On The Quality Of Cryptocurrency Markets

Centralized Versus Decentralized Exchanges

Andrea Barbon, Angelo Ranaldo *

First Draft: July 2021. This Draft: October 2022

Abstract

We compare the market quality of decentralized blockchain-based venues (DEXs) to centralized crypto exchanges (CEXs). Our analysis of transaction costs and deviations from the no-arbitrage condition suggests that liquidity in CEXs and DEXs is similar but DEX prices are less efficient. While the main frictions for DEXs are high exchange fees and the gas cost of blockchain transactions, CEXs involve significant risks and latency associated with delegated custody. We propose and empirically validate a stylized model of DEX liquidity provision, linking volume, fees, and liquidity in equilibrium. Our theory identifies the quantitative conditions for DEXs to overtake CEXs.

Keywords: Market Quality, Decentralized Exchanges, Automated Market Making, Blockchain, Decentralized Finance, Limit Order Book

---

*Andrea Barbon (andrea.barbon@unisg.ch) is Assistant Professor of Finance at the University of St. Gallen, Switzerland. Angelo Ranaldo (angelo.ranaldo@unisg.ch) is Full Professor of Finance and Systemic Risk at the University of St .Gallen, Switzerland. We thank Agostino Capponi, Alfred Lehar, Thomas Moser, Engin Iyidogan, Laurence Daures, Igor Kozhanov, Peter O'Neill, and participants of the SNB-CIF Conference on Cryptoassets and Financial Innovation, the Central Bank Research Association (CEBRA) Annual Meeting, SNB Technology and Finance Webinars, the Bank of Canada finance seminars, the World Federation of Exchanges Webinars, the FiRe Research Seminars, the 13th Annual Hedge Fund Research Conference, for advice and stimulating discussions. The authors thank Charles Milliet for his excellent research assistance. All errors are our own.

I. Introduction

In modern financial markets, many asset classes are traded on centralized exchanges (CEXs). This dominant market structure often relies on an electronic limit order book (LOB) that matches end-user orders in a reasonably transparent, efficient, and centralized way. LOB markets have also been widely adopted for trading cryptocurrencies *off-chain* on centralized exchanges. Recently, however, fueled by the wave of innovation brought about by the advent of blockchain technology, decentralized exchanges (DEXs) have emerged as an alternative market structure for crypto assets. These venues are based on smart-contract implementations of automated market makers (AMM) that enable *on-chain* trading. Although recently introduced, they have been attracting increasing attention and trading volumes.[1] One of the questions arising from this development concerns the market quality offered by these new market systems compared to LOB-based CEXs. We address this issue by empirically assessing the market quality of CEXs and DEXs, focusing on market liquidity and price efficiency.

Our first contribution is to lay out the institutional details of trading and settlement, and the main cost components that distinguish CEXs and DEXs: exchange fees, bid-ask spreads, and gas fees. Second, we conduct a thorough empirical analysis of the transaction costs and deviations from the *law of one price*. We find that DEXs operate with similar transaction costs to CEXs, suggesting that their market liquidity is comparable. However, an analysis of arbitrage bounds estimated for the main exchange triplets shows that DEX cryptocurrency prices tend to be less efficient than CEX prices. We show that this is due to the crucial role of gas fees in restoring the "triangular" no-arbitrage condition, a reflection of the cost of recording transactions

[1] Throughout the paper, we use the acronym CEX to indicate a LOB-based centralized exchange and DEX to indicate an AMM-based decentralized one. It should be noted that CEXs depend on a proprietary IT platform (server, databases, account system, security, etc.), while DEXs rely only on smart contracts and blockchain technology. Also, LOB differs from AMM due to the institutional arrangements regarding liquidity provision. As explained in detail in Section II, in the former, quotes are submitted by market makers and recorded in the order book; in the latter, liquidity is posted to pools by market participants, and a mathematical framework determines transaction prices.

on the blockchain. Gas fees, which depend on the dollar price of the native token and the level of network congestion, affect every DEX transaction, thus burdening arbitrage involving multiple trades. On the other hand, CEXs are subject to various forms of risk, including unauthorized access by hackers and misappropriation of client funds. In contrast, the same cryptocurrency traded in DEXs is virtually immune to theft and hacking, thereby requiring a lower return and compensation for such risks (Biais et al., Forthcoming).

To draw possible paths for improving DEX efficiency, our third contribution is to provide a simple equilibrium model that captures the main characteristics of DEX markets and the risk-return trade-off endured by Liquidity Providers (LPs). Using a unique and highly representative dataset, we test the main empirical predictions and quantify the conditions necessary for improving DEX market quality so that it becomes competitive compared to that of CEX.

While academic research on the topic is growing fast, we are the first to assess the present quality of decentralized AMM markets and their future potential, which is important for at least two reasons. First, it is a new market design that could potentially be applied to more traditional financial securities. Thus, understanding DEX characteristics could point to ways of improving the frictions of traditional markets. For instance, the fact that DEXs rely on AMM rather than LOB has at least two important implications[2]: (i) Regarding market participation, anyone, no matter who they are and what degree of sophistication they have, can offer liquidity to the exchange in a completely passive fashion through liquidity pools. (ii) Regarding welfare distribution, transaction fees are redistributed between market participants rather than collected from the exchange. In addition, the custody of assets remains entirely with the user, thus ensuring the highest level of security and censorship resistance. Second, the polit-

[2]It is often asked whether it is possible to envision an LOB-based DEX. An effective implementation of an *on-chain* order book is currently unfeasible due to technical limitations of present blockchain technology, namely the limited transaction speed and significant gas costs. In principle, however, LOB-based DEXs may be available in the future. Further, an AMM-based CEX would be feasible in practice but, to the best of our knowledge, such a solution has not yet been implemented by any prominent exchange.

ical discussion has centered on the need to regulate cryptocurrency markets in order to protect their users and ensure financial stability. This is an important issue, especially since cryptocurrencies have undergone sharp corrections in 2022. A thorough analysis of the quality of DEXs is desirable so as to address these issues properly.

Our empirical work leverages a unique and very granular dataset that comprises three elements: (i) high-frequency LOB snapshots for two of the most liquid centralized crypto exchanges (Binance, Kraken, and Coinbase), (ii) liquidity pool levels and transaction fees for the most prominent DEX (Uniswap v2 and v3), and (iii) historical gas prices for the Ethereum blockchain, computed as the average gas price over the transactions contained in each validated block.[3] This rich information allows us to accurately reconstruct quoted prices and all main cost components for a representative set of exchange pairs of cryptocurrencies at the minute frequency.

We proceed in three steps. First, we outline the crucial aspects of trading in CEX and DEX markets, including the transaction cost components (exchange fees, bid-ask spreads, and gas fees) and point out that delegating the custody of crypto-assets to a CEX involves risk (e.g., hacking and bankruptcy risk) and settlement issues, while DEX users keep their assets in non-custodial wallets. Then, we analyze market liquidity by computing effective transaction costs for each pair, decomposed by cost components, and for different sizes. For CEXs, we measure transaction costs as the volume-weighted realized half-spread based on the available limit orders (implementation shortfall), plus the percentage transaction fees charged by the exchange. Furthermore, we consider that the settlement of a CEX trade involves withdrawal fees charged by the exchange and deposit costs in the form of gas fees paid to miners. For DEXs, we consider the sum of the realized half-spread (based on the available liquidity in the pools),[4] the percentage transaction fees charged by the protocol, and the gas fees paid to miners

[3]Since the average time between consecutive blocks on the Ethereum blockchain is 13 seconds, the historical gas price series is available at relatively high frequency. However, most of the empirical analysis is carried over at the hourly or daily frequency; therefore, gas prices are averaged across multiple blocks.

[4]There are no standing limit orders in AMM-based DEXs, so the concepts of "ask price" and "bid price" are not well defined. As explained in detail in Section IV, we define the "bid/ask spread" for those exchanges as the percentage difference between the execution price and the quoted mid-price.

operating the Ethereum blockchain.

Three main findings stand out: First, DEXs generally feature similar total transaction costs to CEXs. For example, the total transaction costs of exchanging $100,000\$$ worth of ETH against USDC, USDT, Bitcoin (BTC), Chainlink (LINK), and USDC-USDT range from 40$ to 67$, from 13$ to 68$, and from 39$ to 80$ in Uniswap, Binance, and Kraken, respectively. Second, the analysis of the cost components shows that exchange and gas fees weigh relatively heavily on DEX's total transaction costs, making DEX costs less predictable. On the other hand, Uniswap offers very attractive and more stable bid-ask spreads. Third, the transaction cost is U-shaped as a function of trade size. More specifically, it is highest for small transactions ($1,000\$$), and lowest for medium transactions ($10,000\$$ or $100,000\$$). It then grows again for large transactions (1 million dollars), which applies to DEXs and CEXs.

Next, we study price efficiency by examining triangular price deviations, that is, the difference between the price of exchanging one currency for another (e.g., buy USDC against ETH) and the "synthetic" price that replicates this position by switching from another currency implies two more trades (e.g., selling ETH for USDT and then selling the latter to obtain USDC). In addition to being nearly risk-free, this arbitrage condition is an ideal laboratory for assessing price efficiency within a specific market because it captures the frictions of that specific market and not across markets or involving other instruments such as interest rates and FX derivatives as in the case of the covered interest rate parity condition.

Using the above-defined proxy for transaction costs, we compute arbitrage bounds; that is, regions in which the profits from a triangular trade would be lower than the effective transaction costs of executing that trade. We employ the size of those bounds as an inverse proxy for price inefficiency. Our empirical analysis of exchange triplets uncovers that, when accounting for transaction costs, no-arbitrage conditions are less restrictive for DEXs and result in more significant realized deviations from the theoretically efficient price.

In the third part of the paper, we outline a simple theoretical model that consistently captures the trade-off faced by DEX LPs and postulates clear empirical predictions. For a given exchange pair, expected profits arise from collected fees and are a linear function of the expected trading volume. In addition, because fees are shared among LPs proportionally to the percentage share of the owned pool, the expected return on capital is a decreasing function of pool size. On the other side of the trade-off, LPs face the risk of incurring what is known as *impermanent loss* (IL; also referred to as *divergence loss*), the AMM analog to adverse selection cost in LOB markets. Hence, our model predicts that bid-ask spreads depend on the equilibrium level of assets staked in a liquidity pool that, in turn, is proportional to the expected fees earned by the LP and inversely proportional to the expected IL. We test these empirical predictions with our data from Uniswap v2 and find strong corroboration. Furthermore, empirical data from Uniswap v3 – which allows LPs to choose among different levels of exchange fees – directly support the predictions and the main mechanism underlying our model. Furthermore, based on the insights from our model and on planned institutional changes in the DEX markets expected to lead to lower exchange fees (i.e., the introduction of Ethereum 2.0, sharding, and layer 2 solutions), we make quantitative predictions on the possible future evolution of liquidity and price efficiency of DEXs, conditional on expected levels of the trading volume. Our analysis suggests that the market quality of DEXs may soon catch up with and potentially overtake that of CEXs.

Our contribution to the literature is three-fold: First, we investigate market liquidity considering all relevant cost components of centralized and decentralized crypto exchanges. By doing so, we quantify the conditions, cost components, and transaction sizes under which DEXs become competitive with CEXs, and we are the first to inspect the crucial role of gas costs. Second, we provide a systematic analysis of price efficiency by studying the triangular no-arbitrage conditions based on a unique and comprehensive set of cryptocurrencies. The results of arbitrage boundary conditions show that DEX prices tend to be less efficient mainly because of non-trivial gas fees and exchange fees higher than those of CEXs. Third, we propose a simple equilibrium model of liquidity provision in DEX with intuitive economic predictions and solid empirical support as it explains the vast majority of the cross-sectional and time-series variation in observed liquidity levels. Through the lens of our model, expected changes in the DEX environment will improve its market quality and make it a viable and competitive alternative to the classic CEX market structure.

The rest of the paper is organized as follows. Section II presents a high-level introduction and a simplified mathematical treatment of AMM markets. Section III describes our dataset and provides summary statistics. Section IV analyses liquidity based on transaction costs. Section V studies price efficiency based on triangular arbitrage bounds. Section VI outlines our model for DEX liquidity provision, brings it to the data, and presents the resulting forecasts of future efficiency levels. Section VII reports results on the stability of transaction costs and price asymmetry. Section VIII concludes.

II. AMM Markets

A. High-level description of AMM Markets

Most exchanges in contemporary financial markets use a central LOB system that requires a central institution to maintain a record of available and executed buy orders and sell orders. In this type of exchange, the market price is determined by the most recently matched buy order and sell order. Unlike order-book-based exchanges, AMMs rely on an algorithm that automatically determines transaction and market prices based on the liquidity made available by market participants.

Implementing an LOB exchange directly on the blockchain is hardly feasible, as it would be very costly and slow due to the time-consuming mining process and gas fees paid to miners. Furthermore, current blockchain technology enjoys limited throughput, a resource that is sorely needed in order-based exchanges. Crypto exchanges such as Binance or Kraken that provide an LOB mechanism have to operate it off-chain and are thus centralized entities. This comes at the expense of the benefits offered by decentralized networks. Unlike CEXs, AMMs rely on a simple conservation function that algorithmically computes the asset price based on the liquidity available in the exchange. The most common conservation function is the so-called constant product function $xy = k$, which is used by Uniswap. In AMMs, the liquidity comes from the LPs, who deposit their assets into a smart contract representing a liquidity pool. These available reserves determine the market price of the assets and allow users to directly swap assets without having to interact with a counterparty or third party. To incentivize users to provide funds to the liquidity pools, LPs are compensated by a small exchange fee charged on each transaction. Nonetheless, providing liquidity is not free of risk. Price divergence between the time of provision and withdrawal leads to an economic loss. This arises from the fact that, by design, the LP receives more of the depreciating asset and less of the appreciating asset at the time of withdrawal. In other words, the IL is the relative loss with respect to the holding return before accounting for revenues from transaction fees. This source of risk is similar to the adverse selection faced by market makers in a market with information asymmetry; in both cases, losses occur only when flows have a permanent price impact.

