



Contents lists available at [ScienceDirect](https://www.sciencedirect.com)

# Journal of International Financial Markets, Institutions & Money

journal homepage: [www.elsevier.com/locate/intfin](http://www.elsevier.com/locate/intfin)



## Is decentralized always better? How market structure affects trading costs for tokenized assets<sup>☆</sup>

Angelo Aspris<sup>a</sup>, Anne Haubo Dyhrberg<sup>b</sup>, Sean Foley<sup>c,\*</sup>, William Krekel<sup>c</sup>,  
Talis J. Putnins<sup>d</sup>

<sup>a</sup> University of Sydney, Australia

<sup>b</sup> Wilfrid Laurier University, Canada

<sup>c</sup> Macquarie University, Australia

<sup>d</sup> Digital Finance Co-operative Research Centre and UTS, Australia

### ARTICLE INFO

#### JEL classification:

G10  
G12  
G14  
G15  
G23

#### Keywords:

Market microstructure  
Decentralized finance (DeFi)  
Automated Market Makers (AMMs)  
Tokenization  
Real-world assets (RWA)  
Cryptocurrency

### ABSTRACT

Decentralization in financial markets operates at multiple layers, encompassing both the asset level, through tokenization, and the infrastructure level, via novel trading protocols such as Automated Market Makers (AMMs). This paper critically examines these twin dimensions of decentralization, focusing on the introduction of tokenized equities and the adoption of AMM mechanisms in securities market design. We find that the first tokenized equities struggle to gain adoption, resulting in poor liquidity (high price impact) and substantial price dislocations from the underlying shares. In contrast, we show that AMMs can significantly reduce transaction costs, particularly for large size trades in actively traded assets, using both matching and simulation mechanisms between centralized exchanges and AMMs. Our results suggest that while not all decentralized innovations face an easy path to adoption, some have significant potential for application to traditional asset classes.

### 1. Introduction

Blockchains and distributed ledger technologies have enabled the development of new financial infrastructures, offering decentralized frameworks for recording asset ownership and facilitating trade. In less than two decades, the digital asset market has grown to support a vast ecosystem of tokens, traded on innovative trading venues, with daily activity routinely exceeding US\$100 billion.<sup>1</sup> The emergence of decentralized finance (DeFi) promises to enhance market accessibility and efficiency through disintermediation. However, the success of any financial innovation is ultimately determined by its ability to overcome a complex array of frictions, a challenge that has long motivated the design of new securities and market structures (Allen and Gale, 1988). This paper investigates the central question of whether decentralized architectures are inherently superior to established centralized structures by examining two distinct, yet related, applications: the tokenization of traditional equities and the competition between centralized (CEX) and decentralized (DEX) exchanges for crypto-native assets. We find that the answer is nuanced; while

<sup>☆</sup> An early version of this paper was commissioned by the SWIFT Institute and subsequently published as SWIFT Institute Working Paper No. 2021-001 under the title: "Digital Assets and Markets: A Transaction-Cost analysis of market architectures".

\* Corresponding author.

E-mail addresses: [angelo.aspris@sydney.edu.au](mailto:angelo.aspris@sydney.edu.au) (A. Aspris), [adyhrberg@wlu.ca](mailto:adyhrberg@wlu.ca) (A.H. Dyhrberg), [sean.foley@mq.edu.au](mailto:sean.foley@mq.edu.au) (S. Foley), [william.krekel@mq.edu.au](mailto:william.krekel@mq.edu.au) (W. Krekel), [talis.putnins@uts.edu.au](mailto:talis.putnins@uts.edu.au) (T.J. Putnins).

<sup>1</sup> The active trading universe is considerably concentrated, with the top 100 tokens accounting for more than 90% of total market capitalization.

<https://doi.org/10.1016/j.intfin.2026.102302>

Received 29 July 2025; Accepted 7 February 2026

Available online 18 April 2026

1042-4431/Crown Copyright © 2026 Published by Elsevier B.V. This is an open access article under the CC BY license (<http://creativecommons.org/licenses/by/4.0/>).

decentralized innovations demonstrate significant potential, their ability to deliver superior market quality is critically contingent on overcoming specific, context dependent market frictions.

Our analysis proceeds in two parts, each exploring a different facet of decentralization. First, we examine one of the earliest ambitions of DeFi: the tokenization of real-world assets (RWA). The literature posits that tokenization can generate efficiencies in issuance, trading, and post-trade processes (Malinova and Park, 2023b; Carapella et al., 2023; Kreppmeier et al., 2023; Lee et al., 2024). To explore the potential for efficiencies in trading, we empirically examine one of the first market offerings for tokenized equities, which are designed to gain economic exposure to regulated underlying stocks.<sup>2</sup> We confine this exploratory analysis to a comparative examination of transaction costs and price efficiency and show that tokenized equity markets are characterized by illiquid limit order books, high transaction costs, and low levels of price efficiency. Despite the compelling set of potential benefits associated with greater accessibility and interconnectedness, our findings underscore the challenge of establishing liquidity demand and positive network externalities in the absence of established institutional mechanisms and clear regulatory frameworks.

In the second part of our analysis, we examine a different dimension of decentralization through its application to trading mechanisms. Our focus in particular is on the trading of crypto-native assets on Automated Market Maker (AMM) protocols.<sup>3</sup> This contributes to a growing body of literature on this fundamentally novel market structure (Lehar and Parlour, 2025; Capponi and Jia, 2021). We empirically compare the market quality differentials between AMMs and centralized limit order book (CLOB) exchanges and identify conditions under which natively decentralized mechanisms, might succeed. This comparison uncovers a significant apparent contradiction that forms the core contribution of this paper. On one hand, our direct empirical analysis reveals that transaction costs on AMMs often exhibit significantly higher baseline levels, while facilitating average trade sizes nearly three times larger than on CEXs. On the other hand, a counterfactual simulation, in which we route observed CEX trades through their corresponding AMM pools, holding the trade size and timing constant, suggests a substantial potential for cost savings. The median effective spread of 18.63 bps in the CEX environment contrasts markedly with a simulated AMM cost of just 3.09 bps.

The transaction cost gap identified through this counterfactual analysis is importantly not intended to reflect an unexploited efficiency gain to CEX traders. Rather, the simulation is a methodological tool designed to isolate and measure the differential efficiency of the price formation mechanisms themselves. The resulting gap is an empirical quantification of the price impact savings that would be required to overcome the unmodeled frictions that keep that order flow on the CEX. These factors include explicit costs from which our measure abstracts, such as fixed network fees that are disproportionately costly for small trades and variable AMM protocol trading fees, as well as unpriced risks inherent to the decentralized architecture, such as smart contract risk and the execution risk from predatory trading strategies like sandwich attacks.<sup>4</sup> This finding is therefore consistent with theoretical models of endogenous order flow segmentation (see Aoyagi and Ito (2025)), in which the coexistence of CEXs and DEXs represents an efficient equilibrium where heterogeneous traders sort across venues based on their sensitivities to different cost structures.

The paper additionally provides a granular empirical analysis of these competing market structures, revealing several additional key findings. Our trade-level analysis uncovers a distinctive U-shaped pattern in the transaction cost premium for AMMs across trade sizes, adding nuance to prior studies by revealing this pattern is characteristic of less liquid pools. In contrast, for high-liquidity pools, the cost premium increases monotonically with trade size. Most notably, we find that AMM execution costs are remarkably resilient to market volatility. While CEX spreads increase monotonically with volatility, the execution costs in well-capitalized AMM pools remain stable, and our multivariate analysis shows that the AMM cost advantage grows during periods of high volatility. This resilience is consistent with recent theoretical work; although conventional wisdom might suggest that volatility enhances liquidity withdrawal incentives. Lehar and Parlour (2025) and Foley et al. (2025) demonstrate that such tendencies can be counterbalanced by heightened liquidity provision incentives from increased fee generation during volatile periods.

This paper notably makes several contributions to the literature on market microstructure and digital finance. First, we contribute a comparative analysis of DeFi innovations, grounded in the central role of market frictions. By contrasting the failure of tokenized equities with the case of AMMs, we show that market design must be assessed on its ability to address fundamental challenges of liquidity formation and mitigate adverse selection within a given ecosystem. Second, we empirically identify and analyze the “transaction cost gap” as a tool for measuring the implicit, market-clearing price of DEX-specific constraints. By calculating the gap, the paper provides an empirical price of the combined cost of gas, MEV risk, and smart contract risk for small-order traders. Finally, we contribute a comprehensive, publicly available dataset of daily transaction cost metrics for a large sample of assets on both CEX and AMM platforms, providing a valuable resource for future research.<sup>5</sup>

Ultimately, our findings suggest that while not all decentralized innovations face an easy path to adoption, the architectural innovations of AMMs hold significant potential. Their demonstrated ability to provide resilient liquidity, particularly during periods

<sup>2</sup> Tokenized stocks are created through three main architectural models: custodial-backed tokens which represent a 1:1 claim on real shares held in reserve, synthetic derivatives which track the price of an asset without holding it, and natively issued tokens where the blockchain itself serves as the primary ownership ledger for the underlying firm. Each model presents a different risk profile and level of direct ownership, ranging from a claim on a custodied asset to purely economic exposure or, in the native model, the legally recognized share itself. Both (Nasdaq, 2025) and the (Intercontinental Exchange, 2026) have proposed custodial-backed tokenization models that fall within the custodial framework in recent SEC filings, though NYSE’s platform also supports natively issued tokens and offers 24/7 trading with instant settlement.

<sup>3</sup> As of 2025, decentralized exchanges and lending and borrowing platforms are the largest implementations of DeFi (Aspris and Svec, 2025).

<sup>4</sup> By isolating the implicit costs generated by the pricing mechanism itself, our measure provides a stable benchmark for efficiency. It avoids the analytical complications of trying to normalize for explicit fees that can vary dramatically based on network congestion, the use of main-chain versus layer-2 solutions, or the specific fee-tier chosen by a liquidity pool.

<sup>5</sup> Daily transaction cost data for AMMs and CEXs, respectively are made publicly available here: [http://h.cryptodatawarehouse.org/amm\\_dataset.csv](http://h.cryptodatawarehouse.org/amm_dataset.csv) and [http://h.cryptodatawarehouse.org/cex\\_dataset.csv](http://h.cryptodatawarehouse.org/cex_dataset.csv).

of high market volatility, suggests they can function as efficient and competitive alternatives to traditional market structures. The convergence of these distinct models presents a fertile ground for financial innovation and a challenge for regulatory frameworks, holding the potential to reshape global market functions from issuance to secondary trading.

Our analysis provides a nuanced response to the question of the superiority of financial architectures. We substantiate that the value of a decentralized market structure is critically contingent on its ability to solve for specific market constraints more efficiently than incumbent systems. Our primary contribution is the identification and economic interpretation of a persistent transaction cost gap between CEXs and AMMs. We argue this gap represents an equilibrium outcome, reflecting the price traders rationally pay to avoid the AMM architecture's unique frictions. We uncover a key structural advantage of the AMM model, namely its resilience to market volatility. Unlike CEXs, where spreads widen predictably in turbulent conditions, we find that execution costs in well-capitalized AMMs remain generally stable, suggesting they can serve as a valuable source of liquidity precisely when it is most scarce. These findings challenge a simplistic view of AMMs as merely less efficient than traditional structures and open specific avenues for future research, including the potential application of AMM protocols in traditionally illiquid, dealer-dominated markets.

## 2. Digital assets and markets

This section develops the framework for our empirical analysis by examining two distinct applications of decentralized financial technology. We first review the literature on security design to evaluate the theoretical merits and inherent risks of tokenization. We then turn to the competition between centralized limit order book exchanges and automated market maker protocols. This literature provides the foundation for understanding how different market designs offer competing solutions to the problem of liquidity provision. Together, these reviews establish the basis for our empirical investigation into the conditions under which decentralized architectures can deliver better market outcomes.

### 2.1. Security design and tokenization

Financial innovation in security design is fundamentally a response to market incompleteness and frictions (Allen and Gale, 1988). The development of novel financial products is often driven by market segmentation, where distinct investor preferences and limited participation create incentives for new, endogenous asset structures. Such innovations can enhance market efficiency through improved information incorporation and risk-sharing, or cater to specific investor demands, such as the yield-seeking appetite of retail investors (C el erier and Vall e, 2017), the demand for safe assets (Coval et al., 2009), or aversion to specific risks (Calvet et al., 2023).

The tokenization of real-world assets represents a modern manifestation of this innovation drive. The capability to exchange tokens whose value aligns with securities custodied off-chain gained traction in the early 2020s, with the development of products designed to provide investors with digital representations of company shares. Markets to trade tokenized versions of popular U.S.-listed securities, such as Google, Netflix, and Tesla, gained traction through the promise of greater accessibility and reduced market frictions for a global, crypto-native investor base.

