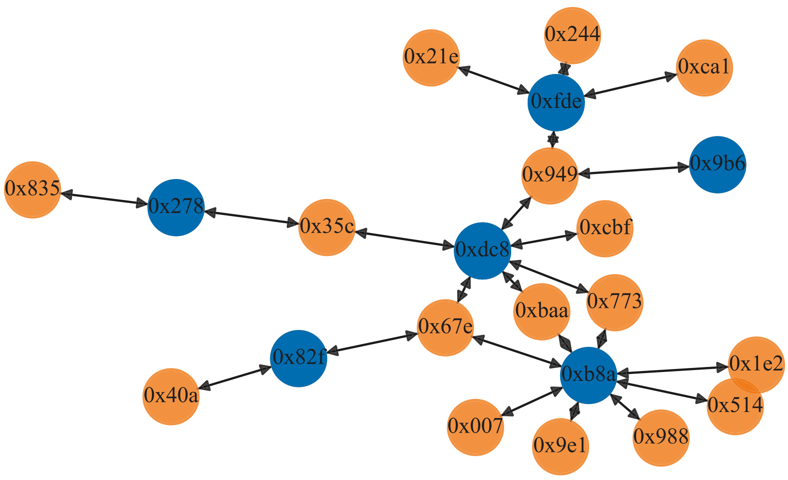

Figure 5.3: Example of wallet cluster

This figure highlights the relationship between a cluster of suspect wallets. The suspect wallets are those wallets that purchase a token in the lead up to a listing announcement and sell the token after it. The suspect wallet nodes are reported in blue, the orange nodes are related wallets, which are strongly connected to the suspect wallets.

Our process to identify insider trading requires the wallet (individual) to have systematically and consistently traded multiple tokens correctly, buying before the listing announcement and selling afterwards, making it implausible to have occurred by chance. This conservative process likely misses cases of insider trading that are one-off or less systematic (type II errors) but reduces the rate of type I errors. For example, consider a fortunate trader who, by chance, purchases a token in the lead up to its listing announcement. Following the price surge, the fortunate trader then sells the token, making a profit. This fortunate trader would be included as part of the suspect wallets but would not be classified as an insider according to our process as the wallet does not display a *systematic* pattern of trading around *multiple* listing announcements.

Consider an alternative scenario where a strategic insider has information about upcoming listing announcements. To obfuscate and conceal her behavior, she transfers funds to a new wallet to trade on each announcement before transferring the profits back. The trading activity of each of the new wallets would appear identical to that of the fortunate trader, where the token was bought prior to the announcement and sold after the listing. But in contrast with the first scenario, there is a systematic pattern of behavior. Being able to detect it and track it is what allows us to distinguish between the two types of traders. We do so by using wallet clustering to identify related wallets that belong to a same individual.

However, there are limitations to clustering using the strongly connected components of a graph. The main limitation is that we cannot capture the relationship of one-way flows between wallets. One-way flows between wallets refers to wallets that transfer funds to the next wallet creating a chain of linked wallets. Using strongly connected components we are unable to identify these wallets because a direct path is not made back into the component.

One solution is to use the weakly connected components of the graph similar to the Union Find algorithm. A challenge with weakly connected components is the noise introduced by including interactions of fraudulent ERC20 tokens, which have transactions to thousands of addresses. These transactions become an edge in the graph representing a connection between nodes. However, the relationship between two addresses based on a connection with a fraudulent token is not likely to be meaningful. A solution to this problem is to apply tighter restrictions on the definition of interactions for the edges of the graph, such as only including reputable tokens, or including a cutoff for the economic significance of a transaction.

We choose to use the strongly connected components approach, as it minimizes the restrictions placed on the interactions between wallets, while still ensuring clusters accurately represent the relationship between our suspect insider trading wallets. We note that, by clustering through the strongly connected components, our estimate of insider trading is likely to approximate a lower bound.[14]

Table 5.2 Panel A, reports descriptive statistics on the overall trading activity of the blatant and semi-concealed insider trading wallets. The blatant (semi-concealed) wallets trade on 71 (118) listing announcements. In the week leading up to the announcement, these wallets accumulate positions totaling $206.5 million, resulting in profits of $30.3 million from trading prior to the listing announcements. The sample median listing volume in the run-up is of $37 million, with insiders contributing around $0.4 million (approximately 1%) to the median volume.

Table 5.2: Descriptive statistics of insider trading wallets

This table presents descriptive statistics on the trading activity of the wallets we classify as blatant and semi-concealed, as well as the total of the two groups. Panel A reports the overall trading activity. Panel B reports a breakdown of the mean and median trading activity by announcement, wallet, and cluster. To estimate the value of profit and outlay, we assume that the insider purchased the token 3 days before the announcement and sold the token 1 day after the listing. All dollar values ($) are in $1,000 units.

Panel A. Total trading activity of insider trading wallets

	Blatant insider traders	Semi-concealed insider traders	Total insider traders
Announcements	71	118	122
Wallets	49	1,161	1,210
Clusters		7	26
Total Position Size ($)	31,102	175,424	206,526
Total Estimated Profit ($)	5,682	24,643	30,325

Panel B. Mean and median trading activity of insider trading wallets

	Blatant insider traders		Semi-concealed insider traders		Total insider traders	
	Mean	Median	Mean	Median	Mean	Median
Announcement						
Position Size ($)	438.1	114.3	1,486.6	374.5	1,692.8	428.4
Estimated Profit ($)	80.0	6.3	208.8	13.5	248.6	15.3
Address						
Position Size ($)	634.7	468.2	151.1	45.7	48.7	48.7
Estimated Profit ($)	116.0	59.3	21.2	–0.2	25.1	–0.1
Cluster						
Position Size ($)			25,060.5	500.0	7,943.3	448.2
Estimated Profit ($)			3,520.5	133.0	1,166.4	83.1

[14]Identification of insider wallets may be subject to potential biases. We therefore take a conservative approach to clustering, based on strongly connected components, and conduct sensitivity analysis on the parameters used in identifying insider traders.

On-chain data is limited by only having information about when a token is transferred to a new account and not when the token is sold. This creates challenges in estimating the profit of insiders that transfer tokens to CEXs in anticipation to selling a token. As it is unclear exactly when a token is sold in these cases, we calculate the profit by assuming that the insider purchases the token three days before the listing announcement and sells one day after the listing.

Table 5.2 Panel B, presents a breakdown of position size and estimated profit by announcement, wallet address, and cluster. The table reports the mean and median values for the blatant and semi-concealed wallet groups, as well as the total when adding both groups. The table shows that, per announcement, semi-concealed wallets have larger position sizes than the blatant wallets. Interestingly, at the address level, the blatant wallets trade up to three times more compared to the semi-concealed group. This result suggests that insiders that conceal their trading are likely to split their trades across multiple wallets to obfuscate their activities. Yet, when looking at the clusters of semi-concealed wallets, we find that their trade size is very similar to that of blatant addresses, showing that while individuals in the two groups differ in their strategies and obfuscation, they trade in a similar size.

The dollar profit earned by insiders in cryptocurrency markets are of a similar magnitude as the profits of insider traders in traditional markets in cases prosecuted by the U.S. SEC. According to Kacperczyk and Pagnotta (2024), prosecuted insiders earned an average of $1.27 million in profits, with a median of $95,109. We find similar profit levels in cryptocurrencies, with an average profit of $1.16 million and a median of $83,071 from insider trading on token listings. This suggests that despite operating in a smaller market, insiders in cryptocurrency markets achieve profit levels comparable to those in traditional markets.

Our findings suggest that there are many instances of insider trading that are not prosecuted in the case by the U.S. DOJ. While the prosecuted case alleges that three individuals purchased and sold at least 25 crypto assets for a profit of more than $1.5

million, our analysis (Table 5.2) identifies 26 distinct clusters of insider trader wallets that profited at least $30.3 million from trading on 122 different listing announcements. Thus, the prosecuted instances reflect about 5-20% of the instances that we identify. Bringing a first of its kind case, such as the prosecuted case of insider trading in cryptocurrency markets, comes with significant challenges. Notably, in classifying the listed tokens as securities, limiting the ability for the SEC to bring the case forward.

Next, we examine the exchanges in which insider traders choose to trade. Table 5.3 reports the interactions of the insider traders' groups with different exchanges (CEXs and DEXs) and DeFi protocols.[15] We use chi-squared tests of proportions to compare the trading of insider traders with a matched sample. The matched sample comprises 10,000 wallets chosen randomly from the list of suspect wallets that transact a token in the week leading up to, and following, the listing announcement. This comparison helps to highlight the differences between typical market behavior and potential insider trading activities.

Overall, the insider trading wallets are significantly more active than the wallets in the matched sample in terms of interactions with CEXs and DeFi protocols. Insider traders interact proportionally more with DEXs, top tier CEXs, NFT exchanges, and DeFi protocols.[16] All but one blatant insider trading wallet are connected to a CEX. Effective Know-Your-Customer (KYC) compliance would imply that these individuals can be identified. We also find that insider traders' wallets make greater use of DEXs (92.7% versus 68.2% in the matched sample), which have no KYC and Anti-money Laundering (AML) requirements. While a large number of the interactions are with Uniswap, at least 75.1% of insider traders interact with other smaller DEXs.

DEXs are also the predominant exchanges used by insider trading wallets to trade on listing announcements. 91.8% (60.4%) of the blatant (semi-concealed) insider trading wallets purchase the tokens from a DEX prior to their listing announcement on Coinbase.

[15] The term "interaction" in our analysis refers to any transaction where a wallet either receives tokens from or sends tokens to a CEX address.

[16] Our analysis indicates that traders have sent/received funds from these exchanges, which means they or someone they know has an account there. This connection does not necessarily imply a direct link to insider information from the exchange but does highlight their interaction with it.

TABLE 5.3: Interactions of insider trading wallets with exchanges

This table presents the proportions of blatant, semi-concealed, and total insider trading wallets that interact with various exchanges and DeFi services. The # column reports the total number of wallets in the group. The % column reports the proportion, in percentage terms, of the total number of wallets in the group. Significance is based on a chi-squared test of proportions comparing the blatant, semi-concealed, and total insider trading wallets to a randomly selected matched sample of other suspect wallets (the "Matched sample") that purchase a token in the week leading up to the listing announcement, and sell the token after the listing announcement. Coinbase, Binance, and FTX are CEXs. Top Tier CEX includes Bitstamp, Bitfinex, Gemini, and Kraken. The low Tier CEX includes all other CEXs. Uniswap, Sushiswap, 1Inch and Other DEX, are DEXs. NFT Exchanges include platforms such as Opensea. DeFi Protocols include well-known DeFi protocols such as Maker and Aave. Tornado Cash is a tumbling service on Ethereum. Significance at the 10%, 5%, and 1% levels is indicated by *, **, and ***, respectively.

	Blatant insider traders		Semi-concealed insider traders		Total insider traders		Matched sample	
	#	%	#	%	#	%	#	%
Coinbase	18	36.7%***	272	23.2%***	290	23.8%***	1,332	13.3%
Binance	38	77.6%***	759	64.9%***	797	65.4%***	5,797	58.0%
FTX	19	38.8%***	241	20.6%***	260	21.3%***	780	7.8%
Top Tier CEX	7	14.3%	250	21.4%***	257	21.1%***	1,014	10.1%
Low Tier CEX	32	65.3%	689	58.9%	721	59.1%	5,780	57.8%
Total CEX	48	98.0%**	997	85.2%	1,045	85.7%	8,525	85.2%
Uniswap	47	95.9%***	1,032	88.2%***	1,079	88.5%***	6,267	62.7%
Sushiswap	32	65.3%***	721	61.6%***	753	61.8%***	1,887	18.9%
1inch	37	75.5%***	627	53.6%***	664	54.5%***	1,371	13.7%
Other DEX	27	55.1%**	889	76.0%***	916	75.1%***	3,760	37.6%
Total DEX	47	95.9%***	1,083	92.6%***	1,130	92.7%***	6,862	68.6%
NFT Exchanges	12	24.5%	617	52.7%***	629	51.6%***	2,192	21.9%
DeFi Protocols	13	26.5%	683	58.4%***	696	57.1%***	2,663	26.6%
Tornado Cash	0	0.0%	31	2.6%***	31	2.5%***	40	0.4%
Total Wallets	49		1,170		1,219		10,000	

Significance: * $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$

The preference for purchasing the tokens from a DEX may be due to the lack of AML and KYC requirements in these exchanges. However, it may also be explained by the accessibility of the tokens on these platforms, where their permissionless nature enables any user of a DEX to list an asset pair for trading.

Semi-concealed insider traders have a significantly larger proportion of interactions with DeFi protocols and NFT exchanges. A small, but significant percentage of these wallets interact with the Tornado Cash smart contract that is known for money laundering and obscuring illicit funds. Tornado Cash pools user funds to break the link between wallet transfers. On August 8, 2022, the U.S. Department of Treasury sanctioned Tornado Cash smart contracts and the associated Ethereum addresses.[17] The semi-concealed wallet interactions with Tornado Cash (2.6%) are highly significant compared to the 0.4% found

[17]See https://bit.ly/3gt7M71 for more information.

for the sample of matched wallets. The increased interactions with Tornado Cash suggest some insiders seek methods of concealment to minimize detection of their trading activity.