The importance of DEXs has continued to increase since their inception. As of January 2021, more than $23 billion were deposited in liquidity pools across Uniswap, Sushiswap, and Pancakeswap combined. This volume is striking, given that Uniswap was created in late 2018 and was then the only DEX system. The following section discusses the advantages and inconveniences of DEXs compared to order-book-based CEXs.

B. Salient Features of CEX and DEX

DEXs based on AMM provide their users with a fundamentally different experience from standard CEXs based on LOB. Below, we discuss a number of relevant advantages and drawbacks of CEXs and DEXs that, in a nutshell, boil down to the user's trade-off between the retention of control over her funds or the benefit from LOB operational efficiency. Let us start with CEXs based on LOB, whose advantage is the ability to provide a pretty efficient and competitive price discovery process and liquidity clustering even in extreme situations (Glosten, 1994). An important drawback is that traders have to deposit their crypto-assets to the exchange to trade. Deposits and withdrawals are associated with fees charged by the exchange, gas costs required to submit the transaction to the blockchain, and significant delays. In particular, to lower the risk of potential double-spending attacks, exchanges require several block confirmations (12 on Binance and 20 on Kraken) for deposits to be accepted, leading to delays of two to four minutes. Moreover, delegated custody implies that exchanges or other central authorities can arbitrarily freeze funds and that exchanges can be hacked and users' funds stolen.[5] Such risks could lead a rational investor to demand higher compensation in equilibrium (Biais et al., Forthcoming). Finally, if exchanges mix their own funds with user funds, this exposes clients to bankruptcy risk.[6]

Regarding DEXs, the custody of assets remains fully with the user, as no third party is required to execute the trade. This benefit arising from the decentralized trust provided by blockchain technology has several important implications. First, users can take full advantage of the censorship-resistant and trustless nature of their crypto-assets (Pagnotta and Buraschi, 2018). Second, it allows users to make use of their crypto-assets in a variety of protocols and to benefit from their utilities.[7] Third, it neutralizes the risk of hackers attacking the exchange and stealing assets. Fourth, it allows users

[5]Since CEXs are not regulated, another risk that has arisen in informal discussions with crypto-asset investors is that the CEX operators themselves engage in arbitrage activities by leveraging their privileged position and inside information. This may be another reason why users prefer DEX markets.

[6]CEXs are subject to at least three sources of risk: (i) unauthorized access to crypto wallets by hackers (e.g., the cases of Poly Network and Japan-based Liquid (Ryder and ORX, 2022)), the likelihood of which increases if fraudsters use sophisticated methods such as quantum computing running encryption protocols; (ii) misappropriation of client funds by CEX managers, see e.g., Thodex and BitConnect (ORX, 2021); inefficient security management, for example, in the case where a coin exchange executive, such as Gerry Cotten of QuadrigaCX, unexpectedly dies leaving the digital vault locked (Mance, 2019).

[7]For instance, ERC20 tokens can be staked to earn interest, used as a mean of payment, posted as collateral in decentralized lending protocols, and can provide access to airdrop events.

to save on the fees commonly associated with depositing and withdrawing assets in CEXs. Finally, but very importantly, in DEXs trade and settlement coincide.

Another major innovation brought by DEXs is that their users can provide liquidity to the exchange in a completely passive fashion. This renders the market fairer given that anyone can provide liquidity, including agents with any level of sophistication, and does not require investing in expensive hardware or developing complex algorithms. By contrast, in LOB-based exchanges, LPs are usually highly specialized, and entry costs are high in terms of both sophistication and capital.[8] Market makers need high-speed computers and state-of-the-art algorithms to update their quotes as quickly as possible and avoid being picked off by high-frequency traders (Foucault et al., 2017).

The DEX market design implies that platform fees charged to each transaction are distributed to LPs in proportion to their shares (Adams et al., 2020). There is thus no welfare reduction stemming from profits accrued by the exchange itself, as there is no limited liability company associated with it. This may translate into economically significant gains for both traders and LPs.

In DEXs, the market can integrate and evolve rapidly in tune with the needs of its participants. For instance, users can quote any pair of ERC20 tokens at any time, immediately, and with no screening procedures. Consequently, new tokens are likely to be tradeable sooner in DEX, while CEX approval procedures may require significant time. Moreover, DEX may allow trading on tokens that are not available on CEXs. On the one hand, this constitutes an advantage by enlarging the space of investment opportunities, improving diversification, and speeding up the process that makes the market more complete. On the other hand, this has the drawback of exposing users to potentially malicious assets.

Finally, since DEX transactions are processed by smart contracts and directly recorded

[8]Market quality is a broad concept that includes concepts such as price efficiency, liquidity, and fairness in the sense that each agent has an equal chance of participating and obtaining a market price that reflects the fundamental value of financial security. This study focuses on the first two aspects, but given the aforementioned aspects, one can argue that the DEX setting is fairer.

on the blockchain, users bear the cost of the non-trivial gas fees required to compensate miners. This fact implies that transactions are subject to an execution delay, the duration of which depends on the speed of the underlying blockchain, the chosen gas price, and the level of network congestion. It is important to highlight, however, that for DEX trades, execution and settlement coincides, meanwhile, trades on CEXs cannot be considered settled as long as the funds are inside the exchange. Thus, if settlement issues are taken into account, trading on CEXs involves even higher fees, longer delays, and risks.

C. *Related Literature*

We contribute to the nascent but growing literature on cryptocurrencies by providing a comprehensive analysis of two main dimensions of market quality: liquidity and price efficiency. Concerning the latter, prior research provides evidence against it focusing on Bitcoin (e.g., Urquhart (2016), Bariviera (2017), and Nadarajah and Chu (2017)). Nadarajah and Chu (2017) explore a larger set of cryptocurrencies and document-wide price variation. Dyhrberg et al. (2018) assess whether and when Bitcoin is investible and at what trading costs. Hautsch et al. (2018) focus on the institutional aspect represented by the distributed ledger technology. They stress that consensus protocols to record the transfer of ownership create settlement latency, exposing arbitrageurs to price risk. Trading activity and arbitrage deviations are also at the core of the analysis in Makarov and Schoar (2020). Using tick data for 34 exchanges across 19 countries, they find arbitrage deviations of Bitcoin prices that were (i) large, persistent, and recurring, (ii) different across countries and regions, and (iii) apparently demand-driven. Using tick-level Bitcoin data from February 2013 to April 2018, Krückeberg and Scholz (2020) provide a detailed analysis of arbitrage spreads among global Bitcoin markets. Arbitrage spreads concentrate during certain periods, such as the early hours of a day and for new exchange market entries.

Regarding market liquidity, Borri and Shakhnov (2018) analyze daily data on Bitcoin prices from 109 exchanges and show that (i) daily returns are widely dispersed, and (ii) temporal variation increases with illiquidity. Brauneis and Mestel (2018) assess the market efficiency of a set of cryptos using unit root tests and by computing some liquidity proxies. They show that less liquid cryptos are less efficient. Brauneis et al. (2021) perform a comprehensive study measuring cryptocurrency market liquidity. They conduct a horse-race comparison among low-frequency transactions-based liquidity measures to ascertain which one was the closest to the actual high-frequency benchmark measure. In addition to Brauneis et al. (2021), a few other studies use order book data to study the market liquidity of cryptocurrencies. For instance, Marshall et al. (2019) find that Bitcoin endures substantial variation in liquidity across different exchanges and that changes in currency liquidity influenced Bitcoin liquidity.

We add to the literature by jointly studying centralized and decentralized crypto exchanges based on innovative blockchain-based venues that serve as a form of AMM. Thus far, only a few papers have studied AMM exchanges. On the theoretical side, Aoyagi and Ito (2021) examine the conditions for the coexistence of such CEX and DEX exchanges, and Evans (2020) outline the pay-offs of LPs on AMM exchanges. Rather than the trader's endogenous choice between CEX and DEX trading venues as in Aoyagi and Ito (2021), we assess the market quality of DEX and CEX, thus analyzing the coexistence of DEXs and CEXs in equilibrium.[9] We examine the possible asymmetric price impact between purchases and sales of cryptocurrencies to provide empirical support for the theoretical predictions of Aoyagi and Ito (2021). Evans et al. (2021) analyze the loss of privacy, worse pricing, and latency of AMM trading. By focusing on Uniswap, Angeris et al. (2019) formalize the common conditions of AMM functioning, including the need for Uniswap prices to follow the reference market price closely. Using a game-theoretic approach, Capponi and Jia (2021) model the impact on

[9]To conduct empirical research on the endogenous choice of traders, one would ideally need data revealing the identity of market participants on both CEX and DEX, which to the best of our knowledge are inaccessible, and to carefully consider what features prevent participants from simultaneously and efficiently exchanging between DEX and CEX, including aspects of interoperability. In a recent work headed towards this direction, Han et al. (2021) analyze a quasi-natural experiment and establish the causal impact of Uniswap liquidity provision on the trading activity on Binance.

utility for LPs and traders of the curvature of the pricing function on Uniswap. Park (2021) provide conditions under which "sandwich attacks" (akin to front-running) can be profitable when AMM relies on the constant product rule. On the empirical side, O'Neill (2022) shows that AMM leans toward being efficient as liquidity flows into more profitable pools, which in turn reduces fee yields and LPs offering liquidity in more volatile pools and earning higher fees.

The closest paper to our study is Lehar and Parlour (2021), which compares AMM and a limit-order market. Our work differs from theirs both in how we study market liquidity and price efficiency. Concerning liquidity, we offer a comprehensive analysis of the three transaction cost components, i.e., exchange fees, bid-ask spreads, and gas fees. In doing so, we provide four novel findings: First, although DEXs generally feature similar total transaction costs to CEXs, exchange fees and gas fees weigh relatively heavily on DEX total costs, and the latter fees are characterized by substantial time-series variation. By correctly accounting for these oscillations, transaction costs on Uniswap turn out to be only partially predictable and stable. Second, Uniswap's spreads are very competitive both in terms of average cost and narrow variation. Third, medium-size trades have a relatively low transaction costs on both DEXs and CEXs, but the transaction costs of large-size trades become more attractive on DEX since the spread's competitiveness outweighs the burden of exchange and gas fees. Fourth, we also find a significant co-movement of the effective transaction costs on Binance and the gas price on the Ethereum blockchain, which is evidence suggestive of gas prices as a proxy for the aggregate demand of immediacy. Concerning price efficiency, Lehar and Parlour (2021) conclude that Uniswap features price efficiency because the market prices for the same cryptocurrencies traded on Uniswap and Binance are very much aligned. Our analysis comprising arbitrage bounds and triangular arbitrage deviations identifies a lower price efficiency in DEX markets. Furthermore, our analysis confirms that the Uniswap and Binance prices are aligned but also shows that the Uniswap mid-price is systematically lower than the equivalent mid-price in Binance and Kraken, and this negative bias increases in downward price movements, supporting the price asymmetry hypothesis.

D. Mathematical Foundations of AMM Markets

During our sample period and at the time of writing, the majority of AMMs relied on the *constant product rule*, which enables an algebraic determination of market price and transaction price based on the available reserves (Adams et al., 2021).[10] The leading example is Uniswap, created and deployed on November 2018 by Hayden Adams, a former mechanical engineer at Siemens. In this section, we provide a brief overview of the mathematics behind constant-product AMMs.

Consider the exchange pair X/Y, and the associated liquidity pool containing x tokens of X and y tokens of Y. The combined amount of both tokens in the pool determines the current market price P_{XY} of X in terms of Y and its inverse P_{YX}, which can be expressed as

$$P_{XY} = \frac{y}{x} \qquad \text{and} \qquad P_{YX} = \frac{x}{y}$$

Let us denote as f the percentage exchange fees charged by the DEX ($f = 0.003$ for Uniswap v2), and let $\varphi = 1 - f$. These fees are immediately applied to the traded amount $\Delta x > 0$, so that the net quantity of token X that goes into the swap transaction is $\varphi \Delta x$. Each trade (swap transaction) is automatically regulated by the constant product rule, which states that the product of the reserves must remain constant before and after any transactions. Hence, when trading an amount $\Delta x > 0$ of token X in exchange for token Y, the output quantity Δy is mathematically determined through the following equation

$$xy = k = (x + \varphi \Delta x)(y - \Delta y) \ ,$$

where k is the constant product invariant. Solving for Δy, one obtains that the output

[10] There exist AMMs based on similar algebraic rules (e.g., *constant sum*). Nevertheless, according to monthly historical snapshots of the CoinGecko DEX ranking, the market share of constant-product AMMs has been above 60% in the period from January 2021 and August 2022. We include Uniswap v3 in such a category even though, formally, the constant-product rule applies only locally within each tick (Adams et al., 2021).

amount is given by

$$\Delta y = y\frac{\varphi \Delta x}{x + \varphi \Delta x}. \tag{1}$$

The transaction price is therefore lower than the quoted price and is given by

$$T_{XY}(\Delta x) \;=\; \frac{\Delta y}{\varphi \Delta x} \;=\; \frac{y}{x + \varphi \Delta x}$$

and the quoted half-spread (as a percentage of the quoted price) can be computed as

$$S_{XY}(\Delta x) = \frac{P_{XY} - T_{XY}}{P_{XY}} = \frac{\varphi \Delta x}{x + \varphi \Delta x}\ . \tag{2}$$

Note that the price impact is an increasing but concave function of the transaction volume Δx, implying that larger volumes have a larger impact on prices but with a marginally decreasing effect. *Ceteris paribus*, a purchase could have a greater impact than a sale as showed in Aoyagi and Ito (2021). We find support for this empirical prediction, as reported in Section VII. Consequently, throughout the paper, we compute the quoted half-spread for both directions ($X \to Y$ and $Y \to X$) for the same traded amount in terms of dollars and consider the average of the two measures.