The architecture of these instruments, however, can be categorized into several distinct models, each with a different risk profile and relationship to the underlying asset. The most common structure is the custodial-backed or wrapped token, where a third-party institution holds the physical shares and issues a corresponding digital token representing a one-for-one claim on the custodied asset. These tokens are representations of claims on financial equities recorded on the blockchain, which require a means to store or provide custody for the reference asset, as well as a mechanism to facilitate redemption (Kreppmeier et al., 2023; Carapella et al., 2023).<sup>6</sup> Alternative structures include synthetic tokens, which are on-chain derivatives that provide purely economic exposure by tracking the price of a stock without holding the asset in reserve, and natively issued tokens, where a company issues its shares directly on a blockchain, making the token the primary, legally recognized share. By aiming to mirror the economic properties of the underlying security, including entitlements to dividends, these tokens provide economic exposure, though the extent to which they confer direct ownership rights varies significantly by their architectural design.

Regardless of the tokenization model, the potential benefits of tokenizing stocks are compelling. The technology could support market liquidity, provide interoperability across platforms, and enable rapid, near-instantaneous settlement (Momtaz, 2023). However, these potential advantages are not without countervailing risks. Concerns have been raised regarding the financial stability implications of mass redemptions from tokenized asset pools (Carapella et al., 2023), and the benefits of eliminating settlement periods may be overstated, as this could impair the ability of investors to manage margin or leverage, potentially increasing systemic risk by removing the netting benefits of longer settlement cycles. Furthermore, as Malinova and Park (2023a) caution, generic token structures can be inferior to traditional equity and result in significant production inefficiencies if not well designed.

This tension between the promise of tokenization and its practical challenges motivates the first part of our paper. The fact that these tokenized equities utilize the same mechanism as their traditional counterparts (CLOB), provides a useful analytical separation. We are among the first to trace the empirical performance of this market, with a particular focus on trading costs and price efficiency. This analysis serves as a crucial case study of the innovation challenges at the nexus of traditional capital markets and distributed ledger technology.

<sup>6</sup> The veracity of FTX's claim that its tokenized-equity products were fully backed by custodied shares is unresolved. Initial disclosures referenced one-for-one custody via CM-Equity AG (a relationship apparently terminated by December 2021, shortly after Binance's July 2021 withdrawal from the segment). After this FTX purportedly routed settlements through Swiss-based Canco GmbH (later "FTX Switzerland"), yet no conclusive documentation has surfaced following the exchange's November 2022 collapse.

## 2.2. Competing for liquidity: Centralized versus decentralized exchanges

Considering the widespread convergence towards the limit order book structure in nearly all globally recognized asset markets over recent decades, the success of the AMM model for trading crypto-assets is noteworthy.<sup>7</sup> The debate to address which market structure provides better execution is part of a comprehensive literature on market design. Early theoretical studies endorse the limit-order book for its efficiency in providing liquidity (Glosten, 1994; Biais et al., 2000; Parlour and Seppi, 2003). Notably, Glosten (1994) demonstrates that under competitive conditions, an electronic open limit order book can be a highly efficient and resilient mechanism for liquidity provision, proving robust to competition from alternative anonymous market structures such as dealer markets. This finding is supported by subsequent work showing that limit order markets can supply sufficient liquidity to discourage the exploration of alternative designs (Biais et al., 2000; Parlour and Seppi, 2003). However, these results often assume environments with limited frictions. Parlour and Seppi (2008) show that the competition-proof nature of the LOB may not be guaranteed in the presence of significant non-informational market frictions, a condition particularly relevant for cryptocurrency markets where technical constraints and transaction costs create substantial impediments to trade.

The AMM model represents a fundamental departure from traditional market making, drawing its origins from prediction markets and market scoring rules (Milionis et al., 2022). Its prominence in the DeFi ecosystem stems from its ability to circumvent the high on-chain costs and latency issues that plague decentralized implementations of the CLOB (Aspris et al., 2021). This innovation has spurred significant academic interest in whether AMMs can effectively compete with established centralized structures.

Recent literature presents a nuanced picture of this competition. Theoretical work by Lehar and Parlour (2025) characterizes the equilibrium in an AMM as a trade-off between fee revenue for liquidity providers (LPs) and the adverse selection costs from being “picked off” by informed arbitrageurs. Their model suggests AMMs may be optimal under specific conditions, particularly when trading volumes are low and asset volatility is moderate. Complementing this, Capponi and Jia (2021) show that high volatility can exacerbate these arbitrage losses to the point of inducing a “liquidity freeze”. From an empirical standpoint, Barbon and Ranaldo (2026) find that explicit trading frictions, particularly network gas fees, are a primary determinant of relative costs, making AMMs cost-effective mainly for large transactions that can amortize these fixed costs.

The coexistence of these distinct market structures also raises important welfare questions. Aoyagi and Ito (2025) model the endogenous choice between CEXs and DEXs, demonstrating that their coexistence can be an efficient market outcome, as different types of traders sort across venues. Extending this logic, recent work has begun to explore the viability of AMMs in traditional finance. Malinova and Park (2024), for instance, find that AMMs could generate substantial transaction cost savings in equity markets, particularly for less liquid securities where traditional market making is less efficient.

Our analysis contributes to and extends this literature in several ways. First, by conducting a detailed transaction cost analysis, we isolate the costs inherent to the competing market mechanisms. Second, a counterfactual simulation, in which we route observed CEX trades through their corresponding AMM pools, holding the trade size and timing constant, allows us to isolate and quantify the cost component arising purely from the price formation mechanism, as well as to measure the implicit, market-clearing price of DEX-specific frictions. Finally, we offer a more granular empirical analysis of how transaction costs vary with market conditions, particularly the interaction among trade size, liquidity, and volatility, providing a more comprehensive assessment of the relative performance of these two market structures.

## 3. Data and metrics

### 3.1. Data source and sample selection

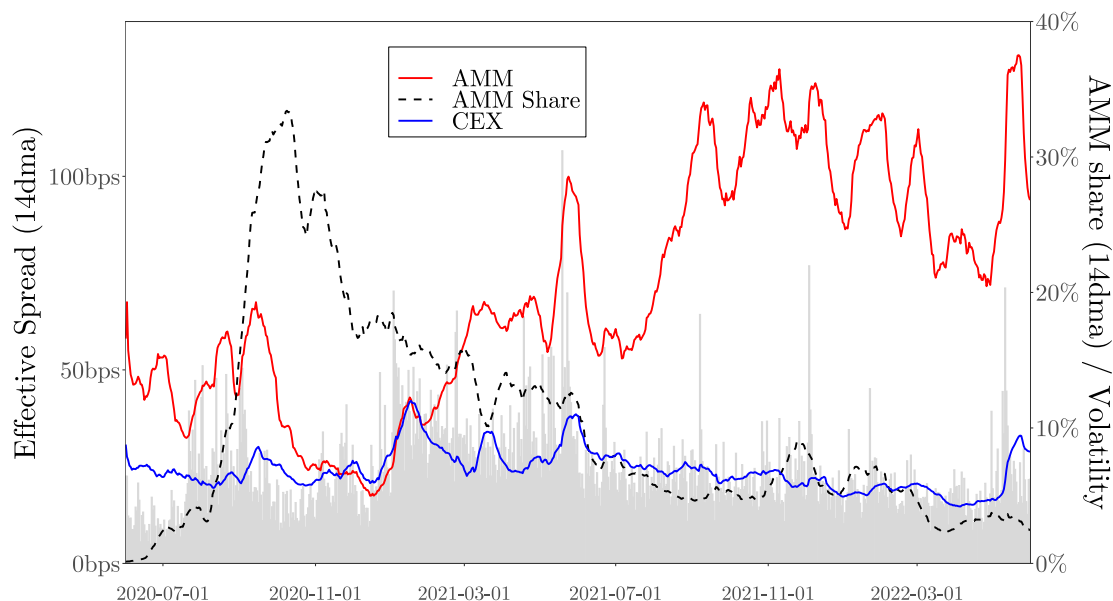
Our analysis combines several data sources for the examination of trading costs across various market structures. We extract centralized cryptocurrency exchange data directly from Tardis.dev (“Tardis”) for the Binance exchange. Binance is the market leader among centralized exchanges with a 66.7% market share as of December 2022.<sup>8</sup> From this same source we are able to extract trade and quote data for tokenized equity instruments traded on the FTX exchange, the principal platform for trading tokenized stocks prior to its collapse in late 2022. The data from Tardis provides the timestamp of each trade (to the millisecond), as well as trade direction (buyer or seller initiated), price, and volume. We also source the best-bid and ask quotes from these markets, which are supplied with every quote revision.

For the examination of the AMM, we extract data directly from the Ethereum blockchain and parse transaction and trade information relating to the Uniswap V2 protocol. Our data identifies the assets and amounts traded within a swap (input/output token), block number and timestamp (of block confirmation), as well as the sender (from address) and receiver (to address) associated with each transaction. The ratio of Asset A traded for Asset B provides an implicit ‘price’ for each transaction. In the case of Uniswap V2, the protocol also tracks the liquidity pool reserves after each state-changing transaction.

To construct our sample, we query the blockchain indexing protocol, TheGraph, and rank liquidity pools by the US\$ market value of their trading volume during the three-month period February 2021 to May 2021. This time window corresponds to a period of

<sup>7</sup> It should be noted that the LOB structure adopted across most jurisdictions is complemented by other important mechanisms, such as call auctions for market openings and closings and the significant volume executed in non-transparent venues like dark pools.

<sup>8</sup> <https://www.coindesk.com/markets/2023/01/04/binance-led-market-share-in-2022-despite-overall-decline-in-cex-volumes>



**Fig. 1. Timeseries transaction costs.** This figure illustrates the evolution of cross-sectional median effective spreads of the AMM (red line) and CEX (blue line). The black dotted line shows the AMM volume share across all asset pairs and the gray bars represent the daily median pair-wise volatility. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

significant growth on Uniswap V2, preceding the introduction of Uniswap V3.<sup>9</sup> We identify trading pairs common to Uniswap and Binance over the period and select 50 of the most liquid token pairs for consideration.<sup>10</sup> Our sample includes a mix of token-to-stablecoin, token-to-token, and stablecoin-to-stablecoin pairs and is designed to capture a broad and representative snapshot of the most actively traded markets on Uniswap V2 during our period of study. Our final sample consists of 13,520,046 transactions on Uniswap and 721,687,027 transactions from Binance for the period 2020-06-01 to 2022-05-31.

This sample covers a period of significant innovation and growth where AMMs and CEXs held meaningful market share across a wide variety of token types (see Fig. 1). The sample period captures various market conditions, offering notable examples of both bull and bear markets. It also includes periods of extreme volatility and relative stability, equally represented in both the equity and cryptocurrency markets. We further use equity market data to benchmark execution costs and trade activity against their tokenized counterparts. This data is retrieved from Refinitiv (TR), for 39 securities for which there were tokenized representations traded on FTX. The equity dataset contains over a billion trades in this time frame dwarfing the 1.25 million tokenized asset trades. We impose several data filters across the different samples. For our CEX and traditional equity samples, we aggregate trade-level data to market orders by combining trades with the same timestamp and direction and calculate a volume-weighted price following the mechanisms outlined in Foley et al. (2023). We further remove observations where the bid/ask price or size are missing, equal to zero, or where buy (sell) orders are executed below (above) the prevailing midquote. The direction of trade is inferred using the (Lee and Ready, 1991) algorithm. For the AMM, we exclude trades that consume more than 10% of the respective token reserve. Unless otherwise indicated, for our analyses we aggregate the data to daily intervals which avoids microstructure noise.

### 3.2. Metrics

The primary objective of our transaction cost analysis is to evaluate the inherent efficiency of the competing price formation mechanisms within the different market structures. To achieve this, we concentrate on the *implicit costs* of trading — those costs that arise directly from the price discovery process. This deliberate focus allows for a direct comparison of the pricing functions of the two market structures, independent of other frictions.

<sup>9</sup> The strategic goal is to construct a static, representative sample of the most important and actively traded pairs on the Uniswap V2 protocol. This specific window represents the high point of V2 market relevance in terms of trading volume, liquidity, and user activity. It concludes immediately prior to the launch of Uniswap V3 on May 5, 2021, an event that, as discussed, began to fragment the AMM liquidity landscape.

<sup>10</sup> In choosing the 50 most liquid overlapping assets, we seek to balance breadth of coverage with active trading and liquidity to ensure our results are both fair and generalizable.