5.3.2 Insider trading wallet examples

To illustrate how insider trading occurs, Figure 5.4 presents an example of the trading activity from one of the blatant insider trader wallets. Our analysis suggests that this wallet took positions in four tokens prior to their listing announcement on Coinbase between May 2021 and November 2021. The example wallet also traded in 2 other tokens, one where the tokens where transferred to Binance before the listing announcement and another which is not related to a Coinbase listing announcement but that is purchased around the same time as other tokens. The wallet only purchased tokens from DEXs, avoiding AML and KYC checks, and interacted with CEXs when moving the funds profited to Binance. We estimate that the wallet made profits of just under $1 million by trading on the listing announcements.

The wallet purchased the tokens GTC and MLN before their listing announcement on June 2021. The wallet also purchased RNDR in this period when there was no Coinbase listing announcement.[18] The position that the wallet took in RNDR was much smaller than the positions it took in GTC and MLN. The wallet then proceeded to purchase KEEP in June 2021, prior to its listing announcement. The wallet did not trade around any other listing announcement for about two months, when it purchased RGT prior to its listing announcement in September 2021. In all these cases, the wallet promptly sold the tokens hours after the listing. In the lead up to the RGT listing announcement, the wallet also took a position in SUKU. SUKU would not be listed until a month later, at which moment, but just prior to the actual listing announcement in late October 2021, the wallet transferred the funds to Binance.[19]

[18]RNDR later listed on Coinbase in February 2022.

[19]It is unclear when the individual behind the wallet sold the tokens on Binance. We take a conservative measure of profit at the time of the transfer.

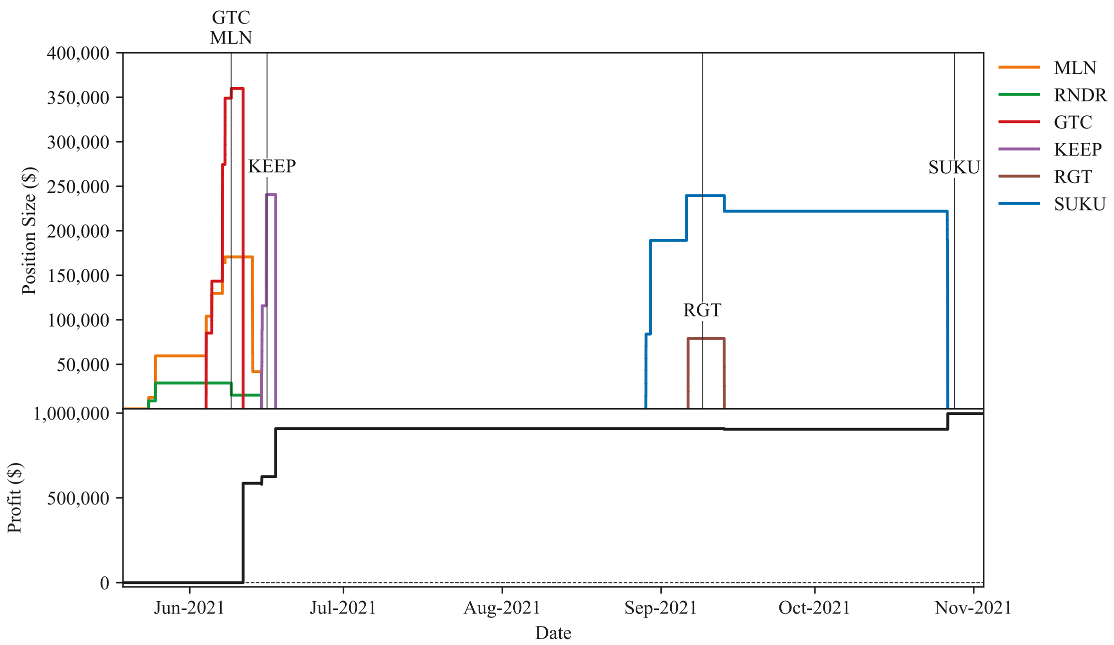

FIGURE 5.4: Example of trading activity of a blatant insider trading wallet

This figure illustrates the positions and profits that a blatant insider trading wallet made by trading ahead of Coinbase listing announcements from May to November 2021. All values are presented in dollars ($). The top panel shows the wallets' position size for trades surrounding the listing dates. The bottom panel presents the total profit that the wallet made on these trades.

The later transaction would not be classified as insider trading by our identification method for two reasons: First, because the token is disposed of before the listing announcement takes place and, therefore, prior to the theoretical price surge. Second, because the purchase occurs in the month before the listing announcement and not in the week leading up to the announcement. While this can be perceived as a limitation of our identification process, it mitigates noise and ensures that only clear cases of insider trading are classified as such.

Figure 5.5 presents an example of the trading activity of a cluster of semi-concealed wallets that we identify as systematically trading prior to Coinbase listing announcements. This cluster of three linked wallets took positions in 11 tokens that were listed between March 2021 to November 2021. The wallets often transferred assets to Binance, but much of their trading occurred on a DEX, avoiding AML and KYC checks. We estimate these wallets profited around $1.6 million in total.

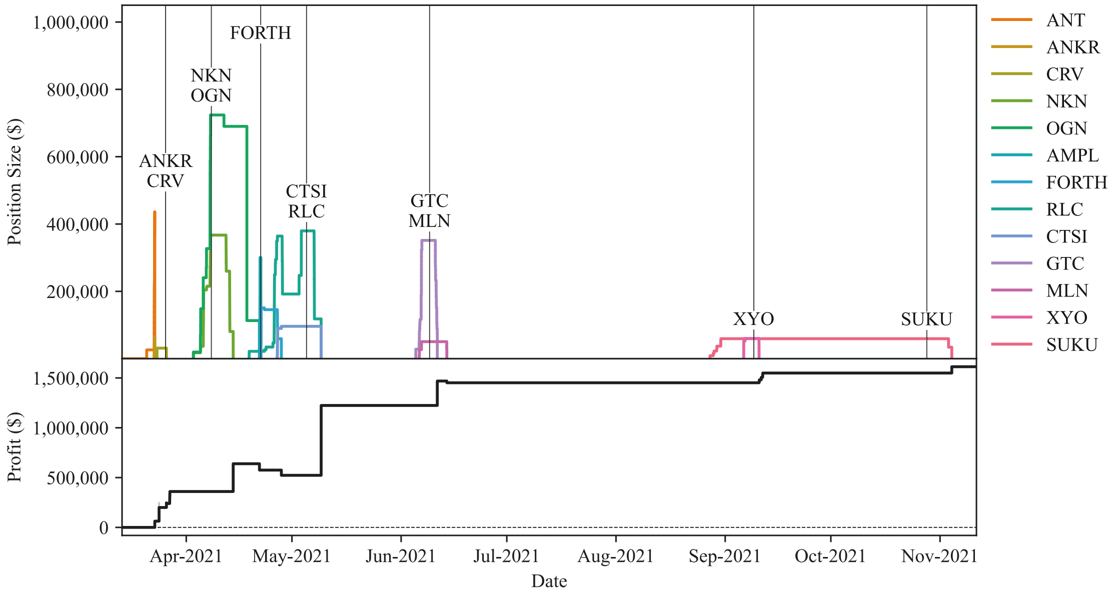

FIGURE 5.5: Example of trading activity of a single cluster of semi-concealed insider trading wallets

This figure illustrates the positions and profits that a set of linked wallets made by trading ahead of Coinbase listing announcements from March to November 2021. The three wallets in this example are linked by direct transfers to the other wallets. All values are presented in dollars ($). The top panel shows the wallets' position size for trades surrounding the listing dates. The bottom panel presents the total profit that the wallets made on these trades.

The first wallet, Wallet 1, took positions in tokens ANKR, NKN, OGN, RLC, CTSI, and FORTH prior to their listing announcements. With the purchase of FORTH, the individual behind the wallet also purchased AMPL. FORTH is the governance token to Ampleforth, which is a decentralized protocol on Ethereum. Ampleforth has both a governance token (FORTH) and the protocol token AMPL. The first listing of FORTH happened on both Uniswap and Coinbase simultaneously, however, AMPL had been actively traded on other exchanges. The individual bought the wrong token prior to the announcement by purchasing the most similar asset. On announcement, the individual rushed to purchase FORTH. This mistake is consistent with the listed assets being tipped by an insider to an external party. Wallet 1 transferred ETH profits to Binance and to another CEX, OKX.

While trading on the ANKR listing announcement, Wallet 1 received ANT from Wallet 2. Wallet 2 only traded on the listing announcements of ANT and ANKR. Similarly, Wallet 1 sent to Wallet 3 roughly 1 ETH over two transactions and then Wallet 3 took positions in CRV, RLC, FORTH, GTC, MLN, SUKU, XYO prior to the Coinbase listing announcements. Similar to the blatant insider trading wallet example, SUKU was purchased in the lead up to a previous listing announcement and it was held until after its subsequent listing announcement. The similar pattern of behavior is consistent with insiders having access to the same information on upcoming token listings. As in the previous example, our identification method does not classify as insider trading the purchase of SUKU given that it occurred over a month before the listing announcement rather than in the week leading up to the announcement. The example wallets would sometimes continue to purchase the token after the listing announcement. This may be to trade on the subsequent price run-up when the listing occurs, or to obfuscate their trading behavior.

5.3.3 Market evidence of insider trading

To support our findings based on blockchain evidence, we next examine price movements in token listings for further evidence of insider trading. We begin by analyzing the price movements in coin listings from 300 hours before the listing announcement, up to 100 hours after the announcement, looking for any abnormal trading patterns that could be potential evidence of insider trading. Figure 5.6 presents the mean cumulative abnormal returns (in excess of market returns) for the 252 newly-listed Coinbase tokens.[20] Market returns are based on the MVDA market index. MVDA is a market-cap weighted index that tracks the top 100 digital assets.[21] Figure 5.6 shows a clear run-up pattern prior to the listing announcement starting approximately 200 hours before the announcement. On average, there is a 19.9% run-up in cumulative abnormal returns prior to the announcement. On the announcement day, the tokens have an average return of 6.1% because of new information entering the market and traders reacting to the news. The run-up pattern we observe is consistent with the run-ups detected in prosecuted cases of insider trading in stock markets (Del Guercio, Odders-White, & Ready, 2017; Patel & Putniņš, 2020).

[20]For robustness, Figure A5.1 in the Appendix presents Figure 5.6 with a longer event window of 60 days pre and 20 days post the Coinbase listing announcement.

[21]For robustness, in Figure A5.2 in the Appendix we use instead an index that focuses on small market capitalization tokens. Our findings indicate a consistent pattern of return run-ups, including excess returns, which align with the original results obtained when using the broader MVDA index.

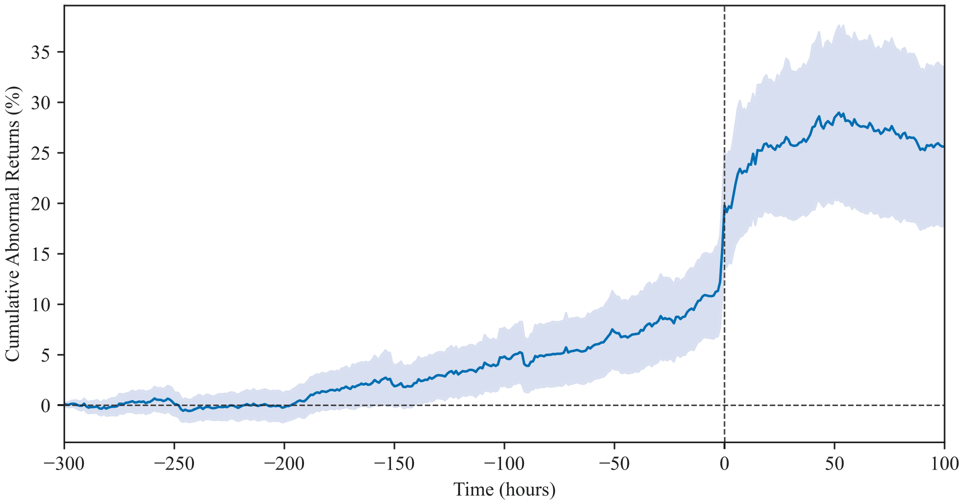

FIGURE 5.6: Cumulative abnormal returns for token listing announcements

This figure plots the mean cumulative abnormal returns for the tokens listed on Coinbase before and after the listing announcements at t = 0. The cumulative returns are measured in hours over a period of 300 hours before the listing announcement, up until 100 hours after it. The shaded region represents the 5% to 95% confidence bounds. The sample contains 252 tokens that were listed on Coinbase between September 2018 to January 2024.

Next, we compare the tokens traded by the blatant and semi-concealed insider trading groups to a null group—the group of tokens that do not fall into the blatant and semi-concealed insider trading groups, respectively. The purpose is to investigate whether the run-up is universal, in which case it may indicate that it is driven by factors other than insider trading, or whether it is specific to the tokens targeted by our identified insider trading group. Figure 5.7 examines the excess cumulative abnormal returns for the tokens traded by the blatant and semi-concealed groups. In analyzing the run-up *difference* of the two groups, we effectively control for factors that could cause a general tendency for run-ups ahead of listings.[22] As evidenced in Figure 5.7, the tokens traded by our identified insider trading groups have an average excess cumulative abnormal return prior to the listing announcement of 21.1% for the blatant, and of 19.9% for the semi-concealed group. The run-up pattern starts approximately 200 hours before the announcement for

[22]For robustness, in Figure A5.3 in the Appendix we present Figure 5.7 instead using a market index that focuses on small market capitalization tokens. The results are consistent with those obtained in the main analysis.

both groups. On the day of the announcement, they experience an excess return of -9.2% (-5.9%).