E. Impermanent Loss

Similarly to liquidity provision in LOB markets, providing liquidity to AMM-based DEXs involves a trade-off between expected profits and adverse selection risk. On the one hand, LPs are compensated by pocketing the transaction fees applied to the trading volume generated by liquidity takers through swaps. On the other hand, a permanent price change leads to an impermanent loss (IL) for the LP. This loss arises from the fact that, as we will show below, providing funds to a liquidity pool is less profitable than simply holding the tokens (Loesch et al., 2021).

Consider a liquidity pool for the exchange pair X/Y, containing x_0 and y_0 units of the two tokens at time $t = 0$. Assume an LP owns a share ψ of the pool and the current quoted price is $P_0 = \frac{y_0}{x_0}$. At $t = 0$, the value of her position in units of y is

$$\psi(x_0 P_0 + y_0) = 2\psi y_0 \,.$$

At $t = 1$, the value of her position in units of y changes to $\psi(x_1 P_1 + y_1) = 2\psi y_1$. The gross percentage change in the value of the deposited liquidity can therefore be expressed as

$$R_{LP} = \frac{2\psi y_1}{2\psi y_0} = \frac{y_1}{y_0} \,.$$

Given the constant product rule $xy = k$, we can rewrite the market price P_0

$$P_0 = \frac{y_0^2}{k} \quad \Rightarrow \quad y_0 = \sqrt{kP_0} \,.$$

Similarly, at $t = 1$, the value in a unit of the LP position is $y_1 = \sqrt{kP_1}$. Hence, the change in value of the LP's liquidity depends solely on the square root of the gross price change $\Delta P = P_1/P_0$ between $t = 0$ and $t = 1$:

$$R_{LP} = \frac{y_1}{y_0} = \frac{\sqrt{kP_1}}{\sqrt{kP_0}} = \sqrt{\Delta P} \,.$$

On the other hand, the return R_H from holding the tokens is simply the average of the returns arising from holding each individual token. This equals to 1 for Y – the accounting unit – and ΔP for X. The total holding return is thus

$$R_H = \frac{1}{2}(\Delta P + 1) \,.$$

The IL, that is, the net opportunity cost from providing liquidity instead of holding the tokens, is therefore given by[11]

$$IL = R_H - R_{LP} = \frac{1}{2}(\Delta P + 1) - \sqrt{\Delta P}\,. \tag{3}$$

By taking the first order derivative with respect to ΔP, one can easily see that IL has a global maximum of 0 for $\Delta P = 1$, while it is strictly positive otherwise. Hence IL represents a cost and highlights that, gross of pocketing the fees, LPs providing liquidity are always worse-off than token holders. We note that IL can be seen as a measure of the level of adverse selection faced by LPs, similar to that faced by market makers in LOB markets. In fact, for any given horizon, $IL = 0$ if the order flow is uninformed and gives only rise to a temporary price impact ($\Delta P = 1$), while it increases in magnitude in the presence of informed order flow, causing a permanent price change ($\Delta P \neq 1$). Intuitively, the provision of liquidity in the AMM framework can be remunerative if the LP provides immediacy primarily to liquidity traders while, at the same time, it may involve net losses when facing a higher fraction of informed traders. Quantitatively, Fukasawa et al. (2022) show that the IL in constant-product AMMs can be hedged through weighted variance swaps.

III. Data and Summary Statistics

Because DEXs are based on smart contracts deployed on blockchains, records of every single interaction with those contracts are available to the public. This rich dataset includes, as primitives, the creation of exchange pairs, the addition and removal of liquidity from LPs, and swap transactions between two quoted tokens. Building on those, one can reconstruct liquidity levels, quoted prices, transaction prices, and trading volume at the pair level at any time. We leverage the application programming interface of TheGraph.com in order to obtain data for Uniswap from the Ethereum MainNet

[11]We define the IL as the *difference* between R_{LP} and R_H, as in Aigner and Dhaliwal (2021) and Fukasawa et al. (2022). The IL can alternatively be defined in percentage terms (Khakhar and Chen, 2022; Heimbach et al., 2022, 2021) by $IL = R_{LP}/R_H - 1 = 2\sqrt{\Delta P}/(\Delta P + 1) - 1$. Our equilibrium model and the empirical results described in Section VI are robust to both definitions.

blockchain. We downloaded data on liquidity pool reserves and volumes at an hourly frequency for the pairs made of the five crypto-tokens that are the subjects of our analysis.

For CEXs, by contrast, data are proprietary. We obtain minute-frequency Open, High, Low, Close, and Volume data and full LOB snapshots from Kaiko for all pairs quoted on the largest crypto exchanges in terms of traded volume, including Binance International and Kraken. For both CEXs and DEXs, our sample period spans from January 2021 to December 2021.[12] For our main empirical analysis of market quality, we focus on five of the most liquid and traded cryptocurrency pairs, namely ETH-USDC, ETH-USDT, ETH-BTC, LINK-ETH, and USDC-USDT. To select these pairs, we first consider the intersection of trading pairs present in Uniswap, Binance, and Kraken. We then compute percentile rankings of average daily volumes in USD over our sample period on each of the three exchanges and take their average. Finally, we select the first five pairs ranked by the resulting metric.[13]

A. Trading Volume

The first panel of Figure 1 reports trading volumes for the LOB-based exchanges in our sample (Binance, Kraken, and Coinbase). Binance is the dominant exchange in terms of volume across the entire sample, rising from roughly $4 billion to $23 billion. Volumes on Coinbase and Kraken are comparable; both present a significant upward trend. The second panel of Figure 1 displays daily volumes for three of the most liquid AMM-based exchanges (Uniswap v2, Pancakeswap, and Sushiswap). For Uniswap v2, which was deployed on the Ethereum MaiNnet on May 2020, the plot shows a 10-fold increase from around $100 million on August 2020 to roughly 1 billion at the end of the sample (December 2021). Similar upward trends are displayed for Sushiswap and

[12]The market has consolidated, and liquidity has reached relatively high levels only since Fall 2020.

[13]Obviously, this method excludes some currencies that are not simultaneously traded on the two types of markets. It does, however, make it possible to compare the same cryptocurrencies with the same fundamental reference value using high-frequency and granular data.

Pancakeswap, which were deployed in September 2020. The bottom panel of Figure 1 displays trading volumes for both the AMM- and LOB-based exchanges in our sample, averaged across the three exchanges in each category. The average DEX volume rises sharply by about two orders of magnitude within the sample period, while the average CEX volume shows roughly a 10-fold increase in the same period. All in all, the data show that trading volume has been increasing sharply for all the CEXs and DEXs in our sample. Even though the DEX increase is significantly steeper, the wedge within the two categories remains around one order of magnitude at the end of our sample.

Table II presents the daily average trading volume for the pairs we consider in millions of USD over the January-December 2021 period, which is the focus of our market quality analysis. These pairs provide a representative sample, as they generate between 15% and one quarter of the volume on each exchange.

B. Gas Prices

The term *gas* refers to the unit of measure of the computational effort required to execute transactions on the Ethereum network. Gas fees are paid in the network's native currency (ETH in the case of the Ethereum network). The aggregate gas fees, summing over all block transactions, are pocketed by the miner validating the block. To trade on a DEX, the user has to pay a number of gas units proportional to the computational complexity of the transaction. The user is free to decide the *gas price* in order to control the degree of priority of execution. Miners choose the set of pending transactions to include in the new block, prioritizing the most profitable transactions, that is, those with the highest gas price. Wallet interfaces automatically suggest an optimal gas price, depending on the current level of network activity and based on the trade-off between the probability of execution within the next 1-2 blocks and the cost. While users can edit the gas price according to their preferences, most transactions in our sample are executed at the suggested gas price.

Gas costs are undoubtedly important in the study of trading in fully decentralized AMM markets since all interactions are *on-chain*. Every trade is triggered by a transaction submitted by the user to the blockchain, representing a function call on the protocol's smart contracts.

However, on a closer inspection, gas fees also matter in CEX trading. In fact, a CEX-based trade is only settled when the trader transfers the asset from the exchange wallet back to their non-custodial wallet. The settlement process thus involves the gas fees paid to miners for the deposit of funds ("deposit fees") and fees charged by the exchange for the withdrawal ("withdraw fees"). The transaction representing the deposit isx executed by the user, who has to pay the gas directly to the network, while the withdrawal operation isx executed by the exchange, which pays for the gas and requires a compensation from the user.

Figure 2 plots the evolution of the gas fees (in USD) required to execute a swap transaction on Uniswap v2. Since the amount of required gas for such an operation is constant over time, the observed time-series variation arises from two factors: (i) the price of ETH relative to the USD and (ii) the prevailing gas price of the network, depending on the level of network congestion. To construct the time series, for each hour, we take the product of the number of required gas units (constant at 110,000) and the average gas price across all transactions validated during that hour. Finally, we convert the amount into US Dollars using hourly ETH prices. The series presents substantial variability, ranging from 1$ in the first part of the sample to 400$ between April and May 2022.[14] The two vertical lines mark the period from January 2021 to December 2021, the period in which our primary analysis of market quality is conducted.

C. Impermanent Loss

We construct hourly realized values of IL based on equation (3), for 100 of the most liquid pairs traded on Uniswap v2 over the period between April 2020 and April 2021.

[14] The exceptional average gas price observed between April 30th and May 1st, 2022 was caused by the NFT drop of the "Otherside". The highly awaited launch of the collection by "Yuga Labs" the company behind Bored Ape Yacht Club and ApeCoin, resulted in more than $150 million spent on gas fees.

Panel C of Table I provides summary statistics on the resulting IL, expressed in basis points. In the first row, we report the distribution across the entire panel, while in the following rows the distribution is collapsed at the pair-level and day-level by taking averages. This panel dataset will be used to proxy for expected IL and bring to the data the equilibrium model outlined in Section VI.

IV. Transaction Costs

One dimension of market quality is market liquidity; that is, the ease with which an asset can be traded at a price close to its consensus value (Foucault et al., 2013b). As a proxy for market illiquidity, we employ the effective transaction costs associated with a single trade, expressed as a percentage of the traded amount. These account for both the price impact associated with a given trade size and any kind of commissions charged by the protocol or the exchange. Due to their fundamentally different mechanics, transaction costs on LOB and AMM markets are modeled using distinct methodologies. Nevertheless, the two measures are based on the same conceptual framework, as they are meant to capture the effective costs incurred by a trader transacting a given amount (in US Dollar terms), including slippage, fees, and settlement costs.

Empirically, we estimate transaction costs TC_{XY} for a trade $X \leftrightarrow Y$ in both AMM and LOB exchanges at the hourly frequency for the five pairs in our sample and different trade sizes $(10^3, 10^4, 10^5, 10^6)$, expressed in USD.

A. CEX Transaction Costs

For CEXs based on LOB, we measure transaction costs of a market order by considering four distinct components: (i) the bid/ask spread implied by the depth of the LOB, (ii) the exchange fees charged by the exchange (taker fees), (iii) the gas fees paid to transfer crypto tokens to the exchange, and (iv) withdrawal fees charged by the exchange. The third and fourth components constitute a measure of settlement costs, motivated by the assumption that the trader does not delegate the ownership of its funds to the exchange.[15] Rather, we assume that the trader holds the crypto tokens in her non-custodial wallet and transfers token X to the exchange whenever she wants to trade. After the transaction $X \to Y$ occurs inside the exchange, she withdraws the resulting units of the Y token by transferring them to her wallet. These deposit and withdrawal operations are expensive and will be discussed below.

We start with the Bid/Ask spread associated with a market order which, since we observe the full depth of ask and bid quotes present in the order book at any point in time, can be computed directly using the volume-weighted bid and ask prices. More specifically, we define the volume-weighted bid price B_{XY} for a sell order of size Δx as

$$B_{XY}(\Delta x) = \frac{\sum_i v_i b_i}{\Delta x} \qquad \text{such that} \qquad \sum_i v_i = \Delta x \ ,$$

where v_i and b_i represent the volume and the price of each filled bid limit order i. The volume-weighted ask price A_{XY} for a buy order of size Δx is defined symmetrically as

$$A_{XY}(\Delta x) = \frac{\sum_j v_j a_j}{\Delta x} \qquad \text{such that} \qquad \sum_j v_j = \Delta x \ ,$$

where v_j and a_j represent the volume and the price of each filled ask limit order j, respectively. We then define the percentage half-spread by

$$S_{XY}(\Delta x) = \frac{A_{XY}(\Delta x) - B_{XY}(\Delta x)}{A_{XY}(\Delta x) + B_{XY}(\Delta x)} \ . \tag{4}$$

Next, the deposit of funds involves the trader paying an amount of gas fees to submit the transaction on the blockchain. The cost of such an operation is fixed in terms of the required gas units (21,000 for the native token ETH and 65,000 for ERC20 tokens). Its dollar value depends on the prevailing gas price on the network.[16] We, therefore, define

[15] As argued in Section I, centralized crypto exchanges are exposed to several operational risk factors. Hence, a trade cannot be considered settled as long as the assets are held by the exchange.

[16] While the trader can choose a custom value for the gas price for each transaction, which determines its

the deposit fees as the product between the number of gas units required to execute the transaction, which is dependent on whether X is the native currency or not, and the prevailing gas price (in USD terms) at the time of the trade.

Furthermore, the withdrawal from the exchange to a custodial wallet incurs a fee charged in units of the withdrawn currency. We collect from Binance and Kraken the value of those fees for each token of interest and express it in USD by considering the dollar value of the token at the time of the trade.

Finally, we define the total transaction costs by adding up the four components defined above. For the sake of simplicity, we condense the deposit and withdrawal fees into one term $DW_{X,Y}/\Delta x$ representing the entirety of settlement costs as a percentage of the traded amount, while we leave the percentage transaction fees f as a separate entity. We thus have

$$TC_{XY}(\Delta x) = S_{XY}(\Delta x) + f + \frac{DW_{XY}}{\Delta x} \ , \tag{5}$$

Notice that the first and last terms in the above expression are time-varying, while exchange fees are constant.