At the trade level ( $t$ ), the effective spread, realized spread, and price impact are decomposed as follows:

$$\underbrace{q_t(p_t - m_t^\tau)}_{\text{Effective Spread}} = \underbrace{q_t(p_t - m_{t+\Delta}^\tau)}_{\text{Realized Spread}} + \underbrace{q_t(m_{t+\Delta}^\tau - m_t^\tau)}_{\text{Price Impact}} \quad (1)$$

where  $\tau$  indicates whether the trade is from a centralized exchange (CEX), an Automated Market Maker (AMM), a traditional equities order-book market (EQX) or tokenized equities order-book market (TEQX). These metrics are expressed in basis points, where  $p_t$  is the transaction price,  $m_t^\tau$  is the prevailing midpoint at the time of the trade and  $m_{t+\Delta}^\tau$  is the midpoint at a reference time (30 s) after the trade. The trade direction  $q_t$  is equal to (+1) for buyer-initiated and (−1) for seller-initiated trades. We adopt a 30-second post-trade horizon for our price impact and realized spread calculations, consistent with established microstructure studies in traditional, continuous-time markets. A single clock-based window ensures that a consistent and direct methodology is applied across both the CEX and AMM samples within our dataset, allowing for a comparable analysis of their relative price dynamics.<sup>11</sup>

Unlike centralized exchanges, AMMs do not display quotes directly. However, the AMM reserves imply a midpoint at which the corresponding liquidity pool stands ready to buy or sell an infinitesimal amount of the respective token. Let  $y$  be the reserves of asset Y and  $x$  the reserves of asset X, then the midpoint is simply the ratio of the corresponding reserves:  $Midpoint_{AMM,xy} = \frac{y}{x}$ . The transaction price and direction can be inferred from the amounts that are sent to and received from the liquidity pool. A trade swapping  $\Delta x$  for  $\Delta y$  implies a transaction price of  $Price_{xy} = \frac{\Delta y}{\Delta x}$  where  $\Delta x$  is the amount of asset X sent to the liquidity pool and  $\Delta y$  the amount of asset Y that is returned to the trader. The signed amount of the traded quantities allows for the identification of a trade's direction. Despite the fundamental differences in market mechanics, the measures composed are comparable and represent the costs incurred by participants in the transaction process.

It is important to note that our effective spread calculations focus on the price impact component of trading costs, excluding the explicit fees charged by AMM protocols. Unlike centralized exchanges where maker-taker fees are typically separate from the bid–ask spread, AMM fees (typically 0.3% for Uniswap V2) are embedded within the trading mechanism and paid directly to liquidity providers. These fees represent a fundamental component of the AMM business model, designed to compensate liquidity providers for the adverse selection risk they bear (Hasbrouck et al., 2026). While our analysis isolates the costs arising from the market's pricing mechanism itself, a comprehensive assessment of total trading costs would necessarily include these protocol fees alongside the price impact measures we examine. Our analysis also excludes network gas fees required for on-chain transaction execution.

While the effective spread and price impact maintain their standard interpretations as the cost of demanding liquidity and the adverse selection component of the spread, respectively, the interpretation of the *realized spread* must be fundamentally adapted for the AMM context. In traditional markets with active market makers, the realized spread is a proxy for liquidity provider profitability. However, the defining feature of an AMM is that liquidity provision is passive; LPs commit capital to a deterministic pricing curve and cannot dynamically adjust quotes in response to new information. As a result, they are systematically exposed to adverse selection from informed arbitrageurs who exploit price discrepancies between the AMM and external reference markets (Lehar and Parlour, 2025). Accordingly, within an AMM the realized spread should be viewed not as a measure of LP profitability but as an empirical gauge of the adverse-selection loss that LPs incur (“impermanent loss”).<sup>12</sup>

#### 4. Trading costs across assets, markets, and time

Our empirical analysis examines trading costs across a diverse range of tokenized assets and market structures. Specifically, we focus on two sets of comparisons: tokenized and traditional equities; and the comparison of tokens and cryptocurrencies traded on CEX and AMMs. The comprehensive instrument sample is detailed in Appendix A.1.

##### 4.1. Tokenized vs centrally registered equities

Table 1 presents summary statistics comparing tokenized equities with their traditional counterparts traded in equity markets. The composition of tokenized equities, detailed in Appendix A.1, reveals a distinct concentration in growth sectors and a notable presence of ‘meme’ stocks. These meme stocks are typically perceived as overvalued based on conventional financial metrics, and the meme retail investors, often considered pro-tech and active on social networks, are likely to also trade in cryptocurrencies (Hasso et al., 2022; Yousaf et al., 2023).

The results show that the emergent tokenized equities are traded far less actively than conventional equities, averaging only 66 trades per security per day versus roughly 47,561 daily trades in traditional equity markets. In terms of transaction costs, the differential between the two markets is substantial. Our analysis shows that spread metrics and price impact for tokenized equities are approximately ten times greater, at 51.42 basis points, despite both markets operating under the same centralized limit-order book structure.

<sup>11</sup> For AMM swaps, the reference midpoint is taken from the first pool state where block time-stamp is greater than  $t + 30$  s. As Ethereum blocks arrive roughly every 13 s, this horizon normally spans two to three block updates, the interval in which most of the permanent price impact in DeFi manifests. Thus, although AMM prices move discretely, the 30-second window still captures the immediate post-trade price discovery process while preserving comparability with centralized exchanges.

<sup>12</sup> We refer to Aquilina et al. (2024a) for a comprehensive review on LP profitability in AMMs.

**Table 1**

Descriptive Statistics: Tokenized Equities. This table reports summary statistics of the analyzed (tokenized) equity sample. Annualized return is the mean annualized return computed using daily log-returns across all asset pairs. Market Cap is the daily market capitalization of the stock in 100 million USD. Volatility is the daily high-low range, scaled by its midpoint, expressed as percentage. Trade size is the daily mean trade size in USD. Trade Count is the total count of trades. Daily Volume is the sum of daily trading volume in million USD. Depth is the time-weighted depth at the best bid and ask in USD. Spread and Price impact measures are daily volume-weighted.

Variable	Tokenized			Equities		
	Mean	Median	SD	Mean	Median	SD
Annualized Return (%)	-40.22	-47.78	41.31	-40.32	-48.03	41.26
Market Cap (USD)	2937.91	582.15	5452.70	2937.91	582.15	5452.70
Volatility (%)	0.32	0.18	0.57	1.06	0.91	0.87
Trade Size (USD)	372.26	30.47	2314.10	7732.72	5775.50	7773.96
Trade Count	65.73	20	196.56	47,561.34	11 155	88,664.07
Daily Volume (USD)	61.86	4.63	231.11	1,694,058.38	155,190.43	3,904,867.06
Depth (USD)	67,244.46	31,679.26	103,329.54	1608.28	781.94	2459.87
Effective Spread (bps)	51.42	29.44	76.57	5.48	2.93	6.93
Realized Spread (bps)	19.28	9.71	81.59	1.84	1.06	3.41
Price Impact (bps)	32.93	17.10	67.73	3.70	1.68	6.09

The high cost of liquidity is associated with relatively small trades, with the average trade size equivalent to US\$372.26, which is a fraction of the average trade size recorded in US equities (US\$7732.72) for the same securities. These results are perhaps unsurprising, given the unclear regulatory status of tokenized equities, which increases their risk premiums relative to traditional forms. Additionally, holders of these tokenized equities faced higher levels of counterparty risk, limited portability of their inventory, and an absence of the traditional protections afforded to investors in regulated traditional markets.<sup>13</sup>

Tokenized equities display sizable and persistent deviations relative to their underlying-market counterpart as shown in [Appendix A.2](#). To examine whether these discrepancies reflect slower information incorporation, we estimate the security-level price-delay statistic of [Hou and Moskowitz \(2005\)](#) and present cross-sectional results in [Table 2](#). The dependent variable is the average price delay for each security, and the key regressor *Tokenized* equals one for tokenized representations and zero for the underlying shares. The coefficient is positive and highly significant, indicating that market-wide information is impounded into tokenized equity prices at a materially slower rate than into the reference securities. Furthermore, our model identifies a separate, additive effect for non-regular trading days, where price delay increases for both tokenized and traditional equities. This is consistent with the additional frictions in price discovery that arise when the primary market for the underlying security is closed. In contrast, price delay is unrelated to the illiquidity measure, suggesting that conventional liquidity considerations do not explain the observed inefficiency. Instead, factors such as exchange-level latency, jurisdictional trading restrictions (U.S. access to FTX equity tokens was barred), and the absence of governance rights in the token specifications are more plausible impediments to timely information integration.

The promised benefits of tokenized equities global reach, round-the-clock trading, fractional ownership, and near-instantaneous settlement remain unrealized. Our analysis reveals instead a market with low trading volumes, uncompetitive transaction costs, and material lags in information incorporation. These findings highlight that even with a familiar limit-order book structure, more than technological novelty is required to overcome foundational frictions related to liquidity and informational efficiency in the area of non-native tokenized assets. For the remainder of this paper, we confine our examination of decentralized trading to native tokens traded between centralized and decentralized exchanges.

#### 4.2. Centralized exchanges vs automated market makers

[Table 3](#) reports summary statistics on the sample of tokens traded concurrently on the CEX and AMM for the period from June 2020 to June 2022. Our univariate analysis yields two key results regarding market activity and structure: first, CEXs exhibit higher trading frequency, and second, AMM transactions exhibit substantially larger average trade sizes than trades executed on centralized exchanges.

Over the sample period, the mean (median) CEX daily trade volume is \$53.2 million (\$0.98 million) compared to \$5.54 million (\$0.16 million) on the AMM. For the average number of daily trades, activity on the CEX exceeds the AMM by a factor of 58 times. This transaction volume gap is comparable to [Barbon and Ranaldo \(2026\)](#) and may be explained by a variety of factors, including the advanced trading features available on CEXs, such as margin trading and stop-loss orders, along with faster execution speeds and familiarity with limit-order book systems. Additionally, differences in stablecoin trading activity contribute to the observed trends. CEXs notably experience higher stablecoin trade activity due to the fee and liquidity structure of early iterations of decentralized protocols like Uniswap V2. The sample period also covers the launch of Uniswap V3, which resulted in a migration of trade activity away from V2, further widening the trade activity gap between the CEX and AMM.

<sup>13</sup> While still in its infancy, the demand for regulation of decentralized markets and DeFi as a whole is increasing, especially due to the growing interconnectedness between traditional and decentralized financial systems ([Aquilina et al., 2024b](#); [Espancia et al., 2023](#)).

**Table 2**

Equity price delays. This table reports the results of a fixed effects regression evaluating the differences in price delays across tokenized and traditional equity markets. The dependent variable price delay is as described in Hou and Moskowitz (2005). The metric is computed separately for equities and tokenized equities by estimating the following two regression models using daily data: (1)  $R_{i,t} = \alpha + \beta_1 R_{m,t} + \epsilon_{i,t}$  and (2)  $R_{i,t} = \alpha + \beta_1 R_{m,t} + \beta_2 R_{m,t-1} + \beta_3 R_{m,t-2} + \beta_4 R_{m,t-3} + \beta_5 R_{m,t-4} + \epsilon_{i,t}$ , where  $R_{i,t}$  is the daily log return of the respective tokenized or traditionally-traded stock  $i$  on day  $t$  and the market return  $R_{m,t}$  being the daily equally-weighted log return of the traditional equity sample. After, the price delay metric is calculated as  $PriceDelay_{i,t} = 1 - \frac{R_{i,t}^{Reg1}}{R_{i,t}^{Reg2}}$ . Tokenized is a binary variable taking the value of 1 if the observation belongs to the tokenized equity market. Amihud is a version of Amihud's illiquidity measure and calculated as return over daily trading volume. OutsideHours is a binary variable taking the value of 1 for observations that occur outside regular trading days, e.g. holidays or weekends. Standard errors reported in parentheses are clustered at the asset pair level and \*\*\*, \*\*, \* denote the statistical significance at the 1, 5, and 10% level, respectively.

	Price delay			
	(1)	(2)	(3)	(4)
Tokenized	0.0777*** (0.0082)	0.0773*** (0.0081)	0.0768*** (0.0085)	0.0764*** (0.0084)
Amihud		1.315 (0.955)		1.272 (0.958)
OutsideHours			0.0469** (0.0204)	0.0470** (0.0204)
Fixed Effects	Stock, Day	Stock, Day	Stock, Day	Stock, Day
N	35,831	35,831	34,416	34,416
Adj. R2	24.26%	24.27%	23.77%	23.78%

**Table 3**

Descriptive Statistics: AMM and CEX Tokens. This table reports summary statistics of the analyzed cryptocurrency sample. Annualized return is the mean annualized return computed using daily log-returns across all asset pairs. Market Cap is the daily market capitalization of the smaller asset of a trading pair in 100 million USD. Volatility is the daily high-low range, scaled by its midpoint, expressed as percentage. Trade size is the daily mean trade size in USD. Trade Count is the daily total count of trades. Daily Volume is the sum of daily trading volume in million USD. TVL is the total-value locked for the AMM in million USD. Depth is the time-weighted depth at the best bid and ask in USD. Spread and Price impact measures are daily medians.