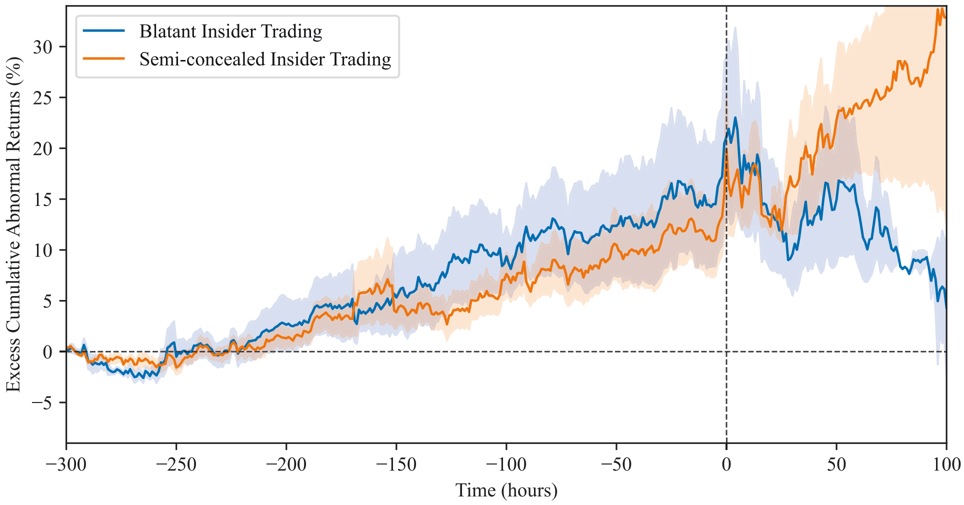

FIGURE 5.7: Excess cumulative abnormal returns of tokens traded by the blatant and semi-concealed insider trading wallets

This figure plots the mean cumulative abnormal returns for the tokens traded by the blatant and semi-concealed insider trading wallets in excess of the cumulative abnormal returns for the null group (token listings that are not traded by the blatant or semi-concealed groups), before and after listing announcements at t = 0. The excess cumulative returns are measured in hours over a period of 300 hours before the listing announcement, up until 100 hours after it. The blue region represents the tokens traded by the blatant wallets, while the orange region depicts the tokens traded by the semi-concealed wallets. The shaded regions display the 5% to 95% confidence bounds. The sample includes 71 tokens traded by blatant wallets and 118 traded by semi-concealed wallets.

To test the statistical significance of the run-ups seen in the plots, we regress the cumulative abnormal returns (CAR) over ranges [-24, -1], [-72, -1] and [-168, -1] hours relative to the listing announcement at time 0 (i.e., 1, 3, and 7-days prior to the listing announcements) on dummy variables for the blatant, semi-concealed, and total insider trading groups:

$$CAR_{[t_1,t_2],i} = \alpha + \beta_1 ITG_i + \gamma Controls_i + \epsilon_i, \tag{5.1}$$

where ITG is a dummy variable that equals 1 if the token is part of the blatant, semi-concealed, or total insider trading groups, or 0 otherwise. *Controls* is a group of control variables including market capitalization, the token's pre-listing announcement volume, the coin's number of previous listings, the coin's number of on-chain transfers, and the coin's on-chain transfer volume in the week leading up to the announcement. Following Liu, Tsyvinski, and Wu (2022) we control for size (market capitalization) and volume which capture the cross section of cryptocurrency returns. Benedetti and Nikbakht (2021) suggest cross listing can lead to value creation for coins, so we control for cross listings by including a coin's number of previous listings. Furthermore, Benedetti and Nikbakht (2021) suggest that network participation increases in the lead up to cross listing events. The increased activity could lead to the anticipation of listings. We therefore control for a coin's on-chain transfers and on-chain volume. All continuous variables are winsorized at 1% and 99%.

Table 5.4 reports the results of the CAR regressions. Regressions 1, 4, and 7 examine the relation between the abnormal return run-up and the tokens traded by the blatant insider trading group. The tokens traded by the blatant group have significant excess returns of 16.4%, 14.8% and 10.9% in the 7, 3, and 1 days before the listing announcement. The tokens traded by the blatant insider trading group see the largest excess returns across all trade periods, consistent with a disregard for concealing their trading activity. Regressions 2, 5, and 8 examine the relation between the abnormal return run-up and the tokens traded by the semi-concealed insider trading group. The tokens traded by the semi-concealed group see lower but more significant excess returns of 14.8%, 12.5% and 9.6% in the 7-, 3-, and 1-day before the listing announcement. Regressions 3, 6, and 9 examine the relation between the abnormal return run-up and the tokens traded by the total insider trading group (the blatant and semi-concealed groups combined). The tokens traded by the total insider trading group see excess returns of 12.7%, 11.4% and 8.8% in the 7-, 3-, and 1-day before the listing announcement. Our findings suggest that the large excess run-up, both in relative and absolute terms, is economically meaningful.

The findings based on market evidence of insider trading provide further support to our results when using blockchain data. The run-up patterns we observe are consistent with the run-ups detected in prosecuted cases of insider trading in stock markets. Our results

TABLE 5.4: Cumulative abnormal return regressions

This table presents estimates for OLS regressions testing if the run-up in abnormal returns can be explained by the likely insider trading groups. The dependent variable in columns 1, 2 and 3 is the abnormal return over a window 7 days prior, up to one hour before, the listing announcement. Similarly, the dependent variable in columns 4, 5 and 6 is the abnormal return over a window 3 days prior, up to one hour before, the listing announcement. The dependent variable in columns 7, 8 and 9 is the abnormal return over a window 1 day prior, up to one hour before, the listing announcement. The key independent variables are the dummy variables fot the insider trading groups *Blatant*, *SemiConcealed*, and $Total$. The dummy variables are equal to 1 if the token listing announcement is traded by the respective insider trading wallets. $PreviousListings$ is the number of listings in other exchanges the token had prior to it being listed on Coinbase. $MarketCap$ is the log of the token's market capitalization on the listing day. $TradeVolume$ is the log total trade volume in the 7 days before the listing announcement. $Transfers$ is the log number of token transfers in the 7 days leading up the the listing announcement. $TransactionVolume$ is the log dollar value of on-chain transfers in the 7 days leading up the the listing announcement. All continuous variables are winsorized at 1% and 99%. Standard errors are clustered by announcement day. t-statistics are reported in parentheses. Significance at the 10%, 5%, and 1% levels is indicated by *, **, and ***, respectively.

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Dependent variable:		$CAR_{[-168,-1]}$			$CAR_{[-72,-1]}$			$CAR_{[-24,-1]}$	
Blatant	0.164*			0.148***			0.109**		
	(1.94)			(2.73)			(2.20)		
SemiConcealed		0.148**			0.125***			0.096**	
		(2.58)			(3.04)			(2.62)	
$Total$			0.127**			0.114***			0.088**
			(2.26)			(2.80)			(2.50)
$PreviousListings$	−0.051	−0.078**	−0.075**	0.016	−0.007	−0.006	0.048***	0.031	0.032*
	(−1.50)	(−2.18)	(−2.11)	(0.65)	(−0.28)	(−0.22)	(2.91)	(1.65)	(1.72)
$MarketCap$	−0.034**	−0.035**	−0.036**	−0.034***	−0.035**	−0.035**	−0.036***	−0.037***	−0.037***
	(−2.37)	(−2.30)	(−2.28)	(−2.77)	(−2.61)	(−2.58)	(−4.03)	(−3.80)	(−3.75)
$TradeVolume$	0.010	0.012	0.011	0.003	0.005	0.005	−0.003	−0.002	−0.002
	(0.81)	(0.96)	(0.95)	(0.39)	(0.60)	(0.59)	(−0.58)	(−0.33)	(−0.33)
$Transfers$	0.002	0.002	0.004	−0.008	−0.007	−0.006	−0.016*	−0.016*	−0.015
	(0.17)	(0.18)	(0.32)	(−0.84)	(−0.79)	(−0.64)	(−1.80)	(−1.77)	(−1.65)
$TransferVolume$	0.002	0.002	0.002	0.004	0.004	0.004	0.005	0.005	0.005
	(0.38)	(0.32)	(0.38)	(1.07)	(1.11)	(1.10)	(1.29)	(1.31)	(1.28)
Intercept	0.606***	0.600***	0.604***	0.608***	0.605***	0.606***	0.717***	0.713***	0.714***
	(3.29)	(3.02)	(2.89)	(3.25)	(2.93)	(2.86)	(5.12)	(4.70)	(4.59)
Observations	203	203	203	203	203	203	203	203	203
Adjusted R^2	11.0%	10.5%	9.4%	18.1%	16.5%	15.6%	28.8%	27.9%	27.3%

Significance: * $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$

have therefore implications for using return run-up patterns to predict insider trading. It shows, without the selection bias of prosecuted cases, that large run-up patterns exist when insider trading activity is present.

While the evidence is consistent with the identified wallets having known of the upcoming listing announcements, we consider whether, instead of using inside information, these wallets could have inferred the upcoming listings from abnormal trading prior to the listings. This seems unlikely for two reasons: First, trade- or price-based signals are noisy and would therefore result in a substantially less consistent pattern of anticipating listings than the one we uncover. Second, we test for abnormal volume prior to the trades belonging to the identified insider traders that could have alerted them of the upcoming listings. To do this, we compare the volume of the tokens traded by the blatant and semi-concealed wallets with the volume traded on other listing announcements for a range of weeks prior to the announcement events.

Table 5.5 presents the results of abnormal trading volume in DEXs and CEXs, for the blatant and semi-concealed insider trading wallets, as well as for the total of both groups. We find a significant difference in means for trading on DEXs in the listing week (week 0) for all groups. In the week leading up to the listing announcement (week -1) there is a significant difference in the CEX volume for the blatant insider trading wallets. It is likely that the abnormal volume can be attributed to the actions of insider traders. Other than that, we find no other significant differences in the abnormal volume measure of the two samples for the weeks further away from the listing announcement (weeks -4 to -2). This finding suggests an absence of abnormal trading volume before that of the suspected insider trader and therefore helps rule out alternative explanations where information is drawn from early trading signals.

TABLE 5.5: Test of abnormal trading volume by insider trading wallets prior to listing announcements

This table reports the test statistics from t-tests of abnormal volume (AV) prior to the likely insider trading event. The test compares differences in the mean AV for the tokens traded by the blatant, semi-concealed, and total insider trading wallets, relative to all other tokens announced for listing by Coinbase. For each of the four weeks leading up to the announcement (columns -4 to -1), and the week of the announcement (column 0), we measure AV as the traded volume on that week minus the traded volume in a benchmark week, being the 5th week prior to the announcement. Test statistics for both CEX and DEX volumes are reported as separate rows. Significance at the 10%, 5%, and 1% levels is indicated by *, **, and ***, respectively.

	Weeks				
	(-4)	(-3)	(-2)	(-1)	(0)
Blatant insider traders					
CEX Volume	0.69	1.03	0.83	2.04^{**}	1.68^{*}
DEX Volume	0.33	0.51	−0.09	1.63	2.32^{**}
Semi-concealed insider traders					
CEX Volume	−0.49	0.05	0.33	0.41	2.69^{***}
DEX Volume	0.87	0.24	−0.53	1.13	0.79
All insider traders					
CEX Volume	−0.49	0.29	0.24	0.39	3.11^{***}
DEX Volume	−0.26	−0.35	−0.53	1.02	2.16^{**}

Significance: * $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$

5.3.4 How prevalent is insider trading?

We estimate that insider trading occurs in 28-48% of listings in our sample.[23] The upper and lower bounds are estimated by taking the occurrence of insider trading by the blatant and semi-concealed insider trading groups. To confirm our estimation based on actual blockchain evidence, we quantify the excess price run-ups. In measuring the excess price run-ups, we control for the abnormal return patterns that are not related to insider trading by contrasting two trading periods. The first period is before the listing announcement, where we expect to see run-ups related to insider trading. The second period is unrelated to the announcement and therefore we should not expect to see run-ups. Specifically, we compare the $CAR_{[-300,-1]}$ to the $CAR_{[-600,-300]}$ to assess the number of listings with abnormal run-up patterns consistent with insider trading in the $CAR_{[-300,-1]}$ range. We take the difference in the number of listings that have an abnormal return greater than 10% between $CAR_{[-300,-1]}$ and the $CAR_{[-600,-300]}$. The level of excess returns greater than 10% is chosen based on an excess return that is consistent with the median return for the $CAR_{[-300,+100]}$ period.

This process results in 53 listings (21% of the sample) that have strong excess run-ups consistent with insider trading, which is close to the lower bound estimation based on blockchain data.[24] Because the blockchain evidence is more direct than that from price run-ups, the lower bound of insider trading is likely closer to the 28% estimate.