B. DEX Transaction Costs

For DEXs based on AMM, we measure transaction costs of a trade by considering three distinct components: (i) the bid/ask spread implied by the depth of the liquidity pools as derived in (2); (ii) the exchange fees charged by the exchange; (iii) the gas fees paid to submit the blockchain transaction. In general, the dollar value of those fees depends on the computational complexity of the smart-contract function being called, the execution priority chosen by the trader, and the prevailing gas price at the execution time. For our purposes, we are interested in the gas required to execute a swap transaction; that is, invoking the `swapExactTokensForTokens` function of the

priority on the network, we assume that the trader picks the gas price suggested by the wallet interface, that is, the prevailing gas price in the network.

relevant router contract.[17] We assume that the quantity of gas required to execute a swap transaction is constant across all currency pairs at $\Gamma = 110{,}000$ gas units.[18] We then approximate the gas cost g of a swap during each hour of our sample period by multiplying Γ by the average gas price paid across all blocks verified during that hour in dollar terms.

The transaction costs on AMM markets are thus computed as the sum of the quoted half-spread S_{XY} defined in (2) and averaged across both directions $X \to Y$ and $Y \to X$, the percentage exchange fee f (30 bps for Uniswap v2), and the gas fee g as a fraction of the trade size:

$$TC_{XY}(\Delta x) = S_{XY}(\Delta x) + f + \frac{g}{\Delta x}. \tag{6}$$

As in the CEX case, the first and last terms in the above expression are time-varying, while exchange fees are constant.

C. *Results*

The main results of the transaction cost analysis are displayed in Figure 3 , showing log transaction costs for different trade sizes on the AMM-based DEX Uniswap and the LOB-based CEXs Binance and Kraken. These are computed at the hourly frequency using (6) and (5), and then averaged across the entire sample period. In addition, Table III provides a breakdown of the total transaction costs into their three separate components (bid-ask spread, gas or DW fees, exchange fees) for different trade sizes. Table IV presents more detailed descriptive statistics, including the time-series variance of each cost component for a $10,000$\$ trade.

Three main findings stand out: First, by comparing the total transaction costs in

[17]Depending on the nature of the token, the exact router function may be different. For instance, for tokens featuring fee re-distribution like SafeMoon, the `swapExactTokensForTokensSupportingFeeOnTransferTokens` function must be used. Nevertheless, the amount of gas required is not significantly different.

[18]We estimate 110,000 by collecting all swap transactions executed on the Uniswap v2 Router contract using the latest 1,000 blocks as of July 1, 2021, and taking their average gas usage. The variation across pairs is minimal for the pairs we consider ($\pm$ 10,000 gas units at most). This figure is significantly larger than the gas required by a simple *transfer* function on an ERC20 contract which costs 65,000 gas units, or a transfer of ETH (the native currency of the Ethereum blockchain) which costs 21,000 gas units.

general, we observe that those on DEXs are comparable to those on CEXs. For instance, the estimates in Table III show that the total transaction costs of exchanging \$100,000 worth of Ethereum (ETH) against USD Coin (USDC), Tether (USDT), Bitcoin (BTC), Chainlink (LINK), and USDC-USDT range from 40\$ to 67\$, from 13\$ to 68\$, and from 39\$ to 80\$ in Uniswap, Binance, and Kraken, respectively. Although a closer look indicates that Binance is the most convenient market venue, these figures suggest that Uniswap offers competitive transaction costs.

Second, the analysis of the cost components in Table III shows that two out of the three components are relatively more expensive in Uniswap. First, Uniswap' s exchange fees are three times larger than those of Binance and slightly higher than those of Kraken. Second, gas fees in Uniswap are on average larger than CEX DW fees. Table IV indicates that they also vary much more than DW fees. On the other hand, Uniswap offers very attractive average bid-ask spreads, comparable to those of Binance and smaller than Kraken's. Uniswap's bid-ask spreads, moreover, are more stable and exhibit lower volatility compared to those of Binance and Kraken.

Third, the total cost follows a U-shape as trade size increases, that is, it is the highest for small transactions ($1,000$\$), lowest for medium transactions ($10,000$\$ or $100,000$\$), and then grows again for large transactions (\$1 million). Figure 3 shows that this pattern applies to both DEX and CEX, but it is more prominent for the former. For instance, a $1,000$\$ trade in Uniswap is remarkably expensive as it costs roughly 300 bps for all considered pairs, while the same transaction on CEXs spans between 150 and 250 bps (top-left panel of Figure 3). This finding does not come as a surprise since, as reported in Table III, gas fees on Uniswap constitute a large percentage (2.67% on average) of the traded amount. Similarly, deposit and withdrawal fees are the most impactful component on CEXs for this trade size, even though they are lower in absolute value on average. As the size of the transaction increases, however, the exchange fees and gas costs become marginal while the spread plays a more predominant role, thus rendering Uniswap more competitive. For instance, while Binance is still the most convenient trading venue for a $100,000\$$ trade in four out of five pairs, Uniswap delivers lower transaction costs for the LINK-ETH pair and similar costs relative to Kraken for the other pairs. This finding is corroborated by Table III showing that gas costs and DW fees are only marginally important, while spreads and exchange fees play a much more critical role. This pattern becomes even more evident when we consider the bottom-right panel of Figure 3 , showing that Uniswap is cheaper than Kraken for most pairs and cheaper than Binance for ETH-USDC and LINK-ETH. Nevertheless, Binance is cheaper for other pairs, particularly ETH-USDT and USDC-USDT.

To sum up, DEXs feature competitive transaction costs, especially as the trade size increases. While the relatively high exchange fees and gas costs weigh more on small transactions, DEXs provides attractive bid-ask spreads, which contribute more to the cost of large transactions. Among the three trading venues analyzed, Binance offers the lowest average transaction costs for all trade sizes, especially for the pairs involving the stablecoins USDT and USDC. At the same time, Uniswap provides the lowest and most stable bid-ask spread for most of the pairs and trade sizes, highlighting the efficiency of the AMM model based on liquidity pools once gas fees and exchange fees are muted.

D. Uniswap v3

A new version of the Uniswap protocol, Uniswap v3, was introduced on May 5th, 2021. The upgrade, deployed through a new set of smart contracts,[19] introduces the possibility for LPs to *concentrate* liquidity on specific price ranges. This innovation implies that the bonding curve is only *locally* defined by a constant-product rule and makes the price impact for a given trade size dependent on the entire distribution of liquidity positions rather than simply on the aggregate liquidity levels. Furthermore, Uniswap v3 allows LPs to choose between four pools for the same trading pair with different levels of exchange fees (1, 5, 30, and 100 bps). This new feature provides us

[19] A comprehensive list of the address of each deployed contract is available on the official Uniswap documentation at https://docs.uniswap.org/protocol/reference/deployments

with a unique laboratory to test the predictions of our theoretical model of liquidity provision, exposed in Section VI, predicting that LPs should accept low exchange fees if the trading volume is high, if the IL is low, or a combination of both.

While we do not have access to comprehensive data from Uniswap v3 in order to estimate transaction costs systematically, we provide a limited analysis based on real-time data. We collected the data leveraging a custom Ethereum node, based on the Geth client. From July 23rd, 2022, to September 11th, 2022, every 30 minutes, we run a script querying data from the Uniswap v3 contracts and computing the effective transaction costs for trade sizes of $(10^3, 10^4, 10^5, 10^6)$ USD for the five pairs we consider in the paper. Similarly, we collected real-time LOB data from Binance and Kraken over the same period and at the same frequency and computed the resulting transaction costs for the pairs and trade sizes of interest. We then apply the same methodology employed in the main analysis, computing transaction costs for Uniswap v3 and the two CEXs using equations (6) and (5), respectively. Whenever multiple pools (with different exchange fees) are available for a trading pair, we choose the one with the highest total volume traded since inception.

Figure 4 presents the results, displaying the estimated transaction costs expressed in basis points on a log scale for the different trade sizes. Differently from Uniswap v2, the new version of the DEX is competitive with the CEXs regarding the cost of trading of *any* size, and it is superior to the centralized counterparts for 13 out of 20 pair-size combinations.

Interestingly, the USDC-USDT pair involving two stablecoins pegged to the US Dollar enjoys extremely low transaction costs on Uniswap v3, one order of magnitude lower than on the CEXs. This is due to the highly concentrated liquidity in the pool, in a tiny interval centered at the exchange rate of 1, allowing for very low spreads even for substantial trade sizes. Furthermore, the pool has exchange fees as low as 1 basis point, justified by the extremely low volatility of the pair. This finding is in line with the predictions of our theoretical model, described in Section VI: LPs accept a very low level of exchange fees, rationally responding to the low level of expected impermanent loss and the high trading volume in this exchange pair. These results, even though they are based on a relatively short sample period, provide suggestive evidence for the future potential of DEX protocols.

While Uniswap v2 is still functional today, after the introduction of Uniswap v3, liquidity may have gradually moved from v2 to the new version of the DEX. Given these dynamics, our results on v2 transaction costs may be biased upward because of the lower liquidity on v2 in the second part of the sample. Therefore, we re-run the above analysis, cutting our sample on May 4th, 2021. We find the results based on the subsample are virtually identical, with no significant differences in transaction costs across all the trade sizes (not reported).

V. Price Efficiency

Finite liquidity and transaction fees constitute frictions limiting arbitrage forces, allowing deviations from efficient prices to persist and blurring the informativeness of transaction prices. We explore deviations from the law of one price by focusing on triangular arbitrage and relating it to liquidity levels. Performed in only one specific market and nearly risk-free, this arbitrage condition is the ideal laboratory for identifying market-specific frictions and for comparing the price efficiency of different market venues. A triangular arbitrage opportunity arises when the law of one price is violated for a closed triplet of currency pairs X/Y, Y/Z, and Z/X. A direct measure of the deviation from price efficiency in this context is the deterministic function of liquidity levels θ, defined as

$$\theta = P_{XY}\, P_{YZ}\, P_{ZX} - 1\,, \tag{7}$$

where P_{AB} is the quoted price of A in units of B. A situation in which $\theta \neq 0$ does not necessarily imply the existence of an arbitrage opportunity since an arbitrageur faces price impact and transaction fees. The idea behind our definition of arbitrage bounds is that, at each point in time, a triangular trade is profitable only if the deviation from the efficient price is sufficiently large. In other words, the net expected profits θ of a triangular trade have to be higher than the costs of executing the three associated transactions. Assuming that arbitrage opportunities do not arise in equilibrium, the observed price levels should never allow for such a triangular trade to be profitable. We can thus derive a mathematical expression for arbitrage bounds by imposing the no-arbitrage condition (in the spirit of Hautsch et al. (2018)).

We first define and compute the cumulative execution cost $K(\Delta x) > 0$ of a triangular trade in a given triplet; that is, executing three transactions: $X \to Y$, $Y \to Z$, and $Z \to X$. Two components of such a cost, regardless of exchange type, are related to the spread and the transaction fees. For AMM markets, we also have to consider a third component, the gas fees, discussed below. The total quoted spread for a triangular trade on X, Y, and Z is given by

$$S_{XYZ}(\Delta x) = 1 - \Big(1 - S_{XY}(\Delta x)\Big)\Big(1 - S_{YZ}(\Delta y)\Big)\Big(1 - S_{ZX}(\Delta z)\Big), \tag{8}$$

where the input quantities for the second and third transaction are, respectively,

$$\Delta y = \Delta x \cdot T_{XY}(\Delta x) \qquad \text{and} \qquad \Delta z = \Delta y \cdot T_{YZ}(\Delta y)\,.$$

Note that equation (8) is simply the sum of the three spreads – associated with each transaction of the triangular trade – appropriately *discounted*; that is, adjusted to account for the fact that input amounts of the second and third trades are smaller than Δx as a result of the spreads of the previous transactions. The total fees charged, as a percentage of the initial amount Δx, are

$$F_{XYZ}(\Delta x) = f\Big(1 + (1 - S_{XY}(\Delta x)) + (1 - S_{XY}(\Delta x) \cdot (1 - S_{YZ}(\Delta y))\Big)\,, \tag{9}$$

where f are the percentage fees charged by the exchange for a single transaction. Given the execution cost $K(\Delta x)$, triangular arbitrage is profitable if and only if

$$\theta > K(\Delta x) \quad \text{or} \quad \theta < -K(\Delta x)\,,$$

and arbitrage bounds for that triplet are defined as $\theta^H, \theta^L = \pm K(\Delta x)$. Since the level and the nature of transaction costs depend on the structure of the exchange, we define empirical proxies for triangular arbitrage bounds separately for AMM and LOB exchanges.

A. Arbitrage Bounds for CEXs

Arbitrage bounds on LOB-based CEXs depend on quoted spreads, defined in (4) and based on the available liquidity in the LOB and the transaction fees charged by the exchange. The total spread (also called implementation shortfall) and the total fees charged as a percentage of the initial amount Δx are defined as in (8) and (9), respectively.

We assume that arbitrageurs keep their arbitrage capital readily available inside the CEX so that they are not subject to deposit and withdrawal fees. The main reason for this assumption is that moving capital from a non-custodial wallet to a CEX takes a significant amount of time.[20] Hence, given a competitive environment, arbitrageurs have an incentive to delegate the custody of their arbitrage capital. Our assumption will be justified a posteriori since, as we show in Section V, the resulting arbitrage bounds are consistent with the observed triangular price deviations.

Thus, the execution cost of a triangular trade of size Δx does not include any deposit or withdrawal fees, and is defined by

$$K(\Delta x) \;=\; S_{XYZ}(\Delta x) \;+\; F_{XYZ}(\Delta x). \tag{10}$$

[20] Deposits from a non-custodial wallet to a CEX take one to five minutes to be executed. The reason is that the exchange initially freezes the assets and requires the user to wait for a predefined number of blocks to be validated on the blockchain (12 for Binance, 20 for Kraken) before the funds are accessible. This measure is in place to decrease the probability of a double-spending attack that would result in a net loss for the exchange.