Variable	AMM			CEX		
	Mean	Median	SD	Mean	Median	SD
Annualized Return (%)	24.54	-9.51	176.37	30.81	-10.08	192.73
Market Cap (USD)	148.22	11.13	548.20	148.22	11.13	548.20
Volatility (%)	10.20	7.76	10.60	7.89	6.16	7.48
Trade Size (USD)	2507.26	1484.12	3247.05	885.60	483.95	1183.66
Trade Count	548.04	59	1621.69	31,557.07	2297.50	100,390.50
Daily Volume (USD)	5.54	0.16	21.81	53.15	0.98	196.04
TVL (USD)/Depth (USD)	30.71	1.53	81.87	393,100.49	2.50	2,433,462.85
Effective Spread (bps)	110.22	46.12	184.76	24.59	18.57	23.85
Realized Spread (bps)	-31.15	27.55	191.70	12.19	5.81	18.78
Price Impact (bps)	192.76	76.08	353.24	11.92	8.89	13.88

In relation to trade size, AMMs facilitate larger trades on average, with mean (median) trade sizes of \$2507 (\$1484) approximately 2.8 (3.1) times greater than on CEXs. This result is consistent with the notion that high fixed explicit transaction costs associated with AMMs, such as gas fees, encourage traders to consolidate their orders to amortize these costs over a larger trade value. Further, the deterministic nature of the AMM pricing function provides traders with greater certainty regarding the price impact of large trades, in contrast to the uncertainty of 'walking the book' in a CLOB market.

A comparison of the implicit cost metrics in Table 3 reveals the fundamental structural differences between the two venues. The mean (median) effective spread is markedly lower on CEXs, at 24.6 (18.6) basis points, compared to 110.2 (46.1) basis points on the AMM.<sup>14</sup> This differential is primarily driven by the large price impact inherent to the AMM constant-product pricing function, which, as demonstrated by Capponi and Jia (2021), exhibits greater convexity than the linear pricing schedules often found in CLOBs. The decomposition of the effective spread provides a deeper insight into the economics of AMM liquidity provision. The combination of a high price impact (mean 192.76 bps) and a large, negative realized spread (mean -31.15 bps) is a key finding of our study. A negative realized spread is consistent with the adverse-selection risk faced by passive liquidity providers in AMMs. This represents the value captured by informed arbitrageurs who trade against stale AMM prices to align them with external market prices. Our

<sup>14</sup> This sample heterogeneity presents certain challenges for transaction cost normalization and interpretation. Standard transaction cost practice favors restricting analysis to assets priced against a common fiat-pegged numeraire (e.g., token-to-stablecoin pools), as this allows for a more direct and unambiguous interpretation of costs measured in basis points. In our mixed sample, particularly for token-to-token pairs, the calculation of basis point costs is inherently sensitive to the choice of the numeraire asset. To mitigate the influence of this heterogeneity on our multivariate results, we include asset-pair fixed effects, which absorb time-invariant, pair-specific characteristics, including the identity of the numeraire in our multivariate analysis.

results provide a large-sample quantification of this cost, demonstrating that the revenue from price impact on AMMs is captured not by the passive LPs, but by the active arbitrageurs who enforce price efficiency. This dynamic is a fundamental distinction from CLOB markets, where active market makers can more effectively manage adverse selection risk and retain a larger portion of the price impact as revenue.

#### 4.3. Trading costs by asset characteristics

We examine the cross-sectional variation in trading costs across tokens by conditioning on several key market characteristics: total value locked (TVL), market depth, volatility, trading volume, and market capitalization. For each trading day, we categorize asset pairs into quintile portfolios based on these characteristics, calculating equally-weighted averages of transaction cost metrics for each portfolio and day. The results, presented in Table 4, include the high-minus-low (Q5-Q1) portfolio differences and their corresponding test statistics in the bottom row of each panel, capturing the economic and statistical significance of the relationship between trading costs and each characteristic.

Panel A sorts portfolios by AMM liquidity reserve (TVL) and CEX depth. Consistent with theoretical predictions on the inverse relationship between trading costs and liquidity provision, we document substantial variation in effective spreads across depth quintiles (Copeland and Galai, 1983; Lehar and Parlour, 2025). For AMMs, we observe a striking 294 basis point differential in effective spreads between the least and most liquid quintiles. This relationship is highly convex, with costs declining rapidly as liquidity increases. Notably, our findings indicate that the most liquid AMM pools (quintile 5) achieve implicit trading costs (2.45 bps) that are comparable to, and even slightly lower than, their most liquid CEX counterparts (3.16 bps), highlighting the competitive potential of well-capitalized AMMs.<sup>15</sup>

Panel B presents our analysis of portfolios sorted by asset volatility, measured as the daily price range normalized by the midpoint price for each trading pair. For CEX portfolios, we document a monotonic increase in transaction costs within a tight band across volatility ranks, with effective spreads and price impact measures displaying differentials of 14.4 and 10.0 basis points, respectively, between extreme quintiles. The AMM portfolios, however, exhibit a distinctive U-shaped pattern in trading costs. While the highest and lowest volatility quintiles show comparable transaction costs, quintiles 2 through 5 demonstrate a positive but more moderate cost gradient. This non-linear relationship can be attributed to two structural features of AMM markets. First, stablecoin pairs in the lowest volatility quintile incur higher fixed execution costs in AMM protocols relative to centralized venues. Second, the elevated trading costs in lower volatility portfolios result from asset pairs with reduced trading activity that trade infrequently. These are characterized by limited liquidity provision and increased price staleness, all of which contribute to the observed U-shaped cost structure.

In Panel C, we document a strong inverse relationship between trading volume and transaction costs across both market structures. The economic magnitude of this relationship is substantial, with effective spreads declining by more than 95% from the lowest to highest volume quintiles in both CEX and AMM venues. Panel D extends these findings by examining the relationship between market capitalization and trading costs, revealing that larger-capitalization tokens exhibit systematically lower spreads, consistent with empirical patterns documented across other asset classes. The elevated trading costs observed in small-capitalization tokens reflect higher fixed costs of intermediation that manifest in wider spreads. Notably, our analysis reveals that transaction costs in centralized cryptocurrency exchanges, and the most liquid AMM pools, achieve levels comparable to those documented in traditional equity markets (Hendershott et al., 2011).

#### 4.4. Time-series variation

Our empirical analysis reveals substantial temporal variation in trading cost metrics across both market structures throughout the sample period. Fig. 1 illustrates the evolution of cross-sectional median effective spreads, with the black dotted line denoting AMM volume share across asset pairs and gray bars representing daily median pairwise volatility. Consistent with the summary statistics presented in Table 3, AMMs exhibit systematically higher effective spreads.

Several distinctive patterns emerge from our analysis. First, the temporal dispersion of transaction costs demonstrates markedly lower variation in the CEX environment (blue line) relative to the AMM setting (red line). Second, our empirical estimates reveal that transaction costs rise sharply with contemporaneous market volatility, and the elasticity of this relationship is materially steeper in the AMM venue than on CEXs. Third, after the initial three-month post-launch window, and amid subdued market volatility, the transaction-cost differential between AMMs and CEXs narrowed steadily, reaching statistical parity by December 2020. This convergence phase was accompanied by a pronounced expansion in AMM market share. From 2021 onward, however, AMM trading costs exhibit a sustained upward trend attributable to two primary forces: the surge in market volatility during the 2021 meme-coin and NFT boom, and a structural shift in the AMM landscape following the introduction of Uniswap V3 in May 2021. The upgrade, which implemented concentrated liquidity ranges and multiple fee tiers, induced a substantial migration of capital away from our benchmark venue, Uniswap V2—particularly in stablecoin pools and other high-turnover pairs.<sup>16</sup> The resulting contraction in AMM market share is reflected in a marked rise in both quoted and effective trading costs.

<sup>15</sup> The AMM result excludes the 30 bps liquidity-provider fee embedded in pool reserves. Incorporating that explicit component raises the all-in AMM spread to 32.45 bps, a level that remains competitive for mid-sized trades (US\$50k) but dominates CEX execution only under fee waiver regimes or when gas prices are unusually low.

<sup>16</sup> For a detailed comparison of Uniswap V2 and V3 architectures and the ensuing liquidity migration, see Foley et al. (2025).

**Table 4**

Descriptive Statistics: Portfolios. This table reports summary statistics of portfolios based on multiple asset pair characteristics. For each day, we classify asset pairs according to the respective characteristic into five portfolios from low (1) to high (5). After, we compute equal averages across days that we show in the table. The second last row of each panel shows differences in means between the fifth (highest) and first (lowest) portfolio. The last row reports the *t*-statistic of a *t*-test testing for differences in means between these two groups. Effective Spread and Price Impact are in bps, Volatility in percentage, Volume in USD thousands, AMM TVL in USD millions and Depth in USD. \*, \*\*, and \*\*\* denote statistical significance at the 10%, 5%, and 1% levels, respectively.

Portfolio	Effective Spread (AMM)	Effective Spread (CEX)	Price Impact (AMM)	Price Impact (CEX)	TVL AMM (USD)	Depth CEX (USD)
Panel A: Sorted on TVL/Depth						
1	296.21	38.17	526.09	18.11	0.28	0.85
2	129.11	30.32	213.99	14.66	0.88	1.66
3	61.59	26.60	105.96	13.59	2.52	16.09
4	35.63	21.90	67.53	9.75	17.40	2739.79
5	2.45	3.16	4.86	2.00	143.99	2,120,496.69
(5-1)	-293.76***	-35.02***	-521.23***	-16.10***	143.71***	2,120,495.84***
	(-65.2)	(-90.52)	(-57.39)	(-77.01)	(66.6)	(24.66)
Panel B: Sorted on Volatility						
1	137.46	16.32	266.70	7.23	3.05	2.22
2	74.66	24.18	132.16	11.14	6.39	5.19
3	80.91	26.12	139.77	11.73	8.64	6.99
4	108.97	26.92	179.60	13.03	11.65	9.46
5	151.42	30.77	245.75	17.24	23.08	16.94
(5-1)	13.96***	14.44***	-20.95**	10.01***	20.04***	14.72***
	(2.79)	(28.05)	(-2.13)	(29.57)	(73.09)	(79.00)
Panel C: Sorted on Volume						
1	262.57	43.69	481.82	18.72	0.05	0.16
2	131.98	37.00	219.86	17.40	0.14	0.47
3	78.78	25.47	128.86	12.61	0.40	1.75
4	47.52	11.07	78.95	7.50	2.13	15.77
5	7.17	1.58	12.75	1.57	27.13	268.44
(5-1)	-255.41***	-42.11***	-469.08***	-17.16***	27.08***	268.28***
	(-57.19)	(-129.72)	(-51.81)	(-84.66)	(37.73)	(42.87)
Panel D: Sorted on Market Capitalization						
1	193.57	42.59	352.41	20.26	112.66	
2	159.28	39.68	273.36	15.90	502.38	
3	92.22	24.32	152.67	13.22	1430.26	
4	59.44	9.58	104.05	6.02	5467.84	
5	29.20	2.59	50.13	2.34	72,172.26	
(5-1)	-164.37***	-40.00***	-302.28***	-17.92***	72,059.59***	
	(-43.44)	(-102.4)	(-41.04)	(-71.99)	(40.45)	

## 5. Determinants of trading costs and comparative analysis

### 5.1. What drives trading costs?

In this section, we analyze the determinants of transaction cost measures for tokens traded on the AMM and CEX. We examine how market and asset-pair characteristics impact these costs and develop a simple regression model to describe them. We regress trading cost components for each token-pair *i*, day *t* across markets *j* in our sample using the specification below:

$$\begin{aligned} \text{TxCost}_{i,j,t} = & \beta_0 + \beta_1 \text{InversePrice}_{i,j,t} + \beta_2 \text{DailyVolume}_{i,j,t} + \beta_3 (\text{TVL}_{i,j,t} \times \text{AMM}_j) \\ & + \beta_4 \text{Volatility}_{i,j,t} + \beta_5 \text{TS}_{i,j,t} + \beta_6 \text{TS}_{i,j,t}^2 + \beta_7 \text{Stablecoin}_i + \beta_8 V3_t + \epsilon_{i,j,t} \end{aligned} \quad (2)$$

where TxCost is the transaction cost variable of interest, the natural logarithm of the effective spread, price impact, and realized spread, for the CEX and AMM, respectively. InversePrice is the inverse of the asset price, DailyVolume measures the daily trading volume, TVL represents the total value of assets locked in a liquidity pool, Volatility measures the price volatility of assets within a pool-pair, TS is the daily mean trade size, Stablecoin is a dummy variable that takes the value 1 if the token is a stablecoin pair and 0 if it is a token/token or stable/token pair. V3 is a dummy variable that takes the value of 1 following the launch of Uniswap V3 on May 5, 2021, and is 0 before this date. Measures of price, volume, pool and trade size are recorded in US\$. The model also includes fixed effects for both asset-pair and day, controlling for unobserved heterogeneity across different pools and days that could influence the variable of interest. To permit day fixed effects, we remove one daily dummy variable to include the dummy variable V3, indicating the time period after the introduction of Uniswap V3.