The findings suggest that insider trading in cryptocurrency markets is more prevalent than in stock markets. We estimate that insider trading occurs in 28-48% of listings in our sample. For comparison, in stock markets insider trading is estimated to occur between one in 20 earnings announcements or one in five merger and acquisition events (Patel & Putniņš, 2020).

[23] An industry report by Solidus Labs (2023) supports our findings suggesting that insider trading occurs in 56% of token listing announcements.

[24] The estimation is robust to different levels of excess returns. At a 5% and 20% level, the lower bound is instead 17% and 15% respectively.

5.4 How do insider traders respond to regulatory scrutiny?

We investigate how regulatory scrutiny impacts an insider's willingness to trade on private information, as well as the use of concealment methods. To this end, we identify two exogenous shocks. The first shock is the tip-off from a prominent twitter user on April 12, 2022, which publicly suggested that insider trading was occurring. This tip-off sparked investigations internally at Coinbase and by the DOJ. We hypothesize that, after the tip-off, insiders are likely to become more aware of the risk of being caught and the potential punishability of their actions.

The second shock that we analyze is the DOJ indictment of the former Coinbase employee for alleged insider trading on July 21, 2022, which signals to insiders that the risks of prosecution are real and have legal consequences. Given that this is the first case of insider trading in cryptocurrency markets to be prosecuted, prior to the indictment, an insider could only make assumptions about the risks of their activities. There was a perception among many market participants that insider trading laws did not apply in cryptocurrency markets. That changed with the indictment, as insiders gained a clear understanding of how regulators view their actions and the risks of prosecution.

Figure 5.8 illustrates how the prevalence of insider trading changes over time. The figure presents the 5-announcement moving average of a dummy variable that is equal to 1 if the listed token is traded by the blatant (or the semi-concealed) wallets and 0 otherwise.[25] The figure suggests that, around the tip-off shock and up until the end of the sample period, instances of blatant insider trading cease to occur.

Interestingly, we see the rates of insider trading by semi-concealed wallets remain mostly unchanged as regulatory scrutiny increases with the shocks of the tip-off and the

[25]For robustness, Figure A5.4 presents the 3- and 7-announcement moving average for the prevalence of insider trading similar to Figure 5.8. The level of blatant insider trading appears to decline slightly before the tip-off. This decline could be attributed to several factors, including anticipation of the upcoming investigation. The main takeaway from our analysis is that the levels of blatant insider trading does not increase after the tip-off and DOJ indictment.

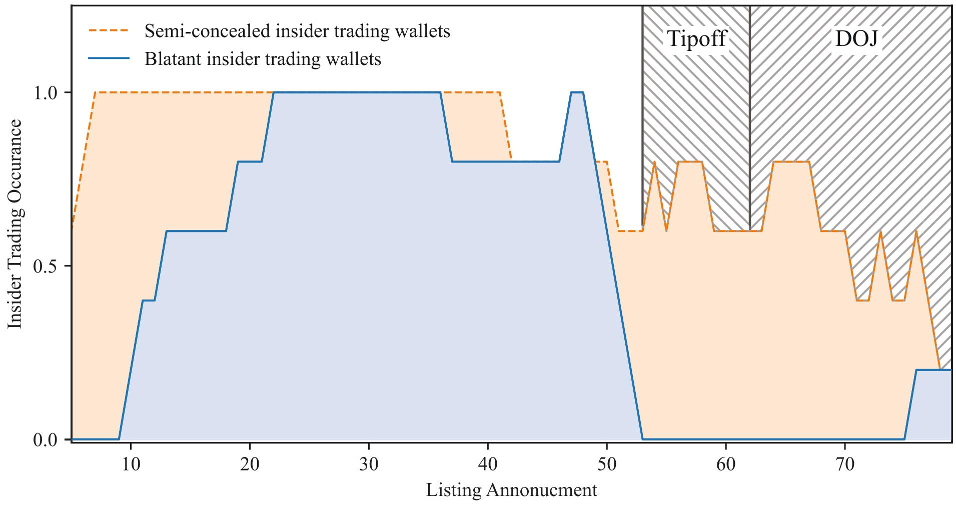

FIGURE 5.8: Prevalence of insider trading activity over time

This figure presents the 5-announcement moving average of a dummy variable that is equal to 1 if the listing announcement was traded by a blatant insider trading wallet (blue area), or only traded by a semi-concealed insider trading wallet (orange area). The downward sloping background region represents the listing announcements that were made after the tip-off. The upward sloping background region represents the listing announcements that were made after the DOJ indictment.

DOJ indictment. As a result, the total amount of insider trading activity does not substantially decline or disappear over time, but it only shows a slight decrease towards the end of the sample period. The results show that rather than ceasing their illegal actions, insiders choose obfuscation methods to continue profiting from trading on confidential information.

To test the significance of the pattern of trading presented in Figure 5.8 we estimate the following regression model to test the impact of the shocks (the tip-off and the DOJ indictment) on the prevalence of insider trading:

$$ITM_i = \alpha + \beta_1 Tipoff_i + \beta_2 DOJ_i + \beta_3 time_i + \beta_4 time_i^2 + \gamma Controls_i + \epsilon_i, \qquad (5.2)$$

where i is an index for each token announcement and ITM is one of the following measures of insider trading: (i) trading by insider trading wallets, (ii) abnormal return patterns, or (iii) volume traded by the insider trading wallets. $Tipoff$ is a dummy variable that equals 1 if the token is listed after the tip-off on twitter on April 12, 2022 and 0 otherwise. DOJ

is a dummy variable that equals 1 if the token is listed after the DOJ indictment on July 21, 2022, and 0 otherwise. $time$ and $time^2$ are clock time variables (days since the start of the sample) that control for time-series effects. Including these time variables effectively provides a regression discontinuity design framework that captures the discrete changes in levels of insider trading after each shock stripping out the effects of linear and quadratic time trends. $Controls$ is a group of control variables from Table 5.4 including market capitalization, the token's pre-listing announcement volume, the coin's number of previous listings, the coin's number of on-chain transfers, and the coin's on-chain transfer volume in the week leading up to the announcement. All continuous variables are winsorized at 1% and 99%.

Table 5.6 reports the results from Equation 5.2 estimated separately for insider trading by the blatant, semi-concealed, and total wallet groups. Regressions 1 and 2 focus on the occurrence of insider trading by the blatant insider trading group. Our findings suggest that the occurrence of blatant insider trading is reduced after the shock from the twitter tip-off. After the public allegations in the tip-off period, insiders may anticipate the future prosecution. This would explain why we see no effect from the DOJ indictment directly. These results are consistent with Figure 5.8, and are in line with our hypothesis that, after the tip-off, insiders become aware they may get caught and worry about the legal risks associated with their actions.

Regressions 3 and 4 focus on the occurrence of insider trading by the semi-concealed insider trading group. In contrast to the results we find from the blatant insider trading group, we do not see a significant decrease in insider trading in both the tip-off and DOJ indictment periods. Similarly, in Regressions 5 and 6, we also do not find a significant decrease in insider trading from the total insider trading group (the blatant and semi-concealed groups combined) in both the tip-off and DOJ indictment periods. These results suggest insiders shift their trading behavior towards using concealment measures to still profit from their private information.

Table 5.6: Regressions to assess the impact of regulatory shocks on the prevalence of insider trading

This table presents estimates from OLS regressions that examine the impact on the prevalence of insider trading of the twitter tip-off and the DOJ indictment. The dependent variables for *Blatant* (Regressions 1 and 2), *SemiConcealed* (Regressions 3 and 4), and *Total* (Regressions 5 and 6) are dummy variables that are equal to 1 if the token listing announcement is traded by the respective insider trading wallets. The key independent variables are $Tipoff$ and DOJ. $Tipoff$ is a dummy variable that is 1 after the twitter tip-off on April 12, 2022, and 0 otherwise. DOJ is a dummy variable that is 1 after the DOJ indictment on July 21, 2022, and 0 otherwise. *time* is days from the first listing in the sample. $time^2$ is adjusted by 10e6. *PreviousListings* is the number of listings in other exchanges the token had prior to it being listed on Coinbase. *MarketCap* is the log of the token's market capitalization on the listing day. $TradeVolume$ is the log total trade volume in the 7 days before the listing announcement. $Transfers$ is the log number of token transfers in the 7 days leading up to the listing announcement. $TransactionVolume$ is the log dollar value of on-chain transfers in the 7 days leading up to the listing announcement. All continuous variables are winsorized at 1% and 99%. Standard errors are clustered by announcement day. t-statistics are reported in parentheses. Significance at the 10%, 5%, and 1% levels is indicated by *, **, and ***, respectively.

	(1)	(2)	(3)	(4)	(5)	(6)
Dependent variable:	*Blatant*		*SemiConcealed*		*Total*	
$Tipoff$	−0.381***	−0.088	−0.074	0.264**	−0.107	0.237*
	(−3.30)	(−0.65)	(−0.54)	(2.10)	(−0.75)	(1.69)
DOJ	0.168**	0.465***	0.120	0.449***	0.127	0.466***
	(2.02)	(3.99)	(0.83)	(3.14)	(0.89)	(3.27)
PreviousListings	−0.051	−0.069	0.044	0.020	0.043	0.020
	(−0.91)	(−1.43)	(0.75)	(0.40)	(0.73)	(0.39)
MarketCap	−0.043	−0.025	−0.032	−0.008	−0.032	−0.009
	(−1.33)	(−1.14)	(−1.02)	(−0.43)	(−1.06)	(−0.50)
$TradeVolume$	0.004	0.001	−0.003	−0.007	−0.004	−0.007
	(0.43)	(0.06)	(−0.30)	(−0.65)	(−0.32)	(−0.67)
$Transfers$	0.034	−0.008	0.055	0.004	0.048	−0.004
	(1.03)	(−0.34)	(1.38)	(0.12)	(1.12)	(−0.10)
$TransferVolume$	0.032**	0.036***	0.032**	0.038**	0.036**	0.041**
	(2.39)	(2.75)	(2.09)	(2.36)	(2.22)	(2.49)
time		0.002***		0.002***		0.002***
		(4.02)		(3.62)		(3.78)
$time^2$		−1.297***		−1.401***		−1.458***
		(−4.00)		(−4.81)		(−4.82)
Intercept	0.440	0.143	0.238	−0.069	0.280	−0.045
	(0.73)	(0.36)	(0.44)	(−0.20)	(0.51)	(−0.13)
Observations	203	203	203	203	203	203
Adjusted R^2	15.7%	27.3%	15.8%	29.3%	15.6%	29.8%

Significance: * $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$

Regressions 2, 4, and 6 show that the occurrence of insider trading is significantly decreasing over time when variables for time (*time* and $time^2$) are included. These results suggest that the levels of insider trading are decreasing regardless of the shocks from the tip-off and DOJ indictment.

Overall, our findings highlight a significant challenge for law enforcement agencies. Instead of diminishing the level of insider trading, as seen in stock markets, regulatory scrutiny in the cryptocurrency space tends to drive such activities "underground", with insiders increasingly relying on concealment techniques. To ensure the robustness of our findings, we also examine the CAR run-ups as an alternative measure of insider trading. Prior research has established CAR run-ups as a reliable indirect indicator of insider trading (Del Guercio, Odders-White, & Ready, 2017; Patel & Putniņš, 2020). Moreover, consistent with this literature, we also found evidence in subsection 5.3.3 supporting the use of CAR run-ups as indirect evidence of insider trading in cryptocurrency markets.

Table 5.7 reports the regression results from Equation 5.2 focusing on the CAR run-ups as an indirect measure of insider trading. Regressions 1 and 2 focus on the 7-day CAR prior to the listing, Regressions 3 and 4 focus on 3-day CAR prior to the listing, and Regressions 5 and 6 focus on the 1-day CAR prior to the listing. Across all CAR ranges, we find that after the tip-off period the CAR in the lead up to the listing announcement is decreasing, consistent with our findings that insider trading is decreasing from Table 5.6. Including the time variables ($time$ and $time^2$), again suggests that the run-ups are decreasing overtime. However, this trend is not significant.

Next, we consider the volume traded by the different insider trading groups rather than the frequency of insider trading. Supporting our results from Table 5.6, we find that there is a significant reduction in the levels of blatant insider trading in response to the tip-off shock. However, there is no significant reduction in the levels of semi-concealed and total insider trading in response to both the tip-off and DOJ indictment. The time variables ($time$ and $time^2$) highlight the significant negative time trend, suggesting that insider trading is decreasing. Regardless, our results highlight a major regulatory challenge as insiders turn to concealment methods in response to increased regulatory scrutiny.