We further assume that the trade size is infinitesimal because of the absence of fixed transaction costs, as the marginal arbitrageur may accrue positive profits while performing arbitrage trades of very small dollar amounts. Hence, the lower and upper arbitrage bounds for θ are given by

$$\theta^H, \theta^L = \pm\Big(S_{XYZ}(\Delta x) \ + \ F_{XYZ}(\Delta x)\Big) \quad \text{with} \quad \Delta x \to 0\ . \tag{11}$$

B. Arbitrage Bounds for DEXs

Arbitrage bounds on AMM markets depend on (i) the quoted spread S, defined in (2) and based on the liquidity available in the three pools; (ii) the exchange fees f charged by the exchange; (iii) the gas fees g associated with the interaction with the underlying blockchain (Ethereum MainNet in the case of Uniswap). The total quoted spread and the total fees charged, as a percentage of the initial amount Δx, are defined as in (8) and (9), respectively. A rough approximation for the total gas required to execute a triangular trade is simply the gas cost of a single swap multiplied by 3. However, arbitrageurs are known to use custom smart contracts to execute the trade within a single function call and to condition the execution on the relevant set of conditions. Hence, to better estimate the quantity of gas required for such an operation, we deployed a basic smart contract capable of executing a triangular trade on the Ethereum MainNet and found that the cost is approximately $g_3 = 241,822$ gas units.[21] Notice that the gas cost of executing the triangular trade through the smart contract is roughly 30% lower than the cost $3g$ executing the three transactions separately. Thus, we describe the total execution cost as (see Appendix A for more details)

$$K(\Delta x) \ = \ S_{XYZ}(\Delta x) \ + \ F_{XYZ}(\Delta x) \ + \ g_3/\Delta x. \tag{12}$$

[21]The deployed contract and its source code are available on Etherscan at this address. A triangular trade was executed and recorded on the blockchain, which is available here.

As is the case with models that have entry costs, arbitrageurs face a trade-off between the cost of gas fees and the price impact. The former is reduced (in %) by increasing Δx, while the latter increases with Δx. Assuming rationality, they choose the optimal trade size $\Delta x^\star$ for which the cost $K(\Delta x)$ is minimized. We solve the optimization problem numerically, finding the optimal $\Delta x^\star$ for each situation in our panel. We then compute the percentage loss by making such an optimal trade; that is, $K(\Delta x^\star)$. Hence, the lower and upper arbitrage bounds for θ are given by

$$\theta^H, \theta^L = \pm\Big(S_{XYZ}(\Delta x^\star) \; + \; F_{XYZ}(\Delta x^\star) \; + \; g_3/\Delta x^\star\Big) . \tag{13}$$

C. Arbitrage Bounds and Price Efficiency

The width of the region between the above-defined arbitrage bounds can be thought of as a proxy for the severity of price inefficiencies. More precisely, we consider the half-width, computed as

$$B = \frac{\theta^H - \theta^L}{2} . \tag{14}$$

Wider bounds for a given triplet imply that the relative prices deviate more from the efficient ones before arbitrageurs can make a profitable arbitrage trade and push the prices closer to efficient levels. We construct bounds at the daily frequency for the five triplets in our sample, separately for each exchange. We then compare the bounds to the realized price deviations θ at the hourly frequency and find that the quoted prices are within the bounds for the vast majority of the observations, thus validating the empirical relevance of our proxy. As an example, graphical representations of the resulting bounds for the triple USDC-USDT-ETH are provided in Figure 6.

D. Results

We estimate arbitrage bounds at the hourly frequency for the five triplets in our sample and then take the average over the period from January 2021 to December 2021. The calculation is based on (11) for the LOB-based Binance and Kraken and on (13) for the AMM-based Uniswap. Figure 5 presents the results, displaying the log-levels of price inefficiency for each triplet, as proxied by the size of arbitrage bounds defined in (14). It is evident that Uniswap is far less price-efficient than its centralized counterparts. For the most liquid triplets (ETH-USDC-USDT and BTC-ETH-USDC), the width of Uniswap's arbitrage bounds sits between 100 and 200 bps, while it rises above 700 bps for the less liquid ones. These estimates are significantly higher than those for CEXs, which are below 100 bps for almost all of the considered triplets. Binance is particularly dominant regarding price efficiency, with tight bounds ranging from 30 to 50 bps.

There are two main reasons for the wider arbitrage boundaries in DEXs. First, such a significant discrepancy between DEXs and CEXs relies on the different levels of exchange fees. The 30 bps charged by Uniswap are higher than those charged by CEXs (10 bps for Binance and for 26 bps for Kraken). As triangular arbitrages require three transactions, these wedges add up to a significant determinant of the net profitability of the trade.

The second – and quantitatively most important – reason for such a low level of price efficiency affecting Uniswap is the high level of gas fees required to submit transactions on the Ethereum blockchain. In order to make up for such significant fixed costs, triangular arbitrage on ETH-based AMM markets requires trading sizeable amounts in terms of USD. This, in turn, means that arbitrageurs have to bear significant trading costs arising from their temporary price impact. Such a limit to arbitrage is a direct consequence of proof-of-work, that is, the cryptographic consensus protocol currently employed by the Ethereum network. Miners have to cover the costs of expensive hardware and significant energy consumption, thus requiring high gas prices in equilibrium.

On the contrary, no fixed costs are charged on CEXs, since transactions are recorded in their internal databases rather than on the blockchain. This allows arbitrageurs to exploit triangular arbitrage opportunities by transacting even infinitesimally small amounts. Arbitrage bounds are only slightly larger than transaction fees multiplied by

3, suggesting that trading costs arising from quoted spreads are not as relevant.

We note that, while CEXs enjoy higher levels of price efficiency as measured by triangular arbitrage bounds, they may not offer a better experience, from an operational perspective, with respect to DEXs. In fact, taking the perspective of a CEX user keeping her funds on a non-custodial wallet, the significant time delay affecting deposits induces a trade-off between price efficiency and speed of execution. Conversely, even if DEX quoted prices are further away from their efficient levels, the user can settle a transaction with a significantly shorter time delay.[22]

VI. Conditions for DEX Dominance

In this section, we present a simple theoretical model of liquidity provision in AMM markets, highlighting the main economic trade-off faced by LPs. Solving the model reveals a direct link between the level of liquidity available in the pools, the fees charged by the protocol, and the total trading volume by market participants. We show that the derived relationship holds strongly in the data. This result allows us to pin down the quantitative conditions of the required growth rate of future trading volume to make Uniswap competitive with CEXs when it comes to market quality.

A. Equilibrium Liquidity

We model a marginal LP that faces the problem of providing the optimal quantity of liquidity to the exchange pair X/Y. We assume that the LP is risk-neutral and that the market is perfectly competitive as in, e.g., Glosten and Milgrom (1985). At time $t = 0$ the total liquidity in the pools is equal to x, and the LP can add or remove a quantity ξ of liquidity. At time $t > 0$, users start to trade on the pair until the trading stops at $t = 1$. Let the random variables V denote the total traded volume (in units

[22]Assuming the chosen gas price is reasonable compared to the prevailing gas price on the network, the transaction is settled within one block, which is significantly less than the 12 blocks required by Binance or the 20 blocks required by Kraken.

of X) and let $\Delta P = P_1/P_0$ denote the gross percentage change in the quoted price, respectively, between $t = 0$ and $t = 1$. As discussed in Section II, the profits and losses of the LP depend on two factors: the fees arising from liquidity takers' trading volume and the impermanent loss due to changes in quoted prices.

B. Equilibrium Model of Liquidity Provision

Let $\mathrm{E}[V]$ denote the expected unsigned trading volume and let $\mathrm{E}[IL]$ denote the expected IL, both estimated at time $t = 0$. The expected fees paid by liquidity takers amount to the product of the exchange fees f and the expected volume $\mathrm{E}[V]$, expressed in units of X. Since this amount is distributed by the protocol to the participating LPs on a pro-rata basis, the marginal liquidity provider depositing ξ units of additional liquidity gains in expectation $\frac{\xi}{x+\xi} f\,\mathrm{E}[V]$ units of X, corresponding to a percentage profit of $\frac{f\,\mathrm{E}[V]}{x+\xi}$. Hence, accounting for both fees and the expected impermanent loss, the net expected percentage profit $\mathrm{E}[R]$ from providing an additional amount ξ of liquidity is equal to

$$\mathrm{E}[R] = \frac{f}{x+\xi}\,\mathrm{E}[V] - \mathrm{E}[IL]\,.$$

The assumption of perfect competition results in zero expected profits for the LP, hence the equilibrium level of total liquidity $x^* = \xi + x$ is given by

$$x^* = \frac{f\,\mathrm{E}[V]}{\mathrm{E}[IL]} \tag{15}$$

showing that total liquidity increases with the expected trading volume and (percentage) exchange fees remunerating the LP, while it decreases with the expected IL. The equilibrium condition (15) has a clear economic interpretation that is conceptually related to standard microstructure models featuring market makers. First, the level of liquidity x^* provided by LPs determines the quoted spread available to traders, as in (2). Second, as noted above, the expected IL can be thought of as a proxy for the level of adverse selection risk faced by LPs. Thus (15) says that spreads are increasing in the level of adverse selection; in other words, LPs require compensation for the losses caused by informed trading.

C. *Empirical Fit*

We use daily liquidity data to test the predictions of our model, proxying for $\mathrm{E}[V]$ with the rolling average of daily traded volume and for $\mathrm{E}[IL]$ with the rolling average of the daily IL, estimated over the previous two weeks.[23] We regress daily log values of empirically observed liquidity on the ones predicted by (15), for 100 exchange pairs from April 2020 to April 2021. Results are reported in Table V and Figure 7, showing a highly significant positive correlation between predicted and observed liquidity levels, with a remarkable R^2 coefficient equal to or higher than 92%. The results are robust to the inclusion of pair- and time-fixed effects. Our results are also robust to changing the size of the rolling windows used to estimate $\mathrm{E}[IL]$ and $\mathrm{E}[V]$. In particular we use 5 days and 20 days, corresponding to the median and average duration of liquidity positions reported in O'Neill (2022), and we find that the resulting R-squared for the baseline specification is 91.38% and 92.50%, respectively. We thus conclude that our simple equilibrium model is empirically relevant, as it is able to capture the main economic trade-off faced by LPs in AMM-based DEXs. At the same time, our findings suggest that LPs behave rationally in the aggregate.

D. *Dominance Conditions*

Re-arranging equation 15, we can link trading volume to liquidity and exchange fees as

$$\mathrm{E}[V] = \frac{-\,\mathrm{E}[IL]x^*}{f} \propto \frac{x^*}{f}. \tag{16}$$

In particular, the equation implies that an exogenous increase in trading volume should lead to an increase in equilibrium liquidity x^*, a decrease in fees f (if allowed by the

[23] We report summary statistics for realized IL in Panel C of Table I.

protocol), or a combination of the two. This makes intuitive sense; since higher trading volume corresponds to more fee proceedings being pocketed by LPs, their incentive to provide liquidity would still be positive after a decrease in f, thus reducing the proceedings to the previous equilibrium level, or after an increase in x^*, thus reducing the expected returns per additional unit of liquidity provided. Therefore, we can use the aforementioned relationship to derive conditions on the time-series dynamics of trading volume under which Uniswap would become as good as Binance in terms of transaction costs and price efficiency. In particular, we focus on scenarios for which the expected increase in trading volume on DEXs is from 3 to 30 times with respect to the trading volume recorded in our sample. Following our model, a given increase ΔV in volume gives rise to a decrease in fees f, an increase of liquidity x^*, or a combination thereof. Moreover, we include three possible values for the dollar value of gas fees g, which is exogenous to the other parameters because it is determined by technological evolution.

E. Results: Transaction Costs

Table VIII presents predicted levels of transaction costs for each combination of parameters. More precisely, hypothetical transaction costs for a 10,000$ transaction executed through Uniswap v2 are reported and expressed in bps. Similarly, each panel of Figure 8 presents the TCs resulting from selected parameter combinations. The current situation is represented by the first row of Panel A, with fees equal to 30 bps and unitary gas and liquidity multipliers. Panel B, which assumes a reduction of gas fees by a factor of 500, shows that transaction costs are roughly halved for most of the pairs. This assumption is reasonable in the context of Ethereum 2.0 and related upgrades.[24] These results highlight that the present high level of gas fees represents a significant friction

[24]The introduction of Ethereum 2.0, with ZK Rollups and data sharding implemented, is expected to allow for around 10 million transactions per second, while the current Ethereum network only supports around 20. In equilibrium, therefore, the gas price is expected to deflate by a factor of 500,000. However, it is fair to expect that the number of active wallets and transactions in the network would also grow significantly at that point, thus positively impacting the gas price. Assuming – as an upper bound – a 1,000-fold increase in the network activity, we thus get to an effective reduction in the gas price by a factor of 500.

for DEX efficiency. As shown in the top-right panel of Figure 8, such a reduction in transaction costs would lead to a situation in which Uniswap is strictly dominated by Binance, but only partially dominated by Kraken. Our second parameter set is depicted in the bottom-left panel of Figure 8 and assumes, in addition to lower gas prices, a six-fold increase in volume leading to a six-fold reduction in exchange fees to 5 bps – a viable assumption in the context of Uniswap v3.[25] Under these assumptions, our equilibrium model predicts that Uniswap would be highly competitive with CEXs, Offering significantly lower transaction costs compared to Binance and Kraken for the majority of pairs. The bottom-right panel of Figure 8 assumes a more sizeable 30-fold increase in trading volume, resulting in a three-fold reduction in exchange fees to 10 bps and a 10-fold increase in pool liquidity. The results show that Uniswap would offer roughly the same transaction costs as Binance for most pairs. This corroborates one of our previous main findings: the most important friction undermining DEXs arises from high levels of exchange fees rather than low levels of liquidity.

All in all, given an increase in trading volume, it would thus be preferable to reduce exchange fees versus attracting more liquidity LPs. Our conclusion applies to the trade size we consider in this analysis (10,000$), but it would likely differ for larger transactions. For those, an increase in available liquidity could provide more benefits to traders with respect to a reduction in the exchange fees.