Table 5 reports parameter estimates and corresponding t-statistics, computed using heteroskedasticity-robust standard errors clustered at both the pool-pair and daily level (Petersen, 2008). The high adjusted R-squared values presented in Table 5, which indicate that our specification explains a substantial portion of the daily variation in transaction costs. Columns 1–3 present AMM

**Table 5**

Regression Analysis - Transaction Cost Comparison. This table reports the results of the regression analysis comparing market quality metrics between the AMM and CEX. Effective Spread, Price Impact, and Realized Spread are evaluated across both exchange types. The explanatory variables are the Inverse Price, Daily Volume (million USD), the daily median TVL (billion USD), the daily average Trade Size (TS), and TS-Squared, both in 10k USD. Volatility is the daily high-low range, scaled by its midpoint. Stablecoin is a binary variable indicating stablecoin pools. V3 is a binary variable taking the value of 1 for dates after the introduction of Uniswap V3. To permit day fixed effects and the variable V3, we remove one daily dummy variable. Standard errors in parentheses are clustered at the asset pair, day level. \*\*\*, \*\*, \* denote statistical significance at the 0.01, 0.05, and 0.10 levels, respectively.

	AMM			CEX		
	Effective Spread	Price Impact	Realized Spread	Effective Spread	Price Impact	Realized Spread
Intercept	2.980*** (0.111)	3.537*** (0.110)	3.378*** (0.212)	3.264*** (0.100)	2.473*** (0.090)	3.983*** (0.544)
Inverse Price	0.000*** (0.000)	0.000*** (0.000)	-0.000*** (0.000)	0.002*** (0.000)	0.003*** (0.001)	0.003* (0.001)
Daily Volume	0.002** (0.001)	0.006*** (0.001)	0.001** (0.000)	-0.000*** (0.000)	-0.000*** (0.000)	0.002*** (0.001)
TVL	-4.463*** (1.408)	-6.443*** (1.299)	-0.078 (0.339)			
Volatility	1.117*** (0.191)	0.782*** (0.172)	-0.612*** (0.149)	0.909*** (0.202)	2.482*** (0.302)	-0.270 (1.192)
TS	1.763*** (0.313)	1.576*** (0.288)	-0.627*** (0.231)	0.879 (1.553)	0.082 (1.687)	-0.135 (3.380)
TS-Squared	-0.343*** (0.093)	-0.300*** (0.083)	0.124* (0.064)	-0.549 (1.008)	0.968 (3.961)	-0.111 (2.095)
Stablecoin	-3.403*** (0.079)	-3.008*** (0.072)	1.001*** (0.042)	-3.573*** (0.281)	-2.307*** (0.122)	-2.843*** (0.622)
V3	1.849*** (0.205)	1.825*** (0.219)	-0.330 (0.231)	-0.158 (0.119)	0.184 (0.126)	-2.567*** (0.886)
N	26,549	26,537	17,641	21,860	19,319	17,040
Fixed Effects	Pair, Day	Pair, Day	Pair, Day	Pair, Day	Pair, Day	Pair, Day
Adj. R2	0.84	0.84	0.33	0.93	0.62	0.15

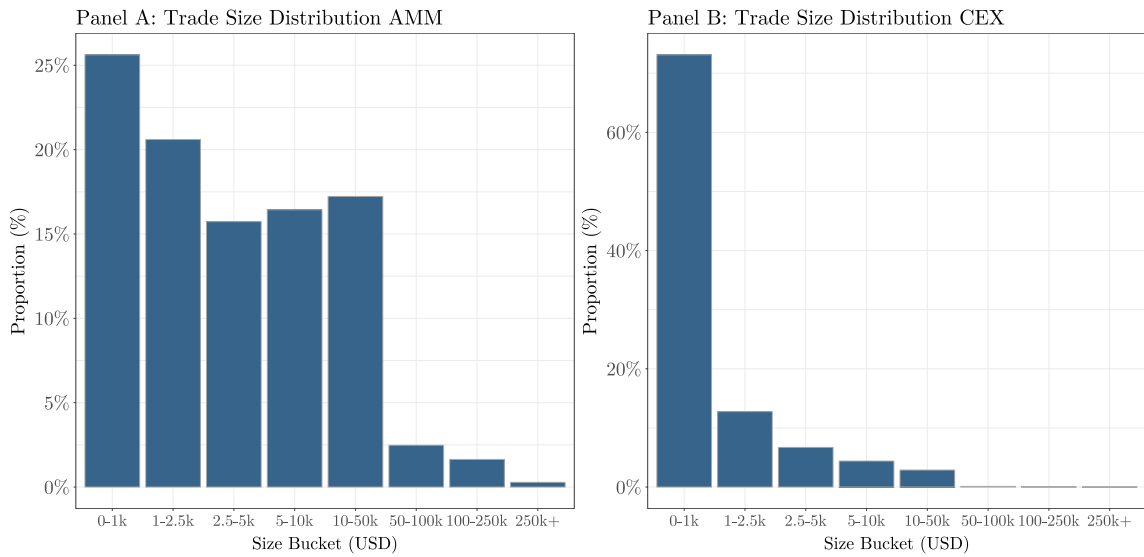
spread metrics, while columns 4-6 document CEX results. In the AMM specifications, we document economically significant liquidity effects, evidenced by the negative TVL coefficient. Our regression estimates describe the conditional mean of the transaction cost measures. Quantitatively, they imply that an increase in TVL of US\$100 million results in an approximately 10% reduction in average effective spreads and price impact, holding other factors constant. This inverse relationship between liquidity depth and transaction costs corroborates our preliminary findings in 4.3. The systematic nature of AMM pricing mechanisms is further evidenced by the robust statistical significance of trade size parameters. Evaluating the marginal effect at the mean of the trade size distribution, our estimates suggest that a US\$1000 increase in trade size corresponds to a 48% increase in average effective spreads and price impact costs. In contrast, CEX specifications exhibit economically insignificant trade size coefficients, indicating heterogeneous optimization of trading strategies across market structures.

Our analysis of volatility reveals a critical distinction between the two market structures. While effective spread coefficients are positive across both venues, the elasticity of price impact to volatility differs significantly. On CEXs, a 1% increase in volatility is associated with a 10% expansion in price impact, reflecting the tendency of active market makers to widen quotes to protect against inventory risk in turbulent periods. In contrast, AMMs exhibit a much lower elasticity, with the same volatility increase leading to only a 1.2% rise in price impact. This reflects a key trade-off inherent in the design of the AMM. The passive pricing function is structurally resilient in that its price curve does not widen in response to market volatility, offering execution certainty to traders. However, this stability for traders comes at a direct cost to liquidity providers, who are systematically exposed to adverse selection from arbitrageurs exploiting these temporarily stale prices, a cost quantified in our analysis by large, negative realized spreads.

Our analysis further reveals that stablecoin pairs exhibit significantly lower transaction costs across both market structures. Our multivariate analysis also captures the structural impact of the Uniswap V3 launch in May 2021. The coefficient on the post-V3 indicator is positive and highly significant for AMM spread measures, confirming that our benchmark venue (Uniswap V2) experienced a material increase in transaction costs after the upgrade. This finding is consistent with a liquidity migration effect, whereby the new, more capital-efficient protocol drew liquidity away from V2 pools, leading to wider spreads for the remaining activity. In contrast, the V3 launch had no discernible statistical effect on transaction costs within the CEX ecosystem, suggesting the impact of this DeFi innovation was localized to the on-chain environment.

## 5.2. The role of trade size

Our empirical analysis reveals that the average trade size on the AMM exceeds that of the CEX by a factor of three, a finding that warrants deeper investigation into the implications of trade size variation for transaction costs across exchange architectures. While centralized venues offer traders the ability to execute against the full depth of the order book, the inherent convexity of AMM pricing functions implies that larger trade sizes incur progressively higher price impact. To examine these dynamics at a more



**Fig. 2. Trade Size Distribution.** This figure illustrates the trade size distribution of transactions on the AMM (Panel A) and CEX (Panel B) during the month of March 2021.

granular level, we analyze the complete set of transactions executed during March 2021, encompassing approximately 800,000 AMM trades and 33.5 million CEX trades. We focus on this four-week window as it predates the deployment of Uniswap V3, thus representing a period of stable and comparable trading activity across both venues.<sup>17</sup>

Fig. 2 presents the distribution of trade sizes for the AMM (Panel A) and CEX (Panel B). The trade size buckets are chosen to facilitate a direct comparison across the two distinct distributions. The distributions exhibit pronounced right-skew characteristics, with trades concentrated in the smallest size category (0–\$1k). On the CEX, nearly 70% of trades fall into this category, compared to just over 25% on the AMM, suggesting fundamental differences in market microstructure and participant behavior. The CEX’s concentration in small trades aligns with established theoretical frameworks of limit order book dynamics, where strategic order splitting serves to minimize price impact (Biais et al., 1995; Viswanathan and Wang, 2002). In contrast, the AMM displays a more uniform distribution across intermediate trade sizes, with substantial activity in the 10–50k USD range (>15% of trades) compared to the CEX’s modest 3%. Moreover, the AMM accommodates a non-trivial fraction (approximately 5%) of large trades exceeding \$50k, a segment virtually absent in CEX trading. This distinctive pattern in the AMM’s trade size distribution reflects its underlying market architecture, where fixed blockchain-related transaction costs and deterministic price impact functions create economic incentives for larger trade sizes.

To analyze the cross-sectional variation in the trade size and transaction cost relationship across different market structures, we estimate the following empirical specification:

$$TxCost_{i,j,t} = \alpha_j + \delta_t + \beta_1 Market_i + \sum_{k=2}^6 \beta_k SizeDummy_k + \sum_{k=2}^6 \gamma_k (Market_i \times SizeDummy_k) + \epsilon_{i,j,t} \quad (3)$$

where  $TxCost_{i,j,t}$  is the transaction cost for trade  $i$  in pool pair  $j$  at time  $t$ ,  $\alpha_j$  and  $\delta_t$  account for both cross-sectional heterogeneity across pool pairs and temporal variation in market conditions.  $Market_i$  is the market type indicator and takes the value of 1 for AMM and 0 for CEX. We examine economically intuitive U.S. dollar size categories which allow us to capture different scales of activity, varying from retail to institutional. These are defined as 0–\$10k, 10–\$50k, 50–\$100k, 100–\$250k, 250k–\$500k, >\$500k. The coefficients of interest are  $\beta_1$  which captures the baseline difference in transaction costs between AMM and CEX markets;  $\beta_k$  ( $k = 2, \dots, 6$ ) which measures the size-specific effects relative to the smallest size category; and  $\gamma_k$  ( $k = 2, \dots, 6$ ) which identifies the differential size effects between AMM and CEX markets. To account for potential heteroskedasticity and within-pool pair correlation in the error terms, we employ robust standard errors clustered at the pool pair level.

The empirical analysis presented in Table 6 shows significant variation in transaction cost patterns across trade size categories and market architectures. The baseline category for the regression is a CEX trade under \$10,000. The coefficient on the AMM indicator variable is statistically insignificant, indicating that for the smallest trade size bucket, there is no significant difference in implicit transaction costs between AMMs and CEXs after controlling for other factors. For larger trades, however, the interaction

<sup>17</sup> Robustness checks conducted using an alternative monthly sample period yield quantitatively consistent results.

**Table 6**

Regression: Trade Sizes. This table reports the results of a fixed effects regression evaluating the differences of transaction costs across trading venues with a focus on trade sizes. The dependent variable is the effective spread in basis points as derived previously. High (low) TVL corresponds to pool, day observations in which the AMM pool TVL is above (below) the daily cross-sectional median. Depending on the model we control for the inverse price of the trade in USD and volatility that is the daily high-low range, scaled by its midpoint on the day the trade is executed. AMM is a binary variable that takes the value of 1 if the observation belongs to the AMM and is 0 for the CEX. Trade size dummy variables are in USD. Standard errors reported in parentheses are clustered at the asset pair level and \*\*\*, \*\*, \* denote the statistical significance at the 1, 5, and 10% level, respectively.