Cohen, Malloy, and Pomorski (2012) note that opportunistic insiders change their behavior in response to increased regulatory intensity. There is also evidence to suggest that insiders substitute trades in peer stocks and exchange traded funds to circumvent detection (Deuskar, Khatri, & Sunder, 2025; Eglīte et al., 2022). Our findings also suggest that insiders opportunistically choose to conceal their trading activity. Prior to the DOJ indictment, there was no reason for an insider to conceal their trading activity if they believed their actions would not be prosecuted. However, as it becomes clear that their

TABLE 5.7: Regressions to assess the impact of regulatory shocks on abnormal return patterns

This table presents estimates from OLS regressions that examine the impact of the twitter tip-off and the DOJ indictment on indirect measures of insider trading (abnormal return patterns). The dependent variable in column 1 is the abnormal return over a window 7 days prior, up to one hour before, the listing announcement. Similarly, the dependent variable in column 2 is the abnormal return over a window 3 days prior, up to one hour before, the listing announcement. The dependent variable in column 3 is the abnormal return over a window 1 day prior, up to one hour before, the listing announcement. The key independent variables are $Tipoff$ and DOJ. $Tipoff$ is a dummy variable that is 1 after the twitter tip-off on April 12, 2022, and 0 otherwise. DOJ is a dummy variable that is 1 after the DOJ indictment on July 21, 2022, and 0 otherwise. $time$ is days from the first listing in the sample. $time^2$ is adjusted by 10e6. $PreviousListings$ is the number of listings in other exchanges the token had prior to it being listed on Coinbase. $MarketCap$ is the log of the token's market capitalization on the listing day. $TradeVolume$ is the log total trade volume in the 7 days before the listing announcement. $Transfers$ is the log number of token transfers in the 7 days leading up the the listing announcement. $TransactionVolume$ is the log dollar value of on-chain transfers in the 7 days leading up the the listing announcement. All continuous variables are winsorized at 1% and 99%. Standard errors are clustered by announcement day. t-statistics are reported in parentheses. Significance at the 10%, 5%, and 1% levels is indicated by *, **, and ***, respectively.

	(1)	(2)	(3)	(4)	(5)	(6)
Dependent variable:	$CAR_{[-168,-1]}$		$CAR_{[-72,-1]}$		$CAR_{[-24,-1]}$	
$Tipoff$	−0.139*	−0.028	−0.111*	0.003	−0.106**	−0.067
	(−1.86)	(−0.28)	(−1.96)	(0.04)	(−2.51)	(−1.15)
DOJ	0.084	0.188**	0.072	0.173**	0.064	0.094*
	(1.23)	(2.24)	(1.25)	(2.40)	(1.47)	(1.68)
$PreviousListings$	−0.050	−0.059	0.014	0.003	0.051**	0.046**
	(−1.29)	(−1.57)	(0.47)	(0.11)	(2.62)	(2.30)
$MarketCap$	−0.043**	−0.034**	−0.041**	−0.031**	−0.042***	−0.038***
	(−2.33)	(−2.07)	(−2.45)	(−2.41)	(−3.42)	(−3.56)
$TradeVolume$	0.010	0.009	0.003	0.003	−0.003	−0.003
	(0.86)	(0.79)	(0.45)	(0.38)	(−0.61)	(−0.56)
$Transfers$	0.004	−0.013	−0.005	−0.024**	−0.015	−0.022*
	(0.29)	(−0.96)	(−0.41)	(−2.01)	(−1.19)	(−1.85)
$TransferVolume$	0.009	0.011*	0.010*	0.013**	0.010*	0.012**
	(1.37)	(1.83)	(1.91)	(2.48)	(1.83)	(2.12)
$time$		0.001		0.000		0.000
		(1.28)		(1.08)		(0.26)
$time^2$		−0.426		−0.395		−0.102
		(−1.60)		(−1.50)		(−0.51)
$Intercept$	0.707**	0.620***	0.688**	0.613***	0.791***	0.778***
	(2.50)	(2.87)	(2.50)	(3.22)	(3.88)	(4.55)
Observations	203	203	203	203	203	203
Adjusted R^2	7.2%	9.6%	11.9%	18.0%	25.0%	25.6%

Significance: * $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$

actions are illegal and may be prosecuted, the use of concealment becomes an important tool for insiders to avoid detection if they continue trading on confidential information.

Methods of concealment are not limited to using multiple wallets to obfuscate trading. An insider can use privacy protocols such as Tornado Cash to break the flow of funds between wallets. Indeed, Table 5.3 shows that our identified insider traders use Tornado Cash significantly more than other traders. The findings in Table 5.3 also show that the insider trader wallets are extremely active and that they trade more than other traders. Another

TABLE 5.8: Regressions to assess the impact of regulatory shocks on the volume of insider trading

This table presents estimates from OLS regressions that examine the impact of the twitter tip-off and the DOJ indictment on the volume of insider trading. The dependent variable for *$Blatant* (Regressions 1 and 2), *$SemiConcealed* (Regressions 3 and 4), and *$Total* (Regressions 5 and 6) is the one plus log of the volume traded by the respective insider trading wallets identified. The key independent variables are *Tipoff* and *DOJ*. *Tipoff* is a dummy variable that is 1 after the twitter tip-off on April 12, 2022, and 0 otherwise. *DOJ* is a dummy variable that is 1 after the DOJ indictment on July 21, 2022, and 0 otherwise. *time* is days from the first listing in the sample. $time^2$ is adjusted by 10e4. *PreviousListings* is the number of listings in other exchanges the token had prior to it being listed on Coinbase. *MarketCap* is the log of the token's market capitalization on the listing day. *TradeVolume* is the log total trade volume in the 7 days before the listing announcement. *Transfers* is the log number of token transfers in the 7 days leading up to the listing announcement. *TransactionVolume* is the log dollar value of on-chain transfers in the 7 days leading up to the listing announcement. All continuous variables are winsorized at 1% and 99%. Standard errors are clustered by announcement day. *t*-statistics are reported in parentheses. Significance at the 10%, 5%, and 1% levels is indicated by *, **, and ***, respectively.

	(1)	(2)	(3)	(4)	(5)	(6)
Dependent variable:	*$Blatant*		*$SemiConcealed*		*$Total*	
Tipoff	−4.077***	−1.206	−0.467	3.366**	−0.981	2.999*
	(−3.42)	(−0.81)	(−0.29)	(2.15)	(−0.58)	(1.72)
DOJ	1.320	4.442***	0.842	4.761***	0.959	5.089***
	(1.66)	(3.81)	(0.47)	(2.66)	(0.54)	(2.83)
PreviousListings	−1.068*	−1.174**	0.021	−0.195	−0.042	−0.249
	(−1.69)	(−2.05)	(0.03)	(−0.28)	(−0.05)	(−0.34)
MarketCap	−0.032	0.109	0.160	0.396*	0.146	0.379*
	(−0.11)	(0.53)	(0.49)	(1.81)	(0.45)	(1.70)
TradeVolume	0.069	0.022	−0.029	−0.077	−0.028	−0.082
	(0.67)	(0.20)	(−0.21)	(−0.59)	(−0.19)	(−0.58)
Transfers	0.491	0.110	0.866*	0.317	0.802	0.242
	(1.39)	(0.46)	(1.89)	(0.80)	(1.60)	(0.59)
TransferVolume	0.331**	0.344**	0.346*	0.393*	0.396*	0.438**
	(2.26)	(2.56)	(1.76)	(1.95)	(1.90)	(2.09)
time		0.020***		0.023***		0.025***
		(4.36)		(4.37)		(4.58)
$time^2$		−0.143***		−0.173***		−0.184***
		(−4.01)		(−5.24)		(−5.23)
Intercept	−3.068	−6.599*	−6.763	−10.767***	−6.340	−10.671***
	(−0.59)	(−1.84)	(−1.25)	(−3.21)	(−1.14)	(−3.03)
Observations	203	203	203	203	203	203
Adjusted R^2	23.7%	32.0%	26.7%	36.8%	27.0%	38.0%

Significance: * $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$

form of obfuscation may be high trading activity in order to hide nefarious transactions.

In summary, we find insiders respond to increased regulatory scrutiny by reducing their blatant insider trading activities, consistent with the existing literature on insider trading in stock markets (Del Guercio, Odders-White, & Ready, 2017; Deuskar, Khatri, & Sunder, 2025). However, we also show that insiders increasingly conceal their trading behavior with increasing levels of regulatory surveillance. These findings have implications for regulators, as increased levels of concealment and obfuscation impact their ability to detect and prosecute insider traders.

While we do not suggest that insider trading cannot be curbed with effective regulation, our findings indicate that insider trading activity did not significantly decrease around the two shocks that we examine. This persistence is likely due to the inherent challenges of detecting insider trading in cryptocurrency markets, where pseudonymity of data and the absence of a unified regulatory jurisdiction complicate enforcement efforts.

5.5 Conclusion

We find evidence suggesting insider trading occurs ahead of 28-48% of cryptocurrency exchange listings. Using blockchain data, we identify wallets that profited an estimated $30.3 million by trading on token listing announcements. These findings point to cryptocurrency markets being susceptible to the same forms of misconduct that regulators have for a long time grappled with in traditional financial markets. We also identify cases of likely insider trading that were not included in the SEC prosecution.

Our findings have implications for regulators and prosecutors. With increased levels of regulatory oversight and enforcement in traditional financial markets, insiders are deterred from trading using inside information (Del Guercio, Odders-White, & Ready, 2017; Deuskar, Khatri, & Sunder, 2025). However, our findings show that in cryptocurrency markets, insiders instead move away from blatant forms of insider trading that are easy to detect, towards using methods of concealment that hide their behavior to continue profiting from inside information. Policing crime—such as insider trading—in cryptocurrency markets is difficult because of the pseudo-anonymity and ability to conceal activity through a variety of methods. It is likely that an arms race will occur between insiders and regulators in which, as methods of detection improve, so too does an insider's ability to conceal their behavior through various measures, creating significant challenges for regulators.

Like other forms of financial misconduct, insider trading is detrimental to the integrity of cryptocurrency markets and its presence may harm investor confidence. If investors believe cryptocurrency markets are prone to widespread insider trading, they may choose not to participate in them, hindering the realization of gains from trade. Perceptions of lawlessness and widespread misconduct can also impede adoption of cryptographically secured ways of representing securities and other financial instruments.

Insider trading in cryptocurrency markets appears to be pervasive with cases of anecdotal evidence still occurring long after the prosecution in January 2023.[26] Cryptocurrency exchanges should have incentives to not turn a blind eye to insider trading, given the risk of losing participants or even facing legal charges. Our study may help regulators and exchanges identify cases of insider trading that so far remain undetected and to increase their awareness of the use of concealment strategies that are currently being used to avoid detection. Future research might investigate whether the insider trading activity subsides with a delay, and the level of insider trading for other exchanges or blockchains.

The recent SEC prosecution signals that regulators are ready to take action to curb insider trading in cryptocurrency markets with U.S. Attorney Damian Williams commenting that "cryptocurrency markets are not lawless."[27] Effective policing of cryptocurrency markets requires surveillance approaches that can overcome the pseudo-anonymity and concealment opportunities in cryptocurrency markets.

[26]See https://bit.ly/3Yf0Xpt.
[27]See https://bit.ly/3ja4Sp4.

Appendix 5.1 Listing announcement difference in means

TABLE A5.1: Listing announcement difference in means

This table reports the difference in means between the full sample of token listings and the ERC20 token listings. *Market Cap* is the token's market capitalization on the listing day reported in $ millions. *Trade Volume* is the sum of volume in the 7 days before the listing announcement, reported in $ millions. *Previous Listings* is the number of listings in other cryptocurrency exchanges the token had prior to being listed on Coinbase. The *CAR* variable denotes cumulative abnormal returns computed over different time ranges—as indicated by the different subscripts—in hours from the time of the listing announcement. All continuous variables are winsorized at 1% and 99%.

	Sample Mean	ERC20 Mean	Δ	t-statistic
Market Cap ($ Millions)	679	327	352	2.49^{**}
Trade Volume ($ Millions)	1,085	785	300	0.78
Previous Listings	7	6	1	1.95^{*}
$CAR_{[-300,+100]}$	0.25	0.28	−0.03	−0.49
$CAR_{[-168,-1]}$	0.14	0.15	−0.01	−0.25
$CAR_{[-72,-1]}$	0.10	0.11	−0.01	−0.29
$CAR_{[-24,-1]}$	0.07	0.07	0.00	−0.17

Significance: * $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$

Appendix 5.2 Abnormal return run-ups

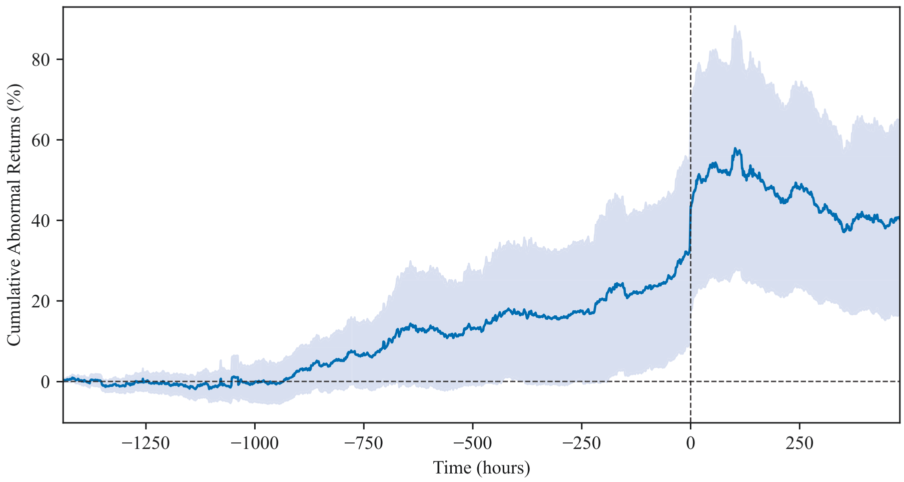

FIGURE A5.1: Cumulative abnormal returns for token listing announcements

This figure plots the mean cumulative abnormal returns for the tokens listed on Coinbase before and after the listing announcements at t = 0. The cumulative returns are measured in hours over a period of 1440 hours before the listing announcement, up until 480 hours after it (60 days before and 20 days after the listing announcment at t = 0). The shaded region represents the 5% to 95% confidence bounds. The sample contains 252 tokens that were listed on Coinbase between September 2018 to January 2024.