F. Results: Price Efficiency

Moving to price efficiency, the results are reported in Table IX, which presents predictions on the degree of price inefficiency for each parameter combination. More precisely, the table presents the size of arbitrage bounds for each exchange triplet, computed as in (14). These are based on (11) for the LOB-based Binance and Kraken and on (13)

[25]We chose 5 bps since it is one of the possible values that market participants can select on Uniswap v3. This new version of the exchange, introduced in late 2021, allows LPs to choose among four distinct liquidity pools for the same exchange pair, featuring 1, 5, 30, and 100 bps. This new feature effectively allows LPs to exploit the trade-off highlighted in our model, lowering the fees required to attract or respond to an increase in trading volume.

for the AMM-based Uniswap. Similarly, each panel of Figure 9 presents the size of arbitrage bounds resulting from selected parameter combinations. The present situation is represented by the first row, Panel A of Table IX, with fees equal to 30 bps and trivial gas and liquidity multipliers. Panel B, which assumes a reduction in gas fees by a factor of 500, shows that price efficiency increases significantly for all triplets. Lower gas fees for USDC-USDT-ETH and USDC-BTC-ETH result in increased efficiency by 20% and 50%, respectively, while the reductions for the other triplets are greater than 80%. These heterogeneous effects of gas fees depend on the diverse size of optimal triangular trades. Since the first two triplets enjoy higher liquidity, the optimal trade size is larger, thus reducing the impact of gas fees, which are a fixed cost, on the profitability of potential triangular arbitrages.

Overall, these results highlight that high levels of gas fees represent a critical friction for price efficiency on DEXs and that this effect is more important for triplets involving low-liquidity pairs. As shown in the top-right panel of Figure 9, such an improvement in price efficiency would lead to a situation in which Uniswap, even though significantly better off relative to the present, is still dominated by CEXs. The second parameter combination is depicted in the bottom-left panel of Figure 9 and assumes, on top of low gas fees, a six-fold increase in volume leading to a six-fold reduction in exchange fees to 5 bps. The plot shows how, under these assumptions, Uniswap would offer more efficient prices with respect to Kraken but would still be dominated by Binance for most of the triplets. The bottom-right panel of Figure 9 assumes a more sizeable 30-fold increase in trading volume, resulting in a three-fold reduction in exchange fees to 10 bps and a 10-fold increase in pool liquidity. Under these assumptions, benefits for Uniswap price efficiency are larger for less liquid triplets, while they are reduced for the most liquid ones. This suggests that price efficiency on DEXs may be limited by both high exchange fees and by low liquidity, depending on which friction is more pronounced for the pairs composing the triplets.

VII. Additional Results

We performed numerous additional and robustness analyses. For brevity, we present here only three of them. The first two analyses are motivated by the importance of gas fees on which we shed light on the previous part of our paper. Specifically, we examine the impact of gas costs on liquidity stability and transaction cost predictability. The third analysis focuses on the dynamic connection between CEX and DEX prices. Although the focus of our work has been on market liquidity and price efficiency *within* the DEX and CEX markets, investigating the joint evolution of cryptocurrency prices is the first step in better understanding the dynamics *across* the two market venues.

A. *Transaction Cost Determinants*

In equation (6), we showed that the DEX transaction cost is the sum of three components: the quoted half-spread, the percentage exchange fee, and the gas fee as a fraction of the trade size. In addition, we found that the dollar value of gas prices exhibits strong time series variation. The natural question is how the bid-ask spread and gas costs impact transaction costs, especially while the exchange fee is exogenously fixed. This question is particularly relevant for market participants who want to assess the stability of market liquidity and predict transaction costs. To shed more light on this issue, we analyze the degree of predictability of effective transaction costs composed of bid-ask spreads and exchange fees (excluding D/W fees) on Binance and Uniswap, focusing on the marginal impact of gas prices. We build two simple proxies for the stability of transaction costs offered by cryptocurrency exchanges based on the performance of a linear forecasting model. These measures can also be interpreted as the degree of predictability in transaction costs; that is, the ability of market participants to forecast the effective cost of trading, conditional on the current level of transaction costs. First, we compute the one-lag auto-correlation coefficient ρ of transaction costs associated with a 10,000 \$ trade for every exchange-pair couple. Second, we run the following time series regression at the hourly frequency for each exchange and each pair:

$$\text{TC}(t) = \alpha + \beta\, \text{TC}(t-1) + \varepsilon(t)\ .$$

For Uniswap, we repeat the exercise and impose a time-invariant gas price equal to the sample average. Figure 10 plots the cross-pair average ρ and average R^2 for each exchange, including the synthetic Uniswap with a fixed gas price (purple bar) and a gas price that changes over time (pink bar). These results clearly show that Uniswap and Binance enjoy a very similar level of stability in transaction costs, while Kraken is less auto-correlated. More importantly, once the time series variation due to fluctuating gas prices is removed, we observe almost perfect predictability of transaction costs on Uniswap. Three considerations emerge from this simple analysis: First, transaction costs are largely, but not completely, predictable on both CEXs and DEXs. Second, their degree of prediction is very similar; failure to consider gas costs leads to the erroneous conclusion that transaction costs and liquidity, in general, are fully predictable and more stable on DEXs than on CEXs. Third, this result confirms the importance of considering gas costs, especially when studying DEX market quality. It also generalizes the findings of Lehar and Parlour (2021), who showed that the liquidity provision on Uniswap is extremely stable but did not account for gas costs.

B. Gas Fees and CEX Trading Cost

We now turn to the analysis of whether gas fees can also affect transaction costs in CEX markets. This is an interesting question since gas levels may capture the aggregate demand for immediacy in the overall cryptocurrency market, which can occur for various reasons, including (i) new fundamental information, which causes traders to adjust their asset positions (Glosten and Milgrom, 1985), and (ii) momentum, which manifests itself with autocorrelated and unbalanced order flows (Foucault et al., 2013a). Regardless of whether it is for informational motives and temporary price movements, bid/ask spreads widen in times when market participants rush into trading and have incentives to pay higher gas prices to gain priority of execution. Empirically, therefore, gas prices can exhibit positive time series correlations not only with DEX transaction costs but also with CEX spreads. To test our hypothesis, we run an hourly frequency regression of the spreads on Binance and Kraken on hourly levels of gas prices. To avoid an obvious source of mechanical correlation, we do not consider the DW-fees component from CEX transaction costs, which depends on gas prices. Rather, we use as a dependent variable the first component of equation (5), that is, B/A spreads. Table VI shows that spreads in both Binance and Kraken are indeed positively related to Ethereum gas prices, thus providing suggestive evidence that the latter serves as a proxy for the overall demand for immediacy. Future research can thoroughly study the mechanisms that determine gas costs and the way that they relate to the supply of and demand for liquidity.

C. Asymmetric Prices

As a final test, we analyze the dynamic connection between the prices of the same cryptocurrency pair traded simultaneously in DEX and CEX markets. Since the mid-price represents the fundamental value of a security, the null hypothesis is that the same cryptocurrency traded in two different markets should have the same mid-price. The alternative hypothesis is that CEX mid-price is slightly higher because informed trading coupled with the convexity of DEX price impact curve stemming from the constant product rule gives rise to asymmetric prices (Aoyagi and Ito, 2021). More specifically, this is due to informed trading that enlarges the CEX spread, encouraging liquidity traders to use more DEX markets. In addition, the convex curve implies a larger price impact for buy orders rather than for sell orders (ceteris paribus), rendering DEX more attractive for informed *sell* orders.

A simple way to test these hypotheses is to compute the ratio between the simultaneous DEX and CEX mid-prices minus one and regress it on constant. The rejection of the null hypothesis in favor of a significantly negative coefficient would support the price asymmetry hypothesis. Furthermore, the negative bias of the DEX mid-price should be larger in case of negative fundamental value movements that we simply capture with a dummy variable equal to one if the lagged mid-price change is below -5%.[26]

We construct mid-price ratios for Uniswap/Binance and Uniswap/Kraken, using hourly prices for the five pairs included in our main analysis, expressed in basis points. The regression results reported in Table VII strongly support the price asymmetry hypothesis. In all specifications, the two significantly negative coefficients indicate that (i) the DEX mid-price is systematically lower than the CEX equivalent and (ii) this negative bias doubles when the fundamental value takes a downward direction. The economic magnitude of the average deviation is relatively weak, indicating that mid-prices on Uniswap are on average 3 bps lower than those on the two CEXs. However, conditioning on negative innovations gives rise to larger and economically significant deviations, between 38 and 49 basis points. The results are robust to the inclusion of pair fixed effects.

VIII. Conclusion

To lay the foundation of our study, we describe the institutional details of AMM compared to centralized exchanges based on limit order books, including (i) the AMM price formation process governed by an algorithm that automatically determines transaction prices based on the liquidity made available by market participants, (ii) the various risks that a CEX user has to deal with when she deposits crypto-assets to the (centralized) exchange and (iii) several advantages of DEX, especially in terms of accessibility, security, censorship resistance, and settlement.

At the heart of the work is the analysis of two key aspects of the market quality of cryptocurrency exchanges: (i) market liquidity, examining all the main cost components to trade on CEXs and DEXs, and (ii) price efficiency, analyzing deviations from the tri-

[26]The direction of our results is robust to choosing a different threshold, including -1%, -2%, and -10%, while the magnitude of the estimated coefficient on the interaction term is threshold-dependent.

angular no-arbitrage condition. The picture that arises is that the DEX environment is competitive in terms of transaction costs, mainly because of relatively competitive spreads that offset high gas fees and exchange fees. On the other hand, CEX price dynamics support price efficiency by following *the law of one price* more closely. Also, in this contest, high levels of gas fees required by proof-of-work blockchains and exchange fees constitute the most significant friction harming the market quality of DEXs. This is particularly relevant for triangular price deviations, which require three distinct transactions to be executed.

Next, we develop an equilibrium model of liquidity provision in DEXs to decipher the risk-return trade-off faced by liquidity providers (LPs), based on Impermanent Loss and the expected profits from trading fees. As the former depends only on the relative volatility of the exchange pair while the latter is decreasing in the total liquidity pool size, our theory implies an optimal level of stacked liquidity in equilibrium. We show that such a stylized model explains most of the empirical variation of liquidity levels in the cross-section of exchange pairs and over time.

The insights provided by our theoretical model allow us to link hypothetical levels of trading volume to the implied amount of deposited liquidity and exchange fees required by LPs in equilibrium. This model can also assess the effects on market quality generated by current and future changes in the blockchain protocol. We conclude that in DEXs, liquidity can further improve, and prices can become as efficient as in CEXs under relatively modest increases in volume, provided gas costs decline thanks to the expected technological advances in blockchain technology. Given these insights, we argue that the innovative market structure of decentralized crypto exchanges could be applied to other asset classes in the future and be competitive with existing CEXs in terms of market quality.

References

H. Adams, N. Zinsmeister, and D. Robinson. Uniswap v2 core. 2020.

H. Adams, N. Zinsmeister, M. Salem, R. Keefer, and D. Robinson. Uniswap v3 core. 2021.

A. A. Aigner and G. Dhaliwal. Uniswap: Impermanent loss and risk profile of a liquidity provider. *arXiv preprint arXiv:2106.14404*, 2021.

G. Angeris, H.-T. Kao, R. Chiang, C. Noyes, and T. Chitra. An analysis of uniswap markets. *arXiv preprint arXiv:1911.03380*, 2019.

J. Aoyagi and Y. Ito. Coexisting exchange platforms: Limit order books and automated market makers. 2021.

A. F. Bariviera. The inefficiency of bitcoin revisited: A dynamic approach. *Economics Letters*, 161:1–4, 2017.

B. Biais, C. Bisiere, M. Bouvard, C. Casamatta, and A. J. Menkveld. Equilibrium bitcoin pricing. *Journal of Finance*, Forthcoming.

N. Borri and K. Shakhnov. Cryptomarket discounts. *SSRN 3124394*, 2018.

A. Brauneis and R. Mestel. Price discovery of cryptocurrencies: Bitcoin and beyond. *Economics Letters*, 165:58–61, 2018.

A. Brauneis, R. Mestel, R. Riordan, and E. Theissen. How to measure the liquidity of cryptocurrency markets? *Journal of Banking & Finance*, 124:106041, 2021.

A. Capponi and R. Jia. The adoption of blockchain-based decentralized exchanges. *arXiv preprint arXiv:2103.08842*, 2021.

A. H. Dyhrberg, S. Foley, and J. Svec. How investible is bitcoin? Analyzing the liquidity and transaction costs of bitcoin markets. *Economics Letters*, 171:140–143, 2018.

A. Evans. Liquidity provider returns in geometric mean markets. *arXiv preprint arXiv:2006.08806*, 2020.

A. Evans, G. Angeris, and T. Chitra. Optimal fees for geometric mean market makers. *arXiv preprint arXiv:2104.00446*, 2021.

T. Foucault, M. Pagano, and A. Röell. *Market Liquidity: Theory, Evidence, and Policy.* Oxford University Press, 2013a.

T. Foucault, M. Pagano, A. Roell, and A. Röell. *Market liquidity: theory, evidence, and policy.* Oxford University Press, 2013b.

T. Foucault, R. Kozhan, and W. W. Tham. Toxic arbitrage. *The Review of Financial Studies*, 30(4):1053–1094, 2017.

M. Fukasawa, B. Maire, and M. Wunsch. Weighted variance swaps hedge against impermanent loss. *Available at SSRN 4095029*, 2022.

L. Glosten. Is the electronic open limit order book inevitable? *Journal of Finance*, 49(4): 1127–1161, 1994.

L. R. Glosten and P. R. Milgrom. Bid, ask and transaction prices in a specialist market with heterogeneously informed traders. *Journal of Financial Economics*, 14(1):71–100, 1985.

J. Han, S. Huang, and Z. Zhong. Trust in defi: an empirical study of the decentralized exchange. *Available at SSRN 3896461*, 2021.

N. Hautsch, C. Scheuch, and S. Voigt. Building trust takes time: Limits to arbitrage in markets with stochastic settlement latency. *SSRN 3302159*, 2018.

L. Heimbach, Y. Wang, and R. Wattenhofer. Behavior of liquidity providers in decentralized exchanges. *arXiv preprint arXiv:2105.13822*, 2021.