	Effective Spread	Effective Spread	Effective Spread Low TVL	Effective Spread High TVL
AMM	6.688 (5.466)	6.082 (5.175)	12.55 (8.623)	-1.748 (1.579)
10k-50k	0.798* (0.393)	0.782** (0.380)	1.547** (0.686)	0.148* (0.0674)
50k-100k	1.555** (0.616)	1.525** (0.599)	2.461** (1.051)	0.649*** (0.102)
100k-250k	0.961** (0.383)	0.974** (0.398)	1.601* (0.825)	0.463** (0.135)
250k-500k	0.611 (0.496)	0.617 (0.482)	1.076 (0.815)	0.0323 (0.0503)
above 500k	0.326 (0.342)	0.272 (0.235)	0.624* (0.342)	-0.185** (0.0634)
AMM × 10k-50k	24.04** (11.03)	24.13** (11.06)	44.45*** (15.71)	2.382*** (0.0728)
AMM × 50k-100k	17.17*** (5.892)	17.36*** (5.846)	27.16** (12.32)	9.580*** (0.378)
AMM × 100k-250k	18.01*** (4.451)	18.24*** (4.316)	21.98*** (7.754)	19.01*** (0.472)
AMM × 250k-500k	21.73*** (4.888)	21.78*** (4.778)	23.35** (10.23)	24.24*** (1.154)
AMM × above 500k	24.97*** (6.048)	24.89*** (5.978)	29.00*** (9.344)	24.60*** (1.142)
Controls		X	X	X
Fixed Effects	Pool, Day	Pool, Day	Pool, Day	Pool, Day
N	34,345,380	34,342,032	17,169,345	17,172,687
Adj. R2	48.71%	53.66%	49.37%	44.28%

terms (AMM × SizeDummy) reveal the additional cost premium for executing on an AMM relative to a CEX. These coefficients trace a distinct U-shaped pattern. The premium is highest for trades in the \$10,000-\$50,000 range, decreases for medium-sized trades (\$50,000-\$250,000), and then rises again for the largest trades (above \$250,000). This finding is consistent with the general cost structure pattern identified by Lehar and Parlour (2021) using daily average data.

We next disaggregate the sample by AMM liquidity, as measured by Total Value Locked (TVL). Columns 3 and 4 reveal diverging patterns. In low-TVL pools (Column 3), the U-shaped relationship persists and is even more pronounced. In these less liquid environments, the inherent convexity of the AMM pricing function is most punitive, leading to high price impact for both moderately sized and very large trades relative to the available liquidity. In stark contrast, for high-TVL pools (Column 4), the U-shaped pattern is not present. Instead, the AMM transaction cost premium increases monotonically with trade size. This observation complements (Barbon and Ranaldo, 2026), who find a monotonic increase in transaction costs in their sample of liquid pairs. This bifurcation in the size-cost relationship underscores the critical role of pool depth in determining the price impact dynamics which determines which cost structure prevails.

### 5.3. Counterfactual transaction cost analysis

Building upon our analysis of trade size distributions and their implications for execution costs, we examine the counterfactual simulation scenario of routing centralized exchange transactions through AMM venues. This analysis serves to isolate the implicit cost component of a transaction. By simulating the execution of CEX trades on their corresponding AMM pools, we can measure the difference in the cost arising purely from the price formation mechanism. This allows us to interpret the resulting “transaction cost gap” not as an unexploited arbitrage, but as an equilibrium measure of the value that traders implicitly place on avoiding other, non-modeled frictions inherent to the AMM architecture, such as network gas fees and smart contract risk.

The deterministic nature of AMM pricing mechanisms and the transparency of the blockchain enable the precise reconstruction of hypothetical execution costs. We accomplish this by reconstructing the temporal evolution of AMM liquidity states using token reserve data, which facilitates the simulation of CEX transactions through their corresponding AMM pools. Formally, the transaction cost gap for a single trade  $i$  executed at time  $t$  is defined as the difference between the actual effective spread on the CEX and the simulated, hypothetical effective spread on the AMM:

$$\text{TransactionCostGap}_{i,t} = \text{EffectiveSpread}_{i,t}^{\text{CEX, actual}} - \text{EffectiveSpread}_{i,t}^{\text{AMM, hypothetical}} \quad (4)$$

**Table 7**

Descriptive Statistics: Facilitating CEX Trades with the AMM. Panel A shows summary statistics of CEX trades that were used in the experiment. Price is the daily median price of the quote asset in USD for matched trades. Trade Size is the daily mean trade size of matched trades in USD. Daily Volume is the sum of daily trading volume of matched trades in USD millions. Volatility is the daily high-low range for all matched trades, scaled by their midpoints. AMM Reserve is the daily median total value locked in the AMM at the time of the matched trades in USD millions. Panel B reports the differences of effective and scaled effective spreads between the ‘actual’ (as it occurred on the CEX) and the ‘hypothetical’ trades (routed through the AMM). They are expressed in basis points. The column ‘Actual’ reports the median effective spread and scaled effective spreads of matched trades that occurred on the CEX in bps. The column ‘Hypothetical’ reports the median effective spread and scaled effective spreads of matched trades that were routed through the AMM in bps. The results are further split across asset pair category and time period. The asset pair category ‘tokens’ includes stable/token and token/token pairs and ‘stablecoins’ stable/stable pairs. The time period ‘pre’ (‘post’) is the time period before (after) the introduction of Uniswap V3. The numbers of asterisks indicate the significance at the 10%, 5% and 1% significance level.

Panel A: Matched Characteristics						
Variable	Mean	Median	SD	Q1	Q3	Q99
Price (USD)	1585.13	1.64	7257.98	0.00	0.01	47,077.02
Trade Size (USD)	784.94	448.31	929.42	54.06	88.61	4414.37
Daily Volume (USD)	51.06	0.91	194.54	0.02	0.03	941.57
Volatility (%)	7.90	6.16	7.49	0.00	0.02	36.55
AMM Reserve (USD)	30.70	1.53	81.86	0.00	0.02	389.20
Panel B: Transaction Cost Comparison						
Sample	Effective Spread			Scaled Effective Spread		
	Actual	Hypothetical	Difference	Actual	Hypothetical	Difference
<i>Full Sample</i>						
all	18.63	3.09	15.54***	0.13	0.03	0.10***
pre	21.73	1.06	20.67***	0.26	0.01	0.25***
post	17.45	5.55	11.90***	0.11	0.03	0.08***
<i>Tokens</i>						
all	20.18	3.16	17.02***	0.14	0.03	0.11***
pre	24.52	1.00	23.52***	0.33	0.02	0.31***
post	18.96	5.55	13.41***	0.12	0.04	0.08***
<i>Stablecoins</i>						
all	1.00	1.70	-0.70***	0.00	0.00	0.00***
pre	1.00	1.70	-0.70***	0.00	0.00	0.00***
post	1.00	1.15	-0.15	0.00	0.00	0.00**

where  $\text{EffectiveSpread}_{i,t}^{\text{CEX, actual}}$  is the observed effective spread for trade  $i$  executed on the CEX at time  $t$ .  $\text{EffectiveSpread}_{i,t}^{\text{AMM, hypothetical}}$  is the simulated effective spread for the same trade  $i$ , calculated as if it were executed against the state of the corresponding AMM liquidity pool at time  $t$ . The transaction cost gap for a specific asset pair on a given day is therefore, the median of all individual transaction cost gaps calculated for that asset pair throughout that day. The explicit algorithm for this counterfactual routing is detailed in [Appendix B1](#).

To ensure the robustness of our analysis and mitigate potential bias from extreme scenarios, we implement filtering criteria that exclude transactions consuming more than 10% of pool liquidity or executing against pools with reserves below US\$50,000.<sup>18</sup> This methodology successfully captures approximately 95% of CEX transactions and 90% of trading volume, with the resulting subsample maintaining statistical equivalence to the original dataset in terms of key trading characteristics.

[Table 7](#) describes the result of the matching analysis. Panel A documents the characteristics of matched CEX transactions, demonstrating statistical equivalence with the broader CEX sample. Panel B reports the comparative transaction cost metrics, disaggregated across token (stable/token and token/token) and stablecoin (stable/stable) pairs. We examine both raw effective spreads (in basis points) and size-normalized effective spreads to account for trade size heterogeneity.

The counterfactual analysis reveals substantially lower execution costs for trades hypothetically routed through AMM venues, with median spreads of 3.09 basis points compared to 18.63 basis points observed on CEX platforms. This cost advantage persists throughout the sample period, though the magnitude diminishes following the introduction of Uniswap V3. The cross-sectional analysis across pair categories reaffirms our earlier findings that stablecoin pair transactions incur higher costs in AMM venues, attributable to the execution of larger trades against relatively constrained liquidity pools in Uniswap V2. The robustness of these findings is further substantiated by size-normalized transaction cost metrics, which demonstrate that the observed cost differentials are not merely artifacts of trade size variation.

<sup>18</sup> This 10% threshold is a conservative filter designed to exclude transactions that are not representative of typical, rational liquidity-demanding behavior. Trades consuming such a large fraction of a pool's reserves induce extreme, non-linear price impact that is not economically comparable to executions on a deep limit order book. These outlier transactions are more likely to represent either user error, part of a complex MEV strategy (such as oracle price manipulation), or an execution in a highly illiquid pool where the resulting price impact is not a meaningful measure of transaction cost. Including them would introduce a significant upward bias to our simulated cost measures, disproportionately affecting the mean and potentially skewing the median

**Table 8**

Descriptive statistics: Portfolios (Hypothetical Transaction Costs). This table reports summary statistics of portfolios based on multiple asset pair characteristics. For each day, we classify asset pairs according to the respective characteristic into five portfolios from low (1) to high (5). We further split the hypothetical sample across 'low', 'medium' and 'high' AMM TVL terciles. After, we compute median transaction cost metrics over the timeseries that we show in the table. The second last row of each panel shows differences in means between the fifth (highest) and first (lowest) portfolio. The last row reports the *t*-statistic of a *t*-test testing for differences in means between these two groups. \*, \*\*, and \*\*\* denote statistical significance at the 10%, 5%, and 1% levels, respectively.

	Effective Spread (Actual)	Effective Spread (Hypothetical)			
Panel A: Sorted on Volatility					
Portfolio	Actual	Low	Medium	High	High-Low
1	16.31	279.64	9.45	1.09	-278.56***
2	24.19	115.96	7.04	0.88	-115.07***
3	26.12	101.93	5.76	1.02	-100.91***
4	26.91	100.57	5.74	1.18	-99.39***
5	30.75	100.49	7.40	1.55	-98.93***
(5-1)	14.44*** (28.08)	-179.16*** (-16.2)	-2.05*** (-3.98)	0.47*** (5.12)	
Panel B: Sorted on Volume					
1	43.85	230.76	4.68	1.76	-229.01***
2	37.55	49.79	8.28	2.03	-47.76***
3	24.77	18.89	7.10	2.25	-16.63***
4	10.78	10.43	5.88	1.56	-8.88***
5	1.80	59.41	6.27	0.57	-58.84***
(5-1)	-42.06*** (-123.13)	-171.35*** (-3.78)	1.59 (1.18)	-1.19 (-1.69)	
Panel C: Sorted on Market Capitalization					
1	42.58	149.56	3.09	1.05	-148.51***
2	39.68	124.67	5.59	2.17	-122.50***
3	24.32	125.48	5.13	1.30	-124.18***
4	9.57	179.73	7.04	1.33	-178.39***
5	2.59	397.95	22.89	0.70	-397.25***
(5-1)	-39.99*** (-102.53)	248.39*** (11.1)	19.80*** (25.45)	-0.35*** (-7.53)	

Our comparative analysis of market architectures under equivalent trade size conditions reveals AMM venues to be substantially more competitive than previously indicated. This analysis deliberately excludes exogenous transaction costs to isolate and contrast the fundamental operational dynamics of limit order books and automated market makers. The observed cost differentials provide compelling evidence of AMM efficiency, though we acknowledge that comprehensive venue selection decisions must consider additional factors such as exchange fees and blockchain network costs.<sup>19</sup>

This equilibrium may stem from the fundamental difference in market architecture, which gives rise to opposing, endogenous execution strategies. Centralized exchanges, operating on a limit order book (CLOB) model, incentivize traders to engage in order splitting—breaking large orders into smaller pieces to minimize price impact and hide trading intentions. Conversely, the AMM cost structure, characterized by a high, fixed network gas fee for every on-chain transaction, encourages order consolidation, where traders bundle demand into larger trades to amortize this fixed cost. Our counterfactual simulation routes the high volume of small CEX trades to an AMM. These are precisely the transactions for which the AMM high fixed costs are most punitive on a relative basis. These fixed costs make AMMs economically non-viable for small orders, thereby incentivizing the order consolidation we observe in AMM trade data. Therefore, the fact that these trades remain on the CEX is a rational economic decision. The trade gap quantifies the implicit spread savings that would have been required to overcome the explicit cost of the gas fee and any perceived differences in platform risk (e.g., custodial vs. smart contract risk). The observed outcome implies that for the marginal CEX trader, these unobserved costs may exceed the potential savings in the price impact dimension.

This interpretation is further reinforced by the portfolio analysis in Table 8. The analysis shows that while CEX spreads are highly sensitive to asset-specific characteristics like volatility and market capitalization, the execution costs for hypothetical trades in well-capitalized (medium and high-TVL) AMM pools are remarkably stable. This structural resilience highlights that AMMs are not merely a substitute for CEXs, but a different type of trading architecture that offers a distinct trade-off: higher fixed costs and deterministic, liquidity-dependent price impact in exchange for execution certainty and resilience to certain market dynamics. The choice of venue is thus a choice between these different bundles of attributes, tailored to the trader's specific execution strategy and cost sensitivities.