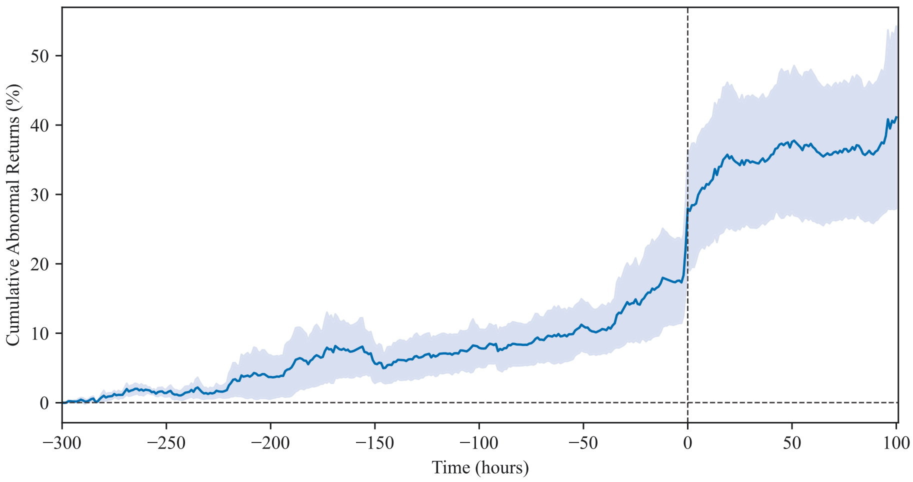

FIGURE A5.2: Cumulative abnormal returns based on small cap index for token listing announcements

This figure plots the mean cumulative abnormal returns for the tokens listed on Coinbase before and after the listing announcements at t = 0. The cumulative returns are measured in hours over a period of 300 hours before the listing announcement, up until 100 hours after it. The calculation of abnormal returns rely on the market returns from the MVIS® CryptoCompare Small Cap Digital Assets 100 Index (MVDASC). The shaded region represents the 5% to 95% confidence bounds. The sample contains 252 tokens that were listed on Coinbase between September 2018 to January 2024.

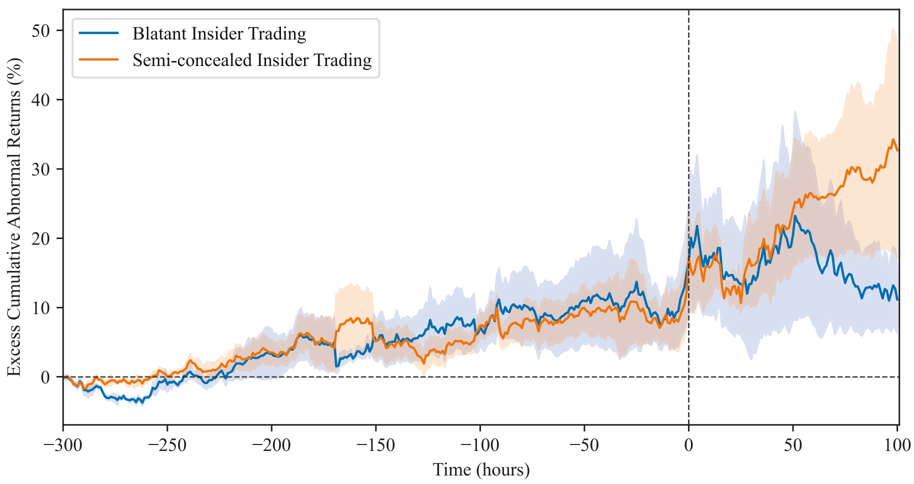

FIGURE A5.3: Excess cumulative abnormal returns based on small cap index of tokens traded by the blatant and semi-concealed insider trading wallets

This figure plots the mean cumulative abnormal returns for the tokens traded by the blatant and semi-concealed insider trading wallets in excess of the cumulative abnormal returns for the null group (token listings that are not traded by either group), before and after listing announcements at t = 0. The excess cumulative returns are measured in hours over a period of 300 hours before the listing announcement, up until 100 hours after it. The calculation of abnormal returns rely on the market returns from the MVIS® CryptoCompare Small Cap Digital Assets 100 Index (MVDASC). The blue region represents the tokens traded by the blatant wallets, while the orange region depicts the tokens traded by the semi-concealed wallets. The shaded regions display the 5% to 95% confidence bounds. The sample includes 71 tokens traded by blatant wallets and 118 traded by semi-concealed wallets.

Appendix 5.3 Prevalence of insider trading through time

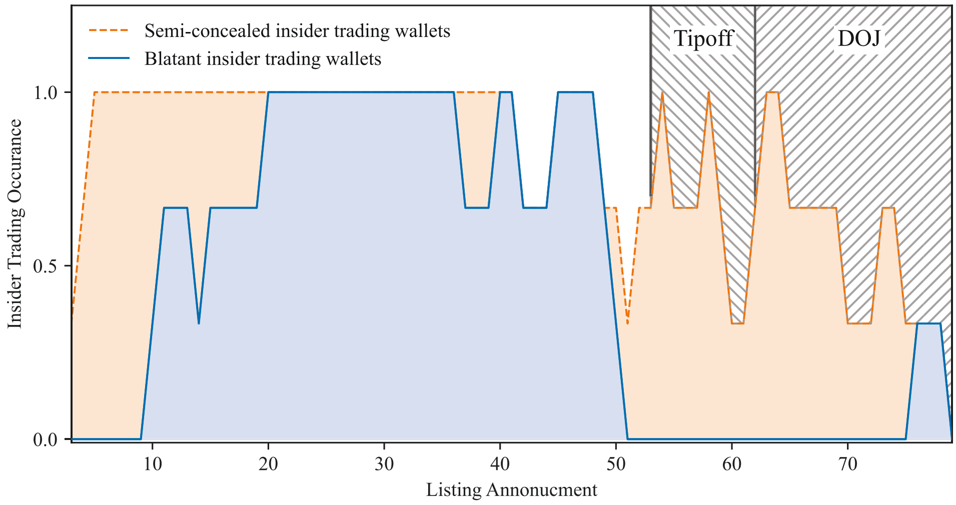

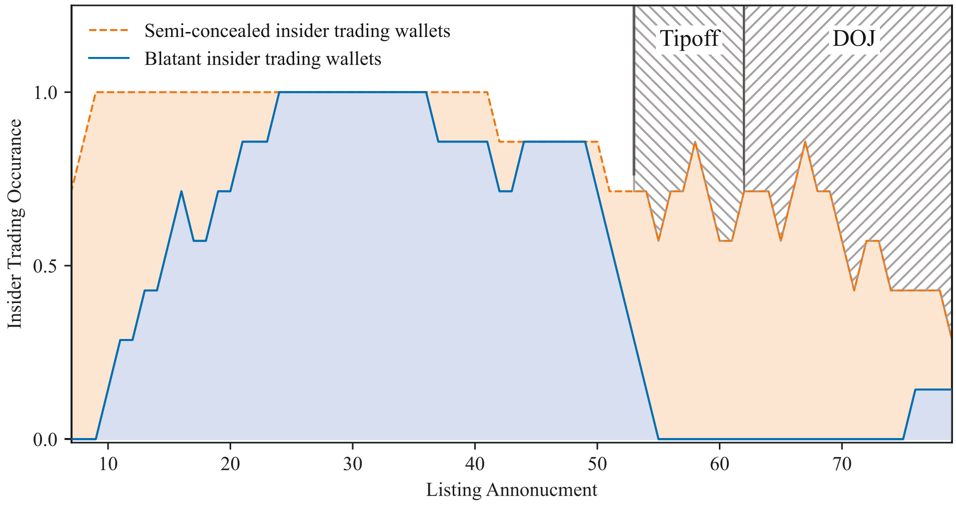

FIGURE A5.4: Prevalence of insider trading activity through time

These figures presents the 3 and 7-announcement moving average of a dummy variable that is equal to 1 if the listing announcement was traded by a blatant insider trading wallet (blue area), or only traded by a semi-concealed insider trading wallet (orange area). The top figure presents the 3-announcement moving average and the bottom figure presents the 7-announcement moving average. The downward sloping background region represents the listing announcements that were made after the tip-off. The upward sloping background region represents the listing announcements that were made after the DOJ indictment.

Chapter 6

Conclusion

Trading makes us better off, and how we trade determines the benefits we realize. Historically, technological improvements have consistently driven welfare gains. Most recently, the rise of electronic markets—driven by advances in communications technology—has made trading faster, cheaper, and more accessible.

This dissertation investigates whether Automated Market Makers (AMMs) represent the next step in the evolution of markets and trading. AMMs automate trading and enable passive liquidity provision, reducing settlement costs and the cost to create new markets. By doing so, AMMs make a wider range of assets tradable and unlock gains from trade that traditional markets cannot facilitate because of high costs. In this thesis, we evaluate AMMs along three dimensions of market quality: liquidity, price discovery, and market integrity.

In Chapter 2, we explore whether AMM liquidity can improve modern markets. AMMs offer passive liquidity provision by automating trading, having prices defined by a pricing function. In contrast, Limit Order Books (LOBs) require constant monitoring with market makers needing to replenish orders in the order book to provide liquidity. Despite the apparent structural differences, we show that measures of AMM liquidity are directly comparable to those used in LOB markets.

Chapter 2 shows that liquidity in AMMs is more stable than liquidity in LOB markets. Stable and resilient liquidity is important in times of high market turmoil, such as flash crashes. Typically, market makers step away in times of market turmoil, leading to less liquidity and worse outcomes for all market participants. We discuss several options to leverage and implement AMM like liquidity, to improve modern markets. The findings highlight the potential of AMMs to improve liquidity in modern markets.

Chapter 3 examines whether AMMs facilitate efficient price discovery, despite relying solely on liquidity demanders to incorporate information into prices. Using a variance decomposition framework, we measure the share of information and noise in AMM prices. We find AMMs can provide efficient price discovery when they are the only market. AMMs can lead the price discovery process when they are more liquid than LOBs. These results show that while structurally different, AMMs can generate information, a key function of financial markets.

Evolutions in trading often aim to reduce costs for traders. With paper money, physical storage costs were reduced; with AMMs, the key advantage lies in reducing the monitoring costs faced by market makers. Modern market makers must balance the cost of monitoring markets against the cost of adverse selection (Khomyn & Putniņš, 2021). AMMs shift this trade-off by lowering the monitoring, automating liquidity provision. As a result, Liquidity Providers (LPs) (market makers) can accept higher adverse selection costs—creating more profit opportunities for informed traders, thus creating greater incentive for information production, reinforcing the wider benefits of AMMs.

Chapter 4 explores market integrity in AMM-based Decentralized Exchanges (DEXs), leveraging the transparency of blockchain data to detect misconduct, we identify approximately $345 billion in DEX volume associated with manipulative or illicit behavior—including both traditional forms, like wash trading, rug pulls and money laundering, along with AMM-specific practices of misconduct, such as sandwich attacks. We estimate that 12% of total DEX volume involves misconduct. Despite this, market integrity issues do not appear to deter users from trading on DEXs.

Chapter 5 examines how DEXs and AMMs shape integrity in the broader crypto market by analyzing the first prosecuted case of insider trading on token listing announcements. We find evidence of systematic insider trading, supported by both on-chain transaction patterns and traditional price run-ups. From our blockchain data, we find that insiders we identify prefer to trade on DEXs, likely because of their permissionless and pseudonymous nature.

In Chapter 5, we examine how bad actors respond to the increase in regulatory oversight. We find insiders shift their behavior in response to regulation aiming to conceal their behavior by using new wallets to continue trading around listing announcements. Through wallet clustering, we identify direct links between these wallets. These findings underscore the regulatory challenges of blockchain-based markets: transparency aids detection, but illicit actors adapt instead of ceasing their actions, highlighting the need for robust surveillance systems to adequately regulate DEXs.

Overall, this dissertation shows AMMs are a viable market structure with the potential to improve liquidity provision in modern markets. AMMs can provide key functions of financial markets while offering more stable and resilient liquidity. The findings of this dissertation provide strong support for the view that AMMs represent the next evolution in trading. However, these benefits come with trade-offs. Chapters 4 and 5 emphasize the challenges of regulating DEXs. AMMs introduce novel forms of market misconduct and on DEXs enforcement and regulation of markets is challenging. Regardless, tackling these challenges would only further highlight the welfare benefits that AMMs could unlock.

6.1 Future research

Chapter 2 highlights the benefit of AMM liquidity being more stable and proposes several hybrid market designs where AMMs could complement traditional LOB markets. One such design involves integrating AMM-like orders into LOBs—orders that do not expire when traded against. This structure could enhance the resilience liquidity in markets.