L. Heimbach, E. Schertenleib, and R. Wattenhofer. Risks and returns of uniswap v3 liquidity providers. *arXiv preprint arXiv:2205.08904*, 2022.

A. Khakhar and X. Chen. Delta hedging liquidity positions on automated market makers. *arXiv preprint arXiv:2208.03318*, 2022.

S. Krückeberg and P. Scholz. Decentralized efficiency? Arbitrage in bitcoin markets. *Financial Analysts Journal*, 76(3):135–152, 2020.

A. Lehar and C. Parlour. Decentralized exchanges. *SSRN 3905316*, 2021.

S. Loesch, N. Hindman, M. B. Richardson, and N. Welch. Impermanent loss in uniswap v3. *arXiv preprint arXiv:2111.09192*, 2021.

I. Makarov and A. Schoar. Trading and arbitrage in cryptocurrency markets. *Journal of Financial Economics*, 135(2):293–319, 2020.

H. Mance. Left high and dry by a crypto founder's demise. *Risk.net*, 2019.

B. R. Marshall, N. N. H. Nguyen, and N. Visaltanachoti. Bitcoin liquidity. *SSRN 3194869*, 2019.

S. Nadarajah and J. Chu. On the inefficiency of bitcoin. *Economics Letters*, 150:6–9, 2017.

P. O'Neill. Can markets be fully automated? evidence from an automated market maker. 2022.

N. ORX. Op risk data: Defi-ant crypto hackersteals $610m. *Risk.net*, 2021.

E. Pagnotta and A. Buraschi. An equilibrium valuation of bitcoin and decentralized network assets. *SSRN 3142022*, 2018.

A. Park. The conceptual flaws of constant product automated market making. *Available at SSRN 3805750*, 2021.

J. Ryder and N. ORX. Top 10 op risk losses for 2021 hog $15bn total. *Risk.net*, 2022.

A. Urquhart. The inefficiency of bitcoin. *Economics Letters*, 148:80–82, 2016.

Appendix

A. Triangular Arbitrage Bounds on AMMs

In this section we derive an algebraic expression for the cost of executing a triangular trade $X \to Y \to Z \to X$. We then define the triangular arbitrage bounds as the degree of absolute-value deviation from the law of one price such that the return from the triangular trade equals the cost of executing it. Let x and y be the number of tokens X and Y be the liquidity pools for the pair X/Y, f the exchange fees, and $\varphi = 1 - f$. We start by calculating the price impact factor for the first trade $X \to Y$. To derive an expression for it, recall that quoted price, the transaction price, and the resulting spread are given by

$$P_{XY} = \frac{y}{x}, \qquad T_{XY} = \frac{y}{x + \varphi\Delta x}, \qquad S_{XY} = \frac{\varphi\Delta x}{x + \varphi\Delta x}.$$

We define the price impact factor ρ_X as follows:

$$\rho_X = 1 - S_{XY} = \frac{x}{x + \varphi\Delta x}$$

As the transaction price can be expressed as the product of the price impact factor and the quoted price, we rewrite the output amount from the first trade as follows:

$$\Delta y = \varphi T_{XY}\Delta x = \varphi\rho_X P_{XY}\Delta x \tag{17}$$

Now let y' be the number of tokens Y in the liquidity pool for the pair Y/Z. We can express the price impact factor ρ_Y for the second trade $Y \to Z$ as a function of Δx using (17):

$$\rho_Y = \frac{y'}{y' + \varphi\Delta y} = \frac{y'}{y' + \varphi^2\rho_X P_{XY}\Delta x}.$$

Similarly, letting z be the number of tokens Z in the liquidity pool for the pair Z/X, we can write the price impact factor ρ_Z as

$$\rho_Z = \frac{z}{z + \varphi\Delta z} = \frac{z}{z + \varphi^3\rho_X\rho_Y P_{XY}P_{YZ}\Delta x}.$$

The cumulative price-impact factor across the three trades is given by $\rho = \rho_X\rho_Y\rho_Z$, while the total spread equals $S_{XYZ} = 1 - \rho$. Next, we define the total exchange fees charged as

$$F_{XYZ}(\Delta x) = f(1 + \varphi\rho_X + \varphi^2\rho_X\rho_Y).$$

Notice that the trade size decreases for the second and third trades due to the price impact and previous exchange fees charged. As the last cost component, we consider the gas cost of executing the three trades within a single smart contract function. In order to estimate the quantity of gas required for such operation, we deployed a simple smart contract capable of executing a triangular trade on the Ethereum MainNet; we found that the cost is approximately $g_3 = 241,822$ gas units. Notice that the gas cost of executing the triangular trade through the smart contract is roughly 30% lower than the cost $3g$ of executing the three transactions separately. We finally represent the total cost of a triangular trade $K(\Delta x)$ as follows:

$$\begin{aligned} K(\Delta x) &= S_{XYZ}(\Delta x) + F_{XYZ}(\Delta x) + g_3/\Delta x) \\ &= 1 - \rho + f(1 + \varphi\rho_X + \varphi^2\rho_X\rho_Y) + g_3/\Delta x\,. \end{aligned}$$

Based on the no-arbitrage assumption, the deviation from the law of one price $\theta = P_{XY}\,P_{YZ}\,P_{ZX} - 1$ should be smaller than the cost $K(dx)$ of executing the triangular trade. We thus define the triangular arbitrage bounds as

$$\begin{aligned} \theta^H, \theta^L &= \pm K(dx) \\ &= \pm\Big(\rho - 1 - f(1 + \varphi\rho_X + \varphi^2\rho_X\rho_Y) - g_3/\Delta x\Big). \end{aligned}$$

Tables and Figures

Table I. Summary Statistics. Panel A reports statistics based on the liquidity pools underlying the 1,000 most liquid exchange pairs in Uniswap v2. We report the total value of liquidity in USD, while we report the number of swap transactions and the time since the pool was initiated in days. Panel B reports statistics for gas prices on the Ethereum blockchain over the January, 2021-December, 2021 period at the hourly frequency. Panel C reports statistics of daily impermanent loss for 100 pairs traded on Uniswap v2 between April 2020 and April 2021, aggregated at different levels and expressed in basis points. The first row refers to the distribution of the realized IL in a pair-day panel, while the second and third rows report the IL averaged by pair and by day, respectively. The pairs are selected by ranking the available pairs in Uniswap v2 by the USD value of liquidity as of September 2021.

	N	Mean	Std	1%	10%	50%	90%	99%
Panel A: Liquidity Pools								
Liquidity (Million USD)	1,000	6.71	41.52	0.01	0.03	0.99	8.32	114.54
Transactions (thousands)	1,000	37.54	121.69	7.24	8.28	18.13	62.77	238.35
Age (days)	1,000	193.05	96.03	2.33	47.04	204.09	333.69	347.26
Panel B: Gas Prices								
Gas Price (GWEI)	8,737	106.72	77.76	14.97	27.05	94.96	191.33	366.91
Cost of a Swap (USD)	8,737	32.68	27.58	3.54	6.60	26.50	65.48	128.96
Panel C: Impermanent Loss								
Pair-Day	18,178	16.16	131.75	0.00	0.02	1.32	23.49	49.61
Pair	100	20.06	21.89	0.21	2.08	13.10	54.44	75.08
Day	356	20.21	37.09	3.26	4.21	10.63	37.67	69.32

Table II. Trading Volume by Pair. The table reports the daily average trading volume in millions of USD over the January, 2021–December, 2021 period for the five trading pairs on which we focus our main empirical analysis. To select these pairs, we start by considering the intersection of trading pairs available in Uniswap, Binance, and Kraken. We then compute percentile rankings of average daily volumes in USD for each of the three exchanges, take their average, and select the first five pairs ranked by the resulting metric. The percentage of the aggregate volume represented by these pairs on each exchange is reported below.

Pair	**Uniswap**	**Binance**	**Kraken**
ETH-USDT	2223.72	2117.46	1267.83
ETH-BTC	538.05	490.38	260.16
USDC-USDT	156.27	153.74	95.14
ETH-USDC	60.39	57.80	35.21
LINK-ETH	9.65	8.90	5.12
Fraction of Total Volume	24.25%	15.34%	17.65%

Table III. Transaction Costs – Breakdown by Trade Size. The table displays average transaction costs for each exchange pair in our sample, for different trade sizes denominated in USD, for Uniswap v2, Binance, and Kraken. The figures are based on equation (6) for DEX and equation (5) for CEXs and it is expressed in basis points. They are first computed at the hourly frequency and then averaged from January 2021 to December 2021. For each pair, we present a breakdown of total transaction costs by their individual components: exchange fees, B/A spread, and gas fees or deposit-withdrawal fees.

Pair		ETH-USDC	ETH-USDT	ETH-BTC	LINK-ETH	USDC-USDT
Panel A: Uniswap						
Exchange Fees		30.000	30.000	30.000	30.000	30.000
B/A Spread	1,000	0.077	0.092	0.103	0.303	0.346
	10,000	0.774	0.925	1.031	3.029	3.465
	100,000	7.739	9.248	10.313	30.251	34.593
	1,000,000	77.114	92.089	102.597	298.139	340.027
Gas Fees	1,000	267.174	267.174	267.174	267.147	267.174
	10,000	26.717	26.717	26.717	26.714	26.717
	100,000	2.671	2.671	2.671	2.671	2.671
	1,000,000	0.267	0.267	0.267	0.267	0.267
Total TCs	1,000	297.251	297.266	297.277	297.450	297.520
	10,000	57.491	57.642	57.749	59.744	60.182
	100,000	40.411	41.920	42.985	62.923	67.265
	1,000,000	107.382	122.356	132.864	328.406	370.294
Panel B: Binance						
Exchange Fees		10.000	10.000	10.000	10.000	10.000
B/A Spread	1,000	3.262	0.318	0.448	7.730	0.504
	10,000	4.015	0.591	0.744	10.505	0.512
	100,000	10.447	2.049	2.911	55.704	0.602
	1,000,000	175.798	10.207	26.165	683.170	2.669
DW Fees	1,000	163.580	133.636	158.055	239.278	225.935
	10,000	16.358	13.363	15.805	23.927	22.593
	100,000	1.635	1.336	1.580	2.392	2.259
	1,000,000	0.163	0.133	0.158	0.239	0.225
Total TCs	1,000	176.842	143.955	168.504	257.008	236.440
	10,000	30.373	23.954	26.550	44.433	33.105
	100,000	22.083	13.385	14.492	68.097	12.861
	1,000,000	185.962	20.341	36.323	693.409	12.895
Panel C: Kraken						
Exchange Fees		26.000	26.000	26.000	26.000	26.000
B/A Spread	1,000	5.466	3.224	1.711	10.548	1.030
	10,000	6.491	4.106	2.125	15.521	1.665
	100,000	27.876	10.689	4.051	52.606	11.167
	1,000,000	460.590	150.860	18.882	691.818	380.783
DW Fees	1,000	195.561	245.553	147.699	161.717	305.945
	10,000	19.556	24.555	14.769	16.171	30.594
	100,000	1.955	2.455	1.476	1.617	3.059
	1,000,000	0.195	0.245	0.147	0.161	0.305
Total TCs	1,000	227.028	274.777	175.410	198.265	332.975
	10,000	52.047	54.661	42.895	57.693	58.259
	100,000	55.832	39.144	31.528	80.223	40.227
	1,000,000	486.785	177.105	45.030	717.980	407.089

Table IV. Transaction Costs – Breakdown by Component. The table displays the time-series distribution of transaction costs for each exchange pair in our sample, for Uniswap v2, Binance, and Kraken, expressed in basis points. The figures are based on equation (6) for DEX and equation (5) for CEXs;, the figures are computed at the hourly frequency from January 2021 to December 2021, assuming a transaction size of 10,000\$. For each pair, we present a breakdown of total transaction costs by their individual components: exchange fees, B/A spread, and gas fees or deposit-withdrawal fees.

Pair		ETH-USDC	ETH-USDT	ETH-BTC	LINK-ETH	USDC-USDT
Panel A: Uniswap						
Exchange Fees		30.000	30.000	30.000	30.000	30.000
B/A Spread	Mean	0.774	0.925	1.032	3.030	3.466
	Std	0.159	0.207	0.376	0.970	1.326
	5%	0.553	0.600	0.526	1.648	1.505
	95%	1.019	1.282	1.595	4.942	5.932
Gas Fees	Mean	26.717	26.717	26.717	26.715	26.717
	Std	24.554	24.554	24.554	24.554	24.554
	5%	4.450	4.450	4.450	4.450	4.450
	95%	63.014	63.014	63.014	63.014	63.014
Total TC	Mean	57.492	57.643	57.749	59.744	60.183
	Std	24.554	24.554	24.455	24.197	24.481
	5%	35.312	35.310	35.939	38.846	37.657
	95%	93.915	93.978	93.996	95.456	95.760
Panel B: Binance						
Exchange Fees		10.000	10.000	10.000	10.000	10.000
B/A Spread	Mean	4.016	0.591	0.745	10.506	0.512
	Std	1.833	0.436	0.503	6.228	0.082
	5%	2.382	0.206	0.204	5.157	0.500
	95%	6.860	1.316	1.587	23.590	0.541
DW Fees	Mean	16.358	13.364	15.806	23.928	22.594
	Std	9.704	9.717	10.147	11.749	14.568
	5%	7.966	4.963	6.344	10.929	9.614
	95%	30.588	27.598	30.441	41.501	43.814
Total TC	Mean	30.374	23.955	26.550	44.434	33.106
	Std	10.118	9.756	10.176	12.740	14.577
	5%	21.136	15.427	16.979	28.533	20.117
	95%	44.761	38.181	41.489	63.139	54.458
Panel C: Kraken						
Exchange Fees		26.000	26.000	26.000	26.000	26.000
B/A Spread	Mean	6.491	4.106	2.125	15.522	1.665
	Std	8.986	4.476	3.039	12.238	1.858
	5%	2.161	1.247	1.132	9.907	0.565
	95%	11.996	8.100	3.432	28.916	4.384
DW Fees	Mean	19.556	24.555	14.770	16.172	30.595
	Std	10.237	10.234	10.261	10.468	14.564
	5%	10.306	15.306	5.478	6.373	17.615
	95%	34.975	39.916	30.287	31.897	51.799
Total TC	Mean	52.047	54.661	42.895	57.693	58.260
	Std	13.895	11.302	10.838	16.613	14.798
	5%	40.705	43.704	33.264	43.645	44.521
	95%	67.702	70.218	58.150	77.483	79.773

Table V. Model Fit. The table reports results from a panel regression of log-observed liquidity levels onto log liquidity levels predicted by the equilibrium model outlined in Section VI and computed as in equation (15). Both the dependent and independent variables are computed at the pair-day level for 100 exchange pairs from April, 2020 to April, 2021. We saturate the regression model with day- and pair fixed effects. T-stats are reported in parentheses, based on robust standard errors double-clustered at the pair- and day-level. Asterisks denote significance levels (***= 1%, **= 5%, *= 10%).