<sup>19</sup> In 2024, the tiered spot trading fees for the largest CEX Binance are 1 bps to 7.5 bps, which is lower than the 30 bps fee of Uniswap V2 but comparable to newer AMM implementations such as Uniswap V3. For trades on AMMs the traders also need to consider the flat network costs ('gas cost') that is paid to block producers and validators. For a further discussion on the implications of gas fees, we refer to Barbon and Rinaldo (2026).

**Table 9**

Regression Analysis: Transaction Cost Gap. This table reports the results of regressions evaluating the effective spread across several gap measures, including token and stablecoin gaps. Volatility is included as an explanatory variable. The AMM Share variable is the AMM market share of daily trading volume in the asset pair, lagged by one period. Price Variance is the relative price deviation of the AMM, lagged by one period in percentage. V3 is a binary variable taking the value of 1 for dates after the introduction of Uniswap V3. To permit day fixed effects and the variable V3, we remove on daily dummy variable. Stablecoin is a binary variable indicating stablecoin pools. The row Controls indicates whether additional control variables (inverse price, TVL, trading volume, average trade size) are included. Robust standard errors are clustered at the pool pair (token) level. \*\*\*, \*\*, and \* denote significance at the 1%, 5%, and 10% levels, respectively.

	Gap	Gap	Gap	Gap (Token)	Gap (Stablecoin)
Intercept	-1.795*** (0.382)	-3.677*** (0.232)	-3.801*** (0.265)	-3.836*** (0.286)	6.858* (1.659)
Volatility		-1.637*** (0.533)	-1.883*** (0.587)	-1.737*** (0.564)	83.958** (19.300)
AMM Share		-2.032*** (0.489)	-1.953*** (0.478)	-1.762*** (0.449)	4.039** (0.673)
Price Variance		31.015*** (6.432)	33.078*** (6.919)	32.898*** (6.747)	-20.038* (5.834)
V3		1.764*** (0.502)	1.791*** (0.430)	1.920*** (0.447)	0.279 (0.128)
Stablecoin		2.188*** (0.568)	2.387*** (0.115)		
Controls	No	Yes	Yes	Yes	Yes
Fixed Effects	Pool, Day	Pool, Day	Pool, Day	Pool, Day	Pool, Day
Adj. R2	69.77%	72.29%	71.89%	70.17%	94.76%
N	18,121	18,121	18,121	17,175	946

### 5.3.1. The transaction cost gap

To examine the drivers of this transaction cost gap between market structures, we model this difference between the AMM hypothetical effective spread and the actual CEX effective spread under a multivariate framework. For analytical purposes, we aggregate this metric to the asset pair-day level using median values.

Our models take the following specification:

$$CostGap_{i,t} = \alpha + \beta_1 Volatility_{i,t} + \beta_2 AmmShare_{i,t-1} + \beta_3 PriceVariance_{i,t-1} + \beta_4 V3_t + \beta_5 Stablecoin_i + \varepsilon_{i,t} \quad (5)$$

where the dependent variable is a log-modular transformation of the gap measure, which allows us to appropriately handle directional differences in transaction costs. Explanatory variables include volatility, the lagged volume market share of the AMM, and a lagged price variance measure that is the relative price deviation of the AMM compared to the CEX. The binary variable V3 takes the value of 1 for observations after the introduction of Uniswap V3 (May 5, 2021) and the binary variable Stablecoin takes the value 1 if the asset pair consists of two stablecoins and is 0 otherwise. We further center the explanatory variables to allow for a more meaningful interpretation of the coefficients and include further controls for TVL, price, trade size and daily volume.

Table 9 reveals systematic patterns in the execution cost differential across the different market structures. The negative intercept coefficients across specifications indicate that AMM venues exhibit lower transaction costs relative to CEX platforms under mean covariate conditions (due to centering). The notable exception emerges in the stablecoin subsample, where the positive intercept suggests a CEX cost advantage, consistent with the structural limitations of Uniswap V2's constant product mechanism in facilitating stablecoin transactions. The analysis reveals that heightened volatility corresponds to a narrowing transaction cost gap, with AMM execution costs becoming increasingly competitive relative to CEX venues. This pattern stems from fundamental differences in price formation mechanisms: while CEX liquidity providers actively adjust quotes in response to market conditions, AMM prices evolve through a mechanical function of pool reserves, where liquidity providers bring no information to the market and have no impact on price (Hansson, 2024). Without trading, price adjustments are unattainable. This structural distinction results in AMMs providing enhanced execution stability and cost efficiency during periods of elevated market volatility.

The negative coefficients of the lagged AMM share variable and the transaction cost differential demonstrate a positive feedback liquidity dynamic, where increased AMM trading activity leads to enhanced execution efficiency, further narrowing or reversing the cost gap between market structures. For AMMs, competition between liquidity providers drives them to provide excess liquidity to capture a larger share of trading fees. This happens because fee revenue is allocated proportionally to the amount of liquidity deposited. The competitive behavior leads to a constant amount of excess liquidity that increases linearly with base demand and helps to reduce transaction costs for participants.<sup>20</sup>

The coefficient on lagged PriceVariance indicates that historical price divergences between venues typically result in widening transaction cost differentials favoring CEX execution, highlighting the critical role of arbitrage in AMM price formation. If

<sup>20</sup> Ma and Crapis (2024) show that higher liquidity provision can lead to welfare loss as more capital becomes exposed to adverse selection costs, and that this loss scales linearly with the number of liquidity providers.

transaction costs on the AMM become too expensive, arbitrageurs are disincentivised to bring the AMM into a new equilibrium state.<sup>21</sup> Notably, stablecoin pairs exhibit contrary dynamics, where increased cross-venue price dispersion corresponds to enhanced AMM cost competitiveness. The temporal and cross-sectional dummy variables corroborate our earlier findings regarding market structure evolution. The positive V3 coefficient reflects diminished AMM cost competitiveness following the introduction of Uniswap V3, consistent with liquidity migration effects. Similarly, the significant positive Stablecoin coefficient reaffirms the structural inefficiencies of the single-fee tier constant product model in facilitating stablecoin transactions. Overall, we show that the AMM becomes a more competitive execution venue in market environments that exhibit higher volatility and a competitively close price alignment with centralized venues.

## 6. Conclusion

This paper examines the value of decentralized financial architectures. Decentralization in financial markets operates at multiple layers, encompassing both the asset level, through tokenization, and the infrastructure level, via novel trading protocols such as Automated Market Makers (AMMs). This paper critically examines these twin dimensions of decentralization, focusing on the introduction of tokenized equities and the adoption of AMM mechanisms in securities market design.

The tokenization of real-world assets (RWA) is an area of considerable promise for the investments industry. Our examination of individual stock equities, however, shows no significant benefits in terms of liquidity provision or price efficiency relative to the same equities traded on incumbent national market systems. Whilst tokenized equities can facilitate round the clock trading and near-instantaneous settlement, a high cost of transaction environment and the slow information incorporation into prices is correlated with this lack of adoption. This case demonstrates that a decentralized “label” does not guarantee a superior market; in this instance, the resulting environment is demonstrably less efficient and riskier than its traditional counterpart.

Our second contribution is the identification and economic interpretation of a transaction cost gap between CEXs and AMMs. We show that while a direct comparison of spreads suggests AMMs are more expensive, a simulated counterfactual analysis reveals significant potential cost savings for CEX-sized trades routed through AMMs. We argue this persistent gap is an equilibrium outcome, reflecting the price traders rationally pay to avoid the AMM architecture’s unique frictions. Namely, high, fixed network fees and unpriced smart contract risk. This provides a framework for understanding endogenous order flow segmentation. Furthermore, we uncover a key structural trade-off in the AMM model related to market volatility. While execution costs for traders in well-capitalized AMMs remain remarkably stable during turbulent periods, in contrast to CEXs where spreads predictably widen, this stability comes at a direct cost to liquidity providers. The passive pricing function associated with the AMM exposes liquidity providers to significant adverse selection from arbitrageurs exploiting temporarily stale prices, a cost quantified in our analysis by the large, negative realized spreads we document. This finding highlights that while AMMs can serve as a valuable source of liquidity when it is most scarce, the cost of providing this stability is borne by the liquidity providers.

Our analysis opens several specific avenues for future research. First, our empirical findings on early, unregulated tokenized equity markets illuminate the microstructure challenges that impede adoption of decentralized non-native assets. These results raise important questions about the conditions under which tokenization can generate efficiency gains, and whether meaningful decentralization is achievable when centralized assets are wrapped in decentralized infrastructure but traded on centralized marketplaces. This requires further inspection of network externalities and the critical idea of whether wrapping a centralized asset in a decentralized token wrapper to be traded on a centralized marketplace is ultimately a meaningful form of decentralization. Our evidence, therefore, may not extrapolate to the emerging generation of exchange-sponsored tokenization initiatives, such as those recently proposed by Nasdaq and the NYSE. A natural direction for future work is to compare trading costs, liquidity, and price discovery between these institutionally backed tokenized equities and the earlier retail-oriented implementations we study here. Second, our interpretation of the transaction cost gap invites a more formal modeling approach and requires careful caveats regarding interpretation. The gap should be understood as measuring the implicit spread savings required to overcome measured execution costs and unmodeled DEX-specific frictions such as gas costs, MEV risk, and smart contract vulnerabilities, conditional on the observed distribution of liquidity. Future work that explicitly controls for this heterogeneity could refine these gap estimates and provide a more comprehensive welfare analysis of venue choice, an avenue we leave to future research. Finally, our finding that well-capitalized AMMs can efficiently accommodate large trades with low sensitivity to volatility also suggests their potential applicability beyond crypto-native assets. A compelling path for future research is to assess the viability of AMM protocols in traditionally illiquid, dealer-dominated markets, such as those for corporate bonds, where sourcing liquidity remains a primary challenge.

## Appendix A

### A.1. List of instruments

Our analysis relies on a comprehensive sample of tokenized equities and digital assets. [Table A.1](#) shows the equity and digital asset pair as well as their categorization.

### A.2. Tokenized price deviations

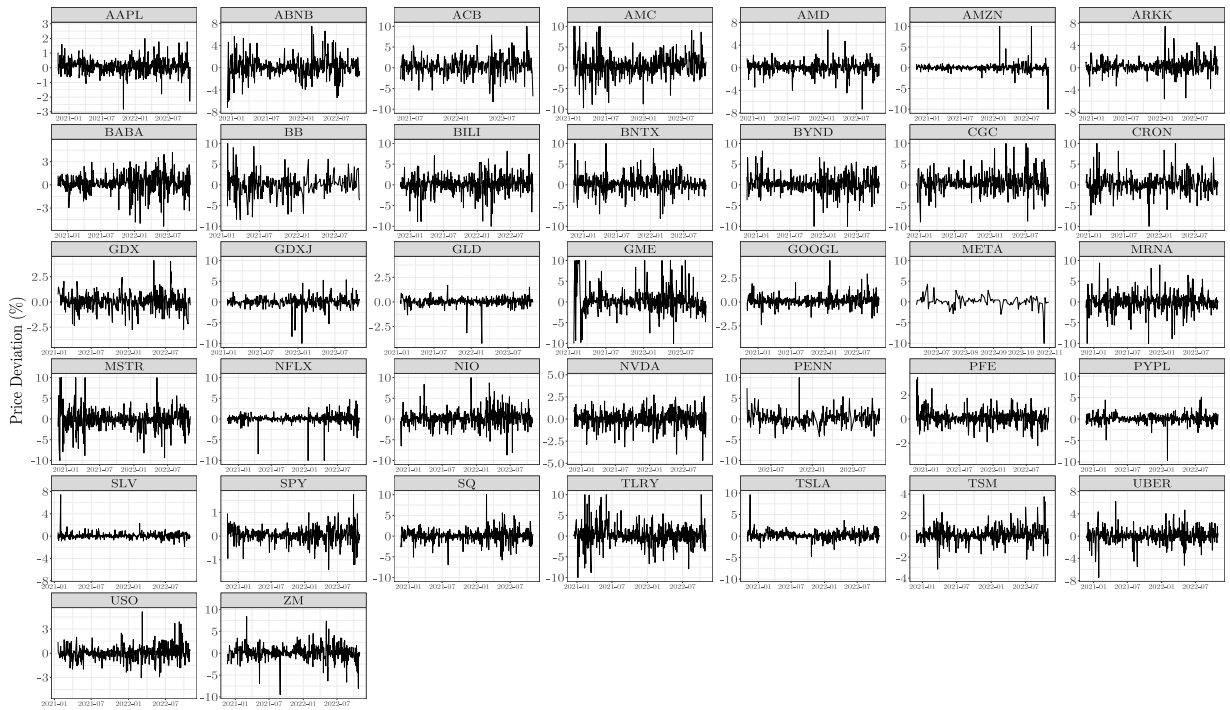
See [Fig. A.1](#).