However, what is unclear is the impact these orders would have on price formation and liquidity in these markets. Future research could explore how introducing new order types could influence price formation and liquidity.

AMM design remains in its infancy, with ongoing industry experimentation driving rapid innovation in how markets operate on-chain through AMMs. As AMMs continue to evolve, future research can investigate how different design choices affect market quality. Understanding how different design choices impact market quality will be key to shaping the future of AMMs.

Future research can explore how AMMs can be used to facilitate trading in asset classes that are infrequently traded. AMMs have low costs to creating markets and enable passive liquidity provision, reducing monitoring costs. Exploring how AMMs can be used to trade illiquid assets like private equity or non/quasi fungible assets is an avenue for future research.

Chapters 4 and 5 reveal the scale of market misconduct in AMM-based DEXs and the challenges of enforcing rules in blockchain-based markets. While transparency enables detection, current regulatory frameworks are not suitable for these markets. Future research can explore solutions to emerging forms of manipulation, such as sandwich attacks, and address the legal complexities of tokenization, market infrastructure, and on-chain regulation. Advancing surveillance and enforcement mechanisms will be essential to ensure integrity as these markets mature.

Bibliography

Acharya, V. V., & Johnson, T. C. (2010). More insiders, more insider trading: Evidence from private-equity buyouts. *Journal of Financial Economics*, *98*(3), 500–523.

Adams, A., Lader, M.-C., Liao, G., Puth, D., & Wan, X. (2023). On-Chain Foreign Exchange and Cross-Border Payments. *Working paper.*

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

Admati, A. R., & Pfleiderer, P. (1988). A Theory of Intraday Patterns: Volume and Price Variability. *The Review of Financial Studies*, *1*(1), 3–40.

Ahern, K. R. (2017). Information networks: Evidence from illegal insider trading tips. *Journal of Financial Economics*, *125*(1), 26–47.

Ahern, K. R. (2020). Do proxies for informed trading measure informed trading? evidence from illegal insider trades. *Review of Asset Pricing Studies*, *10*(3), 397–440.

Akey, P., Grégoire, V., & Martineau, C. (2022). Price revelation from insider trading: Evidence from hacked earnings news. *Journal of Financial Economics*, *143*(3), 1162–1184.

Alexander, C., Chen, X., Deng, J., & Fu, Q. (2023). Impacts of Protocol Updates on Uniswap Decentralized Crypto Exchanges. *Working paper.*

Alexander, C., Heck, D. F., & Kaeck, A. (2022). Price Discovery in Bitcoin: The Role of Limit Orders. *Working paper.*

Aliyev, N., Allahverdiyeva, I., & Putniņš, T. J. (2023). Scam alert: Can cryptocurrency scams be detected early? *Working paper.*

Alizadeh, M., Kabbani, A., Edsall, T., Prabhakar, B., Vahdat, A., & Yasuda, M. (2012). Less is more: Trading a little bandwidth for {ultra-low} latency in the data center. *9th USENIX Symposium on Networked Systems Design and Implementation (NSDI 12)*, 253–266.

Amihud, Y. (2002). Illiquidity and stock returns: Cross-section and time-series effects. *Journal of Financial Markets*, *5*(1), 31–56.

Amihud, Y., & Mendelson, H. (1986). Liquidity and stock returns. *Financial Analysts Journal*, *42*(3), 43–48.

Anderson, J. P. (2019). Insider trading and cryptoassets: The waters just got muddier. *Iowa Law Review*, *104*, 120.

Angeris, G., Chitra, T., Diamandis, T., Evans, A., & Kulkarni, K. (2023). The geometry of constant function market makers. *Working paper.*

Ante, L. (2019). Market reaction to exchange listings of cryptocurrencies. *BRL Working Paper Series.*

Aoyagi, J., & Ito, Y. (2021). Coexisting Exchange Platforms: Limit Order Books and Automated Market Makers. *Working paper.*

Aquilina, M., Foley, S., Gambacorta, L., & Krekel, W. (2024). Decentralised dealers? Examining liquidity provision in decentralised exchanges. *Working paper.*

Aspris, A., Foley, S., Svec, J., & Wang, L. (2021). Decentralized exchanges: The "wild west" of cryptocurrency trading. *International Review of Financial Analysis*, *77*, 101845.

Ates, A., & Wang, G. H. K. (2005). Information transmission in electronic versus open-outcry trading systems: An analysis of U.S. equity index futures markets. *The Journal of Futures Markets*, *25*(7), 679–715.

Augustin, P., Brenner, M., & Subrahmanyam, M. G. (2019). Informed options trading prior to takeover announcements: Insider trading? *Management Science*, *65*(12), 5697–5720.

Austin, J. (2017). What Exactly is Market Integrity? An Analysis of One of the Core Objectives of Securities Regulation. *William & Mary Business Law Review*, *8*(2), 215.

Bagourd, A., & Francois, L. G. (2023). Quantifying MEV On Layer 2 Networks. *Working paper.*

Balaji, A. (2024). Wire Fraud on the Blockchain: Backdoor Prosecutions of Insider Trading. *Criminal Law Bulletin*, *60*(3), 334–377.

Bank for International Settlements. (2023). Project Mariana: Cross-border exchange of wholesale CBDCs using automated market-makers.

Barbon, A., & Ranaldo, A. (2022). On The Quality Of Cryptocurrency Markets: Centralized Versus Decentralized Exchanges. *Working paper.*

Barclay, M. J., & Warner, J. B. (1993). Stealth trading and volatility: Which trades move prices? *Journal of financial Economics*, *34*(3), 281–305.

Benedetti, H., & Nikbakht, E. (2021). Returns and network growth of digital tokens after cross-listings. *Journal of Corporate Finance*, *66*, 101853.

Beveridge, S., & Nelson, C. R. (1981). A new approach to decomposition of economic time series into permanent and transitory components with particular attention to measurement of the 'business cycle'. *Journal of Monetary Economics*, *7*(2), 151–174.

Black, F. (1971). Toward a Fully Automated Stock Exchange. *Financial Analysts Journal*, *27*(4), 28–44.

Bolz, M., Bründler, K., Kane, L., Patsias, P., Tessendorf, L., Gogol, K., Kim, T., & Tessone, C. (2024). Machine Learning-Based Detection of Pump-and-Dump Schemes in Real-Time. *Working paper.*

Booth, G. G., Lin, J.-C., Martikainen, T., & Tse, Y. (2002). Trading and Pricing in Upstairs and Downstairs Stock Markets. *The Review of Financial Studies*, *15*(4), 1111–1135.

Boulatov, A., & George, T. J. (2013). Hidden and Displayed Liquidity in Securities Markets with Informed Liquidity Providers. *The Review of Financial Studies*, *26*(8), 2096–2137.

Brogaard, J., Hendershott, T., & Riordan, R. (2019). Price Discovery without Trading: Evidence from Limit Orders. *The Journal of Finance*, *74*(4), 1621–1658.

Brogaard, J., Li, D., & Yang, J. (2022). Does High Frequency Market Manipulation Harm Market Quality? *Working paper.*

Brogaard, J., Nguyen, T. H., Putniņš, T. J., & Wu, E. (2022). What moves stock prices? the roles of news, noise, and information. *The Review of Financial Studies*, *35*(9), 4341–4386.

Budish, E., Cramton, P., & Shim, J. (2015). The High-Frequency Trading Arms Race: Frequent Batch Auctions as a Market Design Response. *The Quarterly Journal of Economics*, *130*(4), 1547–1621.

Cabrera, J., Wang, T., & Yang, J. (2009). Do futures lead price discovery in electronic foreign exchange markets? *Journal of Futures Markets*, *29*(2), 137–156.

Capponi, A., Jia, R., & Yu, S. (2023). Price Discovery on Decentralized Exchanges. *Working paper.*

Cernera, F., La Morgia, M., Mei, A., & Sassi, F. (2023). Token spammers, rug pulls, and sniperbots: An analysis of the ecosystem of tokens in ethereum and the binance smart chain (bnb). *32nd USENIX Security Symposium (USENIX Security 23)*, 3349–3366.

Chainalysis. (2025). 2025 Crypto Crime Trends from Chainalysis. https://www.chainalysis.com/blog/2025-crypto-crime-report-introduction/.

Chakravarty, S., Gulen, H., & Mayhew, S. (2004). Informed Trading in Stock and Option Markets. *The Journal of Finance*, *59*(3), 1235–1257.

Chang, I., Johnson, L., & Putniņš, T. J. (2025). Misconduct in Decentralized Exchanges. *Working paper.*

Chemaya, N., Liu, D., McLaughlin, R., Ruaro, N., Kruegel, C., & Vigna, G. (2024). The power of default: Measuring the effect of slippage tolerance in decentralized exchanges. *International Conference on Financial Cryptography and Data Security*, 192–208.

Chen, H., & Choi, P. M. S. (2012). Does information vault Niagara Falls? Cross-listed trading in New York and Toronto. *Journal of Empirical Finance*, *19*(2), 175–199.

Chordia, T., Roll, R., & Subrahmanyam, A. (2008). Liquidity and market efficiency. *Journal of Financial Economics*, *87*(2), 249–268.

Cohen, L., Malloy, C., & Pomorski, L. (2012). Decoding inside information. *The Journal of Finance*, *67*(3), 1009–1043.

Collin-Dufresne, P., & Fos, V. (2015). Do Prices Reveal the Presence of Informed Trading? *The Journal of Finance*, *70*(4), 1555–1582.

Comerton-Forde, C., & Putniņš, T. J. (2015). Dark trading and price discovery. *Journal of Financial Economics*, *118*(1), 70–92.

Cong, L. W., Landsman, W., Maydew, E., & Rabetti, D. (2023a). Tax-loss harvesting with cryptocurrencies. *Journal of Accounting and Economics*, *76*(2-3), 101607.

Cong, L. W., Li, X., Tang, K., & Yang, Y. (2023b). Crypto Wash Trading. *Management Science*, *69*(11), 6427–6454.

Cong, W., Harvey, C., Rabetti, D., & Wu, Z.-Y. (2025). An Anatomy of Crypto-Enabled Cybercrimes. *Management Science*, *71*(4), 3622–3633.

Conlon, T., Corbet, S., Hou, G., Hu, Y., & Oxley, L. (2022). Beyond the Noise – Information Discovery in Bitcoin Revisited. *Working paper*.

Daian, P., Goldfeder, S., Kell, T., Li, Y., Zhao, X., Bentov, I., Breidenbach, L., & Juels, A. (2019). Flash Boys 2.0: Frontrunning, Transaction Reordering, and Consensus Instability in Decentralized Exchanges. *Working paper*.

Degryse, H., Jong, F. D., Ravenswaaij, M. V., & Wuyts, G. (2005). Aggressive Orders and the Resiliency of a Limit Order Market. *Review of Finance*, *9*(2), 201–242.

Del Guercio, D., Odders-White, E. R., & Ready, M. J. (2017). The deterrent effect of the securities and exchange commission's enforcement intensity on illegal insider trading: Evidence from run-up before news events. *Journal of Law and Economics*, *60*(2), 269–307.

DeMarzo, P. M., Fishman, M. J., & Hagerty, K. M. (1998). The optimal enforcement of insider trading regulations. *Journal of Political Economy*, *106*(3), 602–632.

Deng, Q., & Zhou, Z.-G. (2024). Daily Liquidity Jump and Diffusion, and Treatment on Crypto Wash Trading. *Working paper.*

Deuskar, P., Khatri, A., & Sunder, J. (2025). Insider trading restrictions and informed trading in peer stocks. *Management Science*, *71*(3), 2390–2412.

Dhawan, A., & Putniņš, T. J. (2023). A new wolf in town? pump-and-dump manipulation in cryptocurrency markets. *Review of Finance*, *27*(3), 935–975.

Eglīte, E., Štaermans, D., Patel, V., & Putniņš, T. J. (2022). Using ETFs to Conceal Insider Trading. *Working paper.*

Eun, C. S., & Sabherwal, S. (2003). Cross-Border Listings and Price Discovery: Evidence from U.S.-Listed Canadian Stocks. *The Journal of Finance*, *58*(2), 549–575.

European Securities and Markets Authority. (2023a). *Decentralised finance – a categorisation of smart contracts.* Publications Office of the European Union.

European Securities and Markets Authority. (2023b). *Decentralised finance in the EU: Developments and risks.* Publications Office of the European Union.

Fama, E. F. (1970). Efficient Capital Markets: A Review of Theory and Empirical Work. *The Journal of Finance*, *25*(2), 383–417.

Félez-Viñas, E., Johnson, L., & Putniņš, T. J. (2025a). Can Automated Market Makers Provide Price Discovery? *Working paper.*

Félez-Viñas, E., Johnson, L., & Putniņš, T. J. (2025b). Insider trading in cryptocurrency markets. *Working paper.*

Félez-Viñas, E., Johnson, L., & Putniņš, T. J. (2025c). Liquidity That Lasts: Learning from AMMs for Market Liquidity. *Working paper.*

Feng, W., Wang, Y., & Zhang, Z. (2018). Informed trading in the bitcoin market. *Finance Research Letters*, *26*, 63–70.