<sup>21</sup> For a discussion on price discovery between CEXs and AMMs, we refer to [Félez-Viñas et al. \(2023\)](#).

**Table A.1**

List of instruments. This table lists the instruments used in this study. The column 'Equities' displays the ticker symbols of traditional and tokenized equities. The column 'Crypto Asset Pair' shows the crypto asset pairs used. The column 'Crypto Asset Pair Category' classifies each crypto asset pair into stable/stable, stable/token, or token/token categories.

#	Equities	Crypto Asset Pair	Crypto Asset Pair Category
1	AAPL	USDC/ETH	stable/token
2	ABNB	ETH/USDT	stable/token
3	ACB	DAI/ETH	stable/token
4	AMC	UNI/ETH	token/token
5	AMD	BTC/ETH	token/token
6	AMZN	AAVE/ETH	token/token
7	ARKK	USDC/USDT	stable/stable
8	BABA	MATIC/ETH	token/token
9	BB	LINK/ETH	token/token
10	BILI	UST/USDT	stable/token
11	BNTX	ETH/CRV	token/token
12	BYND	SNX/ETH	token/token
13	CGC	SAND/ETH	token/token
14	CRON	ETH/GRT	token/token
15	META	BTC/USDC	stable/token
16	GDX	ETH/ENJ	token/token
17	GDXJ	MANA/ETH	token/token
18	GLD	ETH/LUNA	token/token
19	GLXY	HOT/ETH	token/token
20	GME	NPXS/ETH	token/token
21	GOOGL	RUNE/ETH	token/token
22	MRNA	USDT/FRONT	stable/token
23	MSTR	DAI/USDT	stable/stable
24	NFLX	BAT/ETH	token/token
25	NOK	LUNA/USDT	stable/token
26	NIO	LRC/ETH	token/token
27	NVDA	UFT/ETH	token/token
28	PENN	CHZ/USDT	stable/token
29	PFE	FTM/ETH	token/token
30	PYPL	CVP/ETH	token/token
31	SLV	BNT/ETH	token/token
32	SPY	LINA/USDT	stable/token
33	SQ	ETH/AXS	token/token
34	TLRY	BOND/ETH	token/token
35	TSLA	ETH/MFT	token/token
36	TSM	ETH/KNC	token/token
37	UBER	DENT/ETH	token/token
38	USO	ETH/ZRX	token/token
39	ZM	TUSD/ETH	token/token
40		ETH/OMG	token/token
41		HEGIC/ETH	token/token
42		GLM/ETH	token/token
43		STMX/ETH	token/token
44		VGX/ETH	token/token
45		GALA/ETH	token/token
46		BUSD/USDC	stable/stable
47		RLC/ETH	token/token
48		FUN/ETH	token/token
49		FARM/ETH	token/token
50		FORTH/USDT	stable/token



**Fig. A.1. Price deviations of tokenized and traditional equities.** This figure illustrates the price deviation across equities markets. Price deviation is computed at the daily level in percentage by taking the differential of the median tokenized price and median traditional price relative to the median price of the traditional equity.

## Appendix B

### B1. Counterfactual simulation methodology

The counterfactual simulation exercise is designed to estimate the execution cost that a CEX trade would have incurred had it been routed through its corresponding AMM pool at the exact moment of its execution. The analysis excludes protocol fees to isolate the cost of the pricing mechanism. The algorithm proceeds as follows:

For each individual CEX trade  $i$  executed at time  $t$  with volume  $V$  and direction  $d$  (buy or sell of Asset A for Asset B):

- 1. Identify AMM State:** We identify the state of the corresponding AMM liquidity pool for the A/B pair immediately prior to time  $t$ . This state is defined by the token reserves  $(x_t, y_t)$ , where  $x_t$  is the reserve of Asset A and  $y_t$  is the reserve of Asset B. The pre-trade midpoint on the AMM is calculated as  $m_t^{\text{AMM}} = y_t/x_t$ .
- 2. Simulate Fee-less Trade Execution:** We simulate the execution of the CEX trade against the AMM pool's reserves using the constant product formula  $x \cdot y = k$  without applying any protocol fee. If the trade is a **buy** of Asset B (inputting Asset A,  $\Delta x = V_A$ ), the amount of Asset B received is calculated as:

$$\Delta y = \frac{y_t \cdot \Delta x}{x_t + \Delta x}$$

Conversely, if the trade is a **sell** of Asset B (inputting Asset B,  $\Delta y = V_B$ ), the amount of Asset A received is calculated as:

$$\Delta x = \frac{x_t \cdot \Delta y}{y_t + \Delta y}$$

- 3. Calculate Hypothetical Cost:** The hypothetical execution price for the simulated trade is calculated as  $p_t^{\text{sim}} = \Delta y/\Delta x$ . We then compute the hypothetical effective spread using the same methodology as for the observed trades:

$$\text{EffectiveSpread}_i^{\text{sim}} = d \cdot \frac{p_t^{\text{sim}} - m_t^{\text{AMM}}}{m_t^{\text{AMM}}}$$

where  $d$  is +1 for a buy and -1 for a sell.

- 4. Iteration:** This process is repeated for every CEX trade in our sample. Each simulation is an independent counterfactual, meaning each CEX trade is routed against the *actual* historical state of the AMM at that specific time, not a simulated state that accounts for prior hypothetical trades. This ensures we are isolating the cost of the pricing mechanism for each trade individually.

## Data availability

Our dataset has been made publicly available - please see further details inside the paper for access.

## References

- Allen, F., Gale, D., 1988. Optimal security design. *Rev. Financ. Stud.* 1 (3), 229–263.
- Aoyagi, J., Ito, Y., 2025. Coexisting exchange platforms: Limit order books and automated market makers. *J. Political Econ. Microeconomics* 3 (3), 611–648.
- Aquilina, M., Foley, S., Gambacorta, L., Krekel, W., 2024a. Decentralised dealers? Examining liquidity provision in decentralised exchanges. *BIS Working Papers* (1227).
- Aquilina, M., Frost, J., Schrimpf, A., 2024b. Decentralized finance (DeFi): a functional approach. *J. Financ. Regul.* 10 (1), 1–27.
- Aspris, A., Foley, S., Svec, J., Wang, L., 2021. Decentralized exchanges: The “wild west” of cryptocurrency trading. *Int. Rev. Financ. Anal.* 77, 101845. <http://dx.doi.org/10.1016/j.irfa.2021.101845>.
- Aspris, A., Svec, J., 2025. Locked in, levered up: Risk, return, and ruin in DeFi lending. *Br. Account. Rev.* 101691. <http://dx.doi.org/10.1016/j.bar.2025.101691>.
- Barbon, A., Ranaldo, A., 2026. On the quality of cryptocurrency markets: Centralized vs. decentralized exchanges. *Manage. Sci.*
- Biais, B., Hillion, P., Spatt, C., 1995. An empirical analysis of the limit order book and the order flow in the Paris Bourse. *J. Financ.* 50 (5), 1655–1689.
- Biais, B., Martimort, D., Rochet, J.-C., 2000. Competing mechanisms in a common value environment. *Econometrica* 68 (4), 799–837.
- Calvet, L.E., Celerier, C., Sodini, P., Vallee, B., 2023. Can security design foster household risk-taking? *J. Financ.* 78 (4), 1917–1966.
- Capponi, A., Jia, R., 2021. The adoption of blockchain-based decentralized exchanges. *arXiv preprint arXiv:2103.08842*.
- Carapella, F., Chuan, G., Gerszten, J., Hunter, C., Swem, N., 2023. Tokenization: Overview and financial stability implications. In: *Finance and Economics Discussion Series*, vol. 2023-060r1.
- Céli er, C., Vall e, B., 2017. Catering to investors through security design: Headline rate and complexity. *Q. J. Econ.* 132 (3), 1469–1508.
- Copeland, T.E., Galai, D., 1983. Information effects on the bid-ask spread. *J. Financ.* 38 (5), 1457–1469.
- Coval, J., Jurek, J., Stafford, E., 2009. The economics of structured finance. *J. Econ. Perspect.* 23 (1), 3–25.
- Esparcia, C., Escribano, A., Jare o, F., 2023. Did cryptomarket chaos unleash silvergate’s bankruptcy? investigating the high-frequency volatility and connectedness behind the collapse. *J. Int. Financ. Mark. Institutions Money* 89, 101851.
- F el ez-Vi as, E., Johnson, L., Putnin s, T., 2023. How well do automated market makers provide price discovery? Working Paper, University of Technology, Sydney.
- Foley, S., Krekel, W., Kwan, A., 2025. Can simple markets survive? Exploring the tradeoffs of AMM complexity. Available at SSRN 5271981.
- Foley, S., Krekel, W., Mollica, V., Svec, J., 2023. Not so fast: Identifying and remediating slow and imprecise cryptocurrency exchange data. *Financ. Res. Lett.* 51, 103401.
- Glosten, L.R., 1994. Is the electronic open limit order book inevitable? *J. Financ.* 49 (4), 1127–1161.
- Hansson, M., 2024. Price discovery in constant product markets. Available at SSRN 4582649.
- Hasbrouck, J., Rivera, T.J., Saleh, F., 2026. The need for fees at a dex: How increases in fees can increase dex trading volume. *Manage. Sci.*
- Hasso, T., M uller, D., Pelster, M., Warkulat, S., 2022. Who participated in the GameStop frenzy? Evidence from brokerage accounts. *Financ. Res. Lett.* 45, 102140. <http://dx.doi.org/10.1016/j.frl.2021.102140>.
- Hendershott, T., Jones, C.M., Menkveld, A.J., 2011. Does algorithmic trading improve liquidity? *J. Financ.* 66 (1), 1–33.
- Hou, K., Moskowitz, T.J., 2005. Market frictions, price delay, and the cross-section of expected returns. *Rev. Financ. Stud.* 18 (3), 981–1020.
- Intercontinental Exchange, 2026. The New York Stock Exchange develops tokenized securities platform. <https://www.businesswire.com/news/home/20260119300589/en/The-New-York-Stock-Exchange-Develops-Tokenized-Securities-Platform>. (Accessed: 20 January 2026).
- Kreppmeier, J., Laschinger, R., Steininger, B.I., Dorfleitner, G., 2023. Real estate security token offerings and the secondary market: Driven by crypto hype or fundamentals? *J. Bank. Financ.* 154, 106940.
- Lee, M.J., Martin, A., Townsend, R.M., 2024. Zero settlement risk token systems. Federal Reserve Bank of New York, Staff Reports (1120).
- Lee, C.M., Ready, M.J., 1991. Inferring trade direction from intraday data. *J. Financ.* 46 (2), 733–746.
- Lehar, A., Parlour, C., 2025. Decentralized exchange: The uniswap automated market maker. *J. Financ.* 80 (1), 321–374.
- Ma, J., Crapis, D., 2024. The cost of permissionless liquidity provision in automated market makers. *arXiv preprint arXiv:2402.18256*.
- Malinova, K., Park, A., 2023a. Tokenized stocks for trading and capital raising. Available at SSRN 4365241.
- Malinova, K., Park, A., 2023b. Tokenomics: When tokens beat equity. *Manag. Sci.* 69 (11), 6568–6583.
- Malinova, K., Park, A., 2024. Learning from DEFI: Would automated market makers improve equity trading? Available at SSRN 4531670.
- Milionis, J., Moallemi, C.C., Roughgarden, T., Zhang, A.L., 2022. Quantifying loss in automated market makers. In: *Proceedings of the 2022 ACM CCS Workshop on Decentralized Finance and Security*. pp. 71–74.
- Momtaz, P.P., 2023. Security tokens. In: *The Emerald Handbook on Cryptoassets: Investment Opportunities and Challenges*. Emerald Publishing Limited, pp. 61–78.
- Nasdaq, 2025. Q&A: Nasdaq’s new proposal for tokenized securities. <https://www.nasdaq.com/newsroom/qa-nasdaq-new-proposal-tokenized-securities>. (Accessed: 01 November 2025).
- Parlour, C.A., Seppi, D.J., 2003. Liquidity-based competition for order flow. *Rev. Financ. Stud.* 16 (2), 301–343.
- Parlour, C.A., Seppi, D.J., 2008. Limit order markets: A survey. *Handb. Financ. Intermediation Bank.* 5, 63–95.
- Petersen, M.A., 2008. Estimating standard errors in finance panel data sets: Comparing approaches. *Rev. Financ. Stud.* 22 (1), 435–480.
- Viswanathan, S., Wang, J.J., 2002. Market architecture: limit-order books versus dealership markets. *J. Financ. Mark.* 5 (2), 127–167.
- Yousaf, I., Pham, L., Goodell, J.W., 2023. The connectedness between meme tokens, meme stocks, and other asset classes: Evidence from a quantile connectedness approach. *J. Int. Financ. Mark. Institutions Money* 82, 101694.