Foley, S., Karlsen, J. R., & Putniņš, T. J. (2019). Sex, drugs, and bitcoin: How much illegal activity is financed through cryptocurrencies? *Review of Financial Studies*, *32*(5), 1798–1853.

Foley, S., Karlsen, J. R., & Putniņš, T. J. (2022). Measuring global market integrity. *Working paper.*

Foley, S., Meling, T. G., & Ødegaard, B. A. (2023). Tick Size Wars: The Market Quality Effects of Pricing Grid Competition. *Review of Finance*, *27*(2), 659–692.

Foley, S., O'Neill, P., & Putniņš, T. J. (2023). A Better Market Design? Applying 'Automated Market Makers' to Traditional Financial Markets. *Working paper.*

Foley, S., & Putniņš, T. J. (2016). Should we be afraid of the dark? dark trading and market quality. *Journal of Financial Economics*, *122*(3), 456–481.

Fong, K. Y., Holden, C. W., & Trzcinka, C. A. (2017). What are the best liquidity proxies for global research? *Review of Finance*, *21*(4), 1355–1401.

Foucault, T., Kadan, O., & Kandel, E. (2005). Limit Order Book as a Market for Liquidity. *The Review of Financial Studies*, *18*(4), 1171–1217.

Foucault, T., Pagano, M., & Roell, A. (2013). *Market liquidity: Theory, evidence, and policy.* Oxford University Press.

Fricke, C., & Menkhoff, L. (2011). Does the "Bund" dominate price discovery in Euro bond futures? Examining information shares. *Journal of Banking & Finance*, *35*(5), 1057–1072.

Gan, R., Wang, L., Xue, L., & Lin, X. (2024). Exposing stealthy wash trading on automated market maker exchanges. *ACM Transactions on Internet Technology*, *24*(4), 1–30.

Gandal, N., Hamrick, J., Moore, T., & Oberman, T. (2018). Price manipulation in the bitcoin ecosystem. *Journal of Monetary Economics*, *95*, 86–96.

Glosten, L. R., & Milgrom, P. R. (1985). Bid, ask and transaction prices in a specialist market with heterogeneously informed traders. *Journal of Financial Economics*, *14*(1), 71–100.

Gonzalo, J., & Granger, C. (1995). Estimation of common long-memory components in cointegrated. *Journal of Business & Economic Statistics*, *13*(1), 27.

Goyenko, R. Y., Holden, C. W., & Trzcinka, C. A. (2009). Do liquidity measures measure liquidity? *Journal of Financial Economics*, *92*(2), 153–181.

Griffin, J. M., & Mei, K. (2024). How Do Crypto Flows Finance Slavery? The Economics of Pig Butchering. *Working paper.*

Griffin, J. M., & Shams, A. (2020). Is bitcoin really untethered? *The Journal of Finance*, *75*(4), 1913–1964.

Hamrick, J., Rouhi, F., Mukherjee, A., Feder, A., Gandal, N., Moore, T., & Vasek, M. (2021). An examination of the cryptocurrency pump-and-dump ecosystem. *Information Processing & Management*, *58*(4), 102506.

Hanson, R. (2003). Combinatorial Information Market Design. *Information Systems Frontiers*, *5*(1), 107–119.

Harris, L. (2003). *Trading and Exchanges – Market Microstructure for Practitioners*. Oxford University Press.

Hasbrouck, J. (1995). One Security, Many Markets: Determining the Contributions to Price Discovery. *The Journal of Finance*, *50*(4), 1175–1199.

Hasbrouck, J. (2004). Liquidity in the Futures Pits: Inferring Market Dynamics from Incomplete Data. *Journal of Financial and Quantitative Analysis*, *39*(2), 305–326.

Hasbrouck, J. (2007). *Empirical market microstructure: The institutions, economics, and econometrics of securities trading* (1st ed.). Oxford University Press.

Hasbrouck, J. (2021). Price Discovery in High Resolution. *Journal of Financial Econometrics*, *19*(3), 395–430.

Heimbach, L., Pahari, V., & Schertenleib, E. (2024). Non-atomic arbitrage in decentralized finance. *2024 IEEE Symposium on Security and Privacy (SP)*, 3866–3884.

Heimbach, L., Schertenleib, E., & Wattenhofer, R. (2022). Risks and returns of uniswap v3 liquidity providers. *Proceedings of the 4th ACM Conference on Advances in Financial Technologies*, 89–101.

Heimbach, L., Wang, Y., & Wattenhofer, R. (2021). Behavior of Liquidity Providers in Decentralized Exchanges. *Working paper*.

Heimbach, L., & Wattenhofer, R. (2022). Eliminating Sandwich Attacks with the Help of Game Theory. *Proceedings of the 2022 ACM on Asia Conference on Computer and Communications Security*, 153–167.

Hornuf, L., Momtaz, P. P., Nam, R. J., & Yuan, Y. (2025). Cybercrime on the ethereum blockchain. *Journal of Banking & Finance*, *175*, 107419.

Hsieh, W.-L. G., Lee, C.-S., & Yuan, S.-F. (2008). Price discovery in the options markets: An application of put-call parity. *Journal of Futures Markets*, *28*(4), 354–375.

Huang, R. D., & Stoll, H. R. (1996). Dealer versus auction markets: A paired comparison of execution costs on nasdaq and the nyse. *Journal of Financial Economics*, *41*(3), 313–357.

John, K., Monnot, B., Mueller, P., Saleh, F., & Schwarz-Schilling, C. (2025). Economics of Ethereum. *Journal of Corporate Finance*, *91*, 102718.

Kacperczyk, M., & Pagnotta, E. S. (2019). Chasing private information. *Review of Financial Studies*, *32*(12), 4997–5047.

Kacperczyk, M., & Pagnotta, E. S. (2024). Legal risk and insider trading. *The Journal of Finance*, *79*(1), 305–355.

Kallunki, J., Kallunki, J.-P., Nilsson, H., & Puhakka, M. (2018). Do an insider's wealth and income matter in the decision to engage in insider trading? *Journal of Financial Economics*, *130*(1), 135–165.

Kamps, J., & Kleinberg, B. (2018). To the moon: Defining and detecting cryptocurrency pump-and-dumps. *Crime Science*, *7*(1), 1–18.

Khomyn, M., & Putniņš, T. J. (2021). Algos gone wild: What drives the extreme order cancellation rates in modern markets? *Journal of Banking & Finance*, *129*, 106170.

Klein, O., Kozhan, R., Viswanath-Natraj, G., & Wang, J. (2023). Price Discovery in Cryptocurrencies: Trades versus Liquidity Provision. *Working paper*.

Kyle, A. S. (1985). Continuous Auctions and Insider Trading. *Econometrica*, *53*(6), 1315–1335.

Kyle, A. S. (1989). Informed Speculation with Imperfect Competition. *The Review of Economic Studies*, *56*(3), 317–355.

Large, J. (2007). Measuring the resiliency of an electronic limit order book. *Journal of Financial Markets*, *10*(1), 1–25.

Lehar, A., & Parlour, C. (2025). Decentralized Exchange: The Uniswap Automated Market Maker. *The Journal of Finance*, *80*(1), 321–374.

Lehmann, B. N. (2002). Some desiderata for the measurement of price discovery across markets. *Journal of Financial Markets*, *5*(3), 259–276.

Levine, R. (2005, January 1). Chapter 12 Finance and Growth: Theory and Evidence. In P. Aghion & S. N. Durlauf (Eds.), *Handbook of Economic Growth* (pp. 865–934, Vol. 1). Elsevier.

Li, T., Shin, D., & Wang, B. (2021). Cryptocurrency pump-and-dump schemes. *Working paper.*

Liu, Y., Tsyvinski, A., & Wu, X. (2022). Common risk factors in cryptocurrency. *The Journal of Finance*, *77*(2), 1133–1177.

Lo, D. K., & Hall, A. D. (2015). Resiliency of the limit order book. *Journal of Economic Dynamics and Control*, *61*, 222–244.

Lundblad, C. T., Yang, Z., & Zhang, Q. (2022). Detecting insider trading in the era of big data and machine learning. *Working paper.*

Malinova, K., & Park, A. (2023). Learning from DeFi: Would Automated Market Makers Improve Equity Trading? *Working paper.*

Mazorra, B., Adan, V., & Daza, V. (2022). Do not rug on me: Zero-dimensional Scam Detection. *Working paper.*

Milionis, J., Moallemi, C. C., & Roughgarden, T. (2024). Automated market making and arbitrage profits in the presence of fees. *International Conference on Financial Cryptography and Data Security*, 159–171.

Milionis, J., Moallemi, C. C., Roughgarden, T., & Zhang, A. L. (2022). Automated Market Making and Loss-Versus-Rebalancing. *Working paper.*

Muravyev, D., Pearson, N. D., & Paul Broussard, J. (2013). Is there price discovery in equity options? *Journal of Financial Economics*, *107*(2), 259–283.

Nadler, M., & Schär, F. (2023). Tornado Cash and Blockchain Privacy: A Primer for Economists and Policymakers. *Federal Reserve Bank of St. Louis Review.*

Neuder, M., Rao, R., Moroz, D. J., & Parkes, D. C. (2021). Strategic Liquidity Provision in Uniswap v3. *Working paper.*

Nguyen, K., Putniņš, T. J., Nguyen, T. H., & Do, B. L. (2024). Why are cryptocurrencies so volatile? *Working paper.*

Nguyen, K. Q., & Putniņš, T. J. (2023). Stealing Digital Assets: Evidence From Cryptocurrency Exchange Hacks. *Working paper.*

Nuutila, E., & Soisalon-Soininen, E. (1994). On finding the strongly connected components in a directed graph. *Information processing letters*, *49*(1), 9–14.

O'Hara, M. (2003). Presidential address: Liquidity and price discovery. *The Journal of Finance*, *58*(4), 1335–1354.

Othman, A. (2021). New Invarients for Automated Market Making. *Working paper.*

Othman, A., Pennock, D. M., Reeves, D. M., & Sandholm, T. (2013). A Practical Liquidity-Sensitive Automated Market Maker. *ACM Transactions on Economics and Computation*, *1*(3), 14:1–14:25.

Othman, A., & Sandholm, T. (2011). Liquidity-Sensitive Automated Market Makers via Homogeneous Risk Measures. *International Workshop on Internet and Network Economics*, 314–325.

Paquet-Clouston, M., Haslhofer, B., & Dupont, B. (2019). Ransomware payments in the bitcoin ecosystem. *Journal of Cybersecurity*, *5*(1).

Patel, V., & Putniņš, T. J. (2020). How much insider trading happens in stock markets? *Working paper.*

Patel, V., Putniņš, T. J., Michayluk, D., & Foley, S. (2020). Price discovery in stock and options markets. *Journal of Financial Markets*, *47*, 100524.

Putniņš, T. J. (2013). What do price discovery metrics really measure? *Journal of Empirical Finance*, *23*, 68–83.

Putniņš, T. J. (2020). An overview of market manipulation. *Corruption and Fraud in Financial Markets: Malpractice, Misconduct and Manipulation*, *13*.

Qin, K., Zhou, L., & Gervais, A. (2022). Quantifying blockchain extractable value: How dark is the forest? *2022 IEEE Symposium on Security and Privacy (SP)*, 198–214.

Rösch, D. M., Subrahmanyam, A., & van Dijk, M. A. (2017). The Dynamics of Market Efficiency. *The Review of Financial Studies*, *30*(4), 1151–1187.

Roşu, I. (2009). A Dynamic Model of the Limit Order Book. *The Review of Financial Studies*, *22*(11), 4601–4641.

See, K. (2023). The satoshi laundromat: A review on the money laundering open door of bitcoin mixers. *Journal of Financial Crime*, *31*(2), 416–426.

Sila, J., Kočenda, E., Kristoufek, L., & Kukacka, J. (2024). Determinants of wash trading in major cryptoexchanges. *Working paper*.

Smith, A. (1776). An inquiry into the nature and causes of the wealth of nations.

Sokolov, K. (2021). Ransomware activity and blockchain congestion. *Journal of Financial Economics*, *141*(2), 771–782.

Solidus Labs. (2023). Insider Trading Detected in 56% of Token Listings. https://www.soliduslabs.com/reports/crypto-insider-trading.

Verstein, A. (2019). Crypto assets and insider trading law's domain. *Iowa Law Review*, *105*, 1.

Victor, F., & Weintraud, A. M. (2021). Detecting and Quantifying Wash Trading on Decentralized Cryptocurrency Exchanges. *Proceedings of the Web Conference 2021*, 23–32.

Volosnikov, V., Tikhomirov, V., Azhel, I., & Chizh, I. (2021). ALGEBRA Ecosystem: Decentralized exchange. *White paper*.

Voltaire. (1734). *Lettres philosophiques*. E. Lucas, au Livre d'or Amsterdam.

Xu, J., Paruch, K., Cousaert, S., & Feng, Y. (2022). SoK: Decentralized Exchanges (DEX) with Automated Market Maker (AMM) Protocols. *ACM Computing Surveys*, *55*(11), 1–50.

Zhang, Y., Chen, X., & Park, D. (2018). Formal Specification of Constant Product (x * y = k) Market Maker Model and Implementation. *White paper*.

Zornow, D., Strauber, J., & Merzel, D. (2022). Financial crime in the United Stated: Overview.