	(1)	(2)	(3)	(4)
Dependent Variable	Log(Liquidity)	Log(Liquidity)	Log(Liquidity)	Log(Liquidity)
Log(Predicted Liquidity)	0.89***	0.61***	0.89***	0.53***
	(29.94)	(12.39)	(30.14)	(10.24)
Constant	5.92***			
	(21.71)			
Observations	42,293	42,293	42,293	42,293
R-squared	0.92	0.97	0.92	0.98
Pair Fixed Effects	-	Yes	-	Yes
Date Fixed Effects	-	-	Yes	Yes

Table VI. CEX Transaction Costs and Gas Prices. The table reports results from the hourly time series regression of transaction costs on Binance and Kraken (averaged across pairs) on their lagged values and on the average price of gas from Ethereum blocks validated during the same hour. Transaction costs are computed based on equation 5, except that we remove deposit fees so as to avoid mechanical correlation with gas prices. We saturate the regression model with day fixed effects. T-stats are reported in parentheses and are based on robust standard errors. Asterisks denote significance levels (***= 1%, **= 5%, *= 10%).

	(1)	(2)	(3)	(4)
Dependent Variable	Binance Spread	Binance Spread	Kraken Spread	Kraken Spread
Constant	0.263***	1.419***	2.321***	3.383***
	(5.693)	(14.484)	(6.118)	(5.243)
Gas Price	0.002***	0.005***	0.015***	0.020***
	(3.687)	(4.488)	(4.227)	(3.581)
Binance Spread (t-1)	0.864***	0.388***		
	(38.330)	(10.341)		
Kraken Spread (t-1)			0.359***	0.106
			(5.073)	(1.688)
Observations	6,802	6,802	6,808	6,808
R-squared	0.803	0.861	0.229	0.382
Date Fixed Effects	-	Yes	-	Yes

Table VII. Asymmetric Mid-Prices. The table reports results from time series regressions of the ratio between simultaneous DEX and CEX mid-prices on a constant. We construct the ratios at the hourly frequency as $P_{\text{DEX}}/P_{\text{CEX}} - 1$, where P_{DEX} is the quoted price of Uniswap v2, while P_{CEX} is the mid-price of Binance on specifications (1)-(3) and the mid-price of Kraken on specifications (4)-(6). On some specifications, we include a dummy variable as a regressor, indicating whether the price decreased by at least 5 percentage points in the previous hour. Specifications (3) and (6) include pair fixed-effects in order to control for potential heterogeneity across pairs. T-stats are reported in parentheses, based on robust standard errors clustered by pairs. Asterisks denote significance levels (***= 1%, **= 5%, *= 10%).

	(1)	(2)	(3)	(4)	(5)	(6)
Dependent Variable		Price Ratio **Uniswap** / **Binance**			Price Ratio **Uniswap** / **Kraken**	
Constant	-3.327*** (-3.415)	-3.266*** (-3.417)	-3.267*** (-124.350)	-3.393*** (-3.4585)	-3.318*** (-3.4110)	-3.318*** (-254.23)
$\Delta P_{t-1} < 0.05$		-37.922** (-2.349)	-37.553** (-2.294)		-48.622*** (-5.584)	-48.351*** (-5.727)
Observations	28,666	28,666	28,666	29,754	29,754	29,754
Pair FEs	-	-	Yes	-	-	Yes

Table VIII. Uniswap Hypothetical Transaction Costs. The table displays hypothetical transaction costs for a 10,000$ transaction executed through Uniswap, expressed in basis points. They are computed as in equation (6), first at the hourly frequency and then averaged from January, 2021 to December, 2021 for different gas fee reduction factors (*Gas*), platform fees (*Fees*, in basis points), and liquidity multipliers (*Liq*). Each row represents a potential future scenario requiring an increase in trading volume equal to ΔV, according to our model presented in Section VI.

Fees	Liq	Gas	ΔV	BTC ETH	ETH USDT	LINK ETH	USDC ETH	USDC USDT
Panel A: Current gas prices								
30	1	1	1	57.75	57.64	59.74	57.49	60.18
10	1	1	3	37.75	37.64	39.75	37.49	40.19
5	1	1	6	32.75	32.65	34.75	32.49	35.19
30	10	1	10	56.82	56.81	57.02	56.79	57.06
10	10	1	30	36.82	36.81	37.02	36.80	37.06
Panel B: Lower gas prices								
30	1	500	1	31.09	30.98	33.08	30.83	33.52
10	1	500	3	11.09	10.98	13.09	10.83	13.53
5	1	500	6	6.09	5.98	8.09	5.83	8.53
30	10	500	10	30.16	30.15	30.36	30.13	30.40
10	10	500	30	10.16	10.15	10.36	10.13	10.40

Table IX. Uniswap Hypothetical Price Inefficiency. The table displays hypothetical levels of price inefficiency on the Uniswap exchange, expressed in basis points. They are estimated as in equation (13), first at the hourly frequency and then averaged from January, 2021 to December, 2021 for different gas fee reduction factors (*Gas*), platform fees (*Fees*, in basis points), and liquidity multipliers (*Liq*). Each row represents a potential future scenario requiring an increase in trading volume equal to ΔV, according to our model presented in Section VI.

Fees	Liq	Gas	ΔV	USDC USDT ETH	USDC BTC ETH	USDT BTC ETH	LINK USDT ETH	USDC USDT BTC
Panel A: Current gas prices								
30	1	1	1	122.03	189.01	959.83	785.41	967.91
10	1	1	3	62.38	129.55	902.72	727.76	910.70
5	1	1	6	47.43	114.65	888.40	713.31	896.36
30	10	1	10	100.54	121.19	371.67	313.02	374.25
10	10	1	30	40.96	61.53	312.76	253.93	315.34
Panel B: Lower gas prices								
30	1	500	1	91.18	94.18	129.98	121.49	130.35
10	1	500	3	31.42	34.44	70.34	61.83	70.71
5	1	500	6	16.44	19.46	55.40	46.87	55.76
30	10	500	10	90.22	91.14	102.47	99.78	102.59
10	10	500	30	30.45	31.38	42.75	40.05	42.87

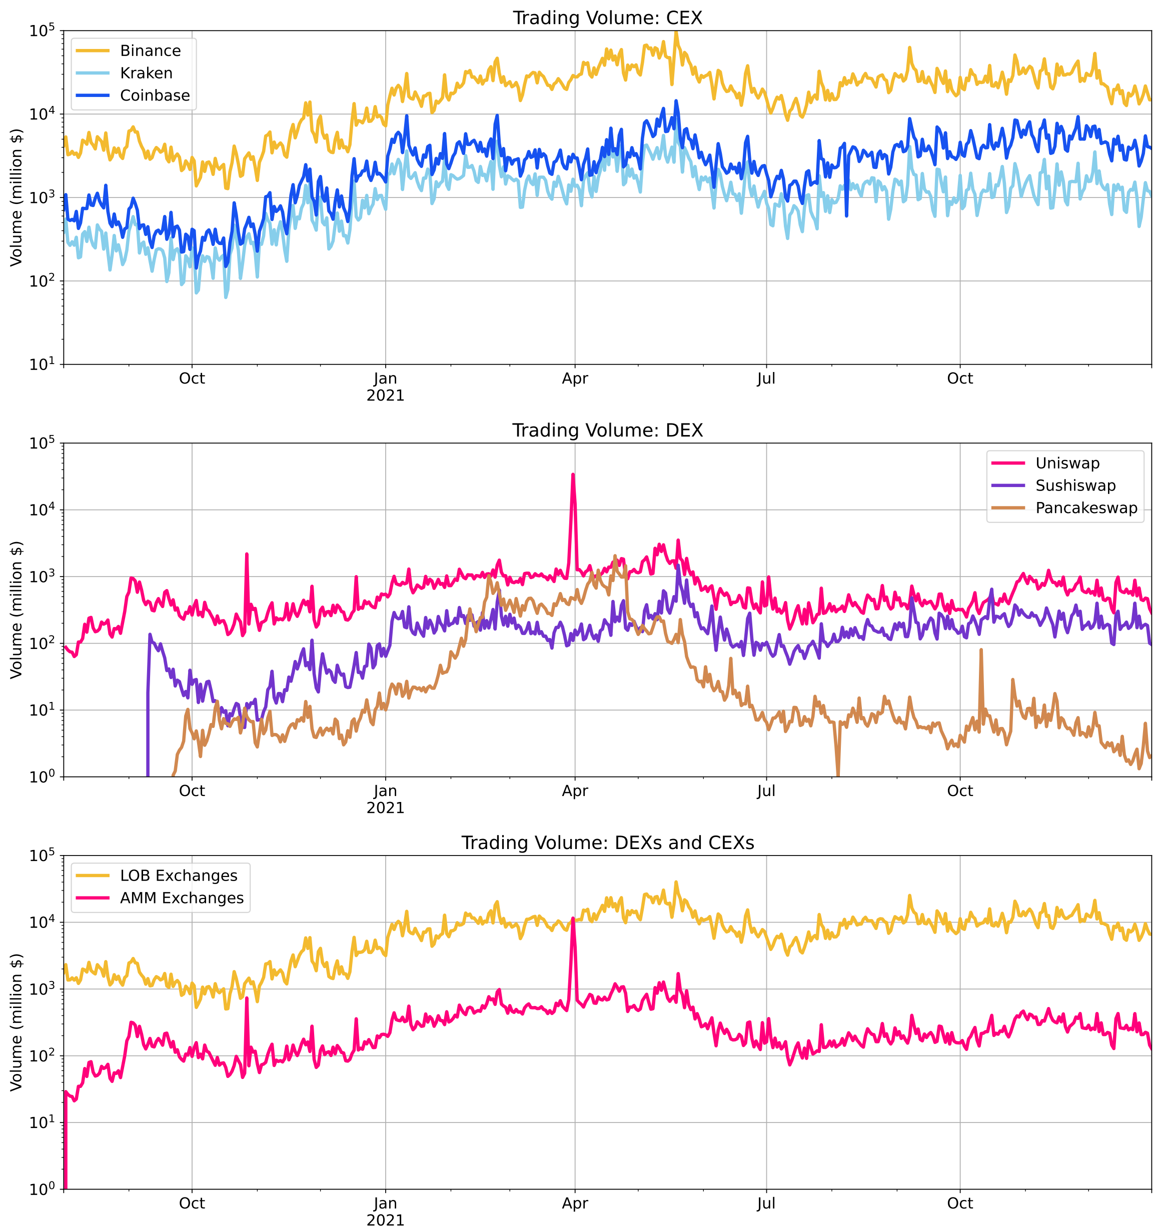

Figure 1. Trading Volume. The figure presents traded volumes for the CEXs and DEXs in our sample, namely Binance, Kraken, Coinbase, Uniswap v2, SushiSwap, and PancakeSwap from August, 2020 to December, 2021. The first panel displays the volume traded in CEXs, as the sum of the volumes for all trading pairs listed on each exchange. Similarly, the second panel displays the aggregate volume traded in DEXs. The last panel presents the volume traded in both CEXs and DEXs, averaged across the three exchanges in each category. In all panels, the vertical axes use the log scale and are expressed in million USD.

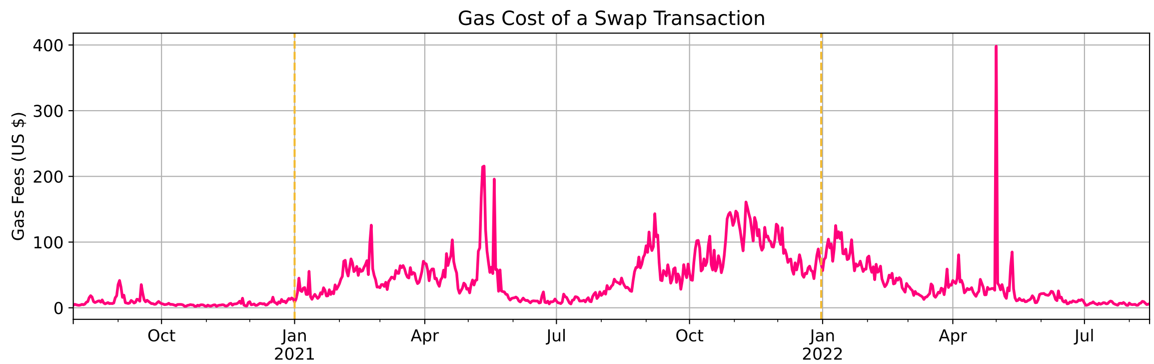

Figure 2. Gas Fees. The figure presents the time series evolution of the gas costs of a swap transaction in our sample in USD. This is computed at the hourly frequency, multiplying the units of gas required to execute a swap, numbering roughly 110,000, by the average gas price associated with transactions in the blocks validated during each hour in USD. Since the number of gas units is constant over time, the time series variation comes from oscillating gas prices in ETH and the fluctuation of the USD/ETH exchange rate. The two yellow vertical lines indicate the period on which our main market quality analysis is conducted, that is, from January, 2020 to December, 2020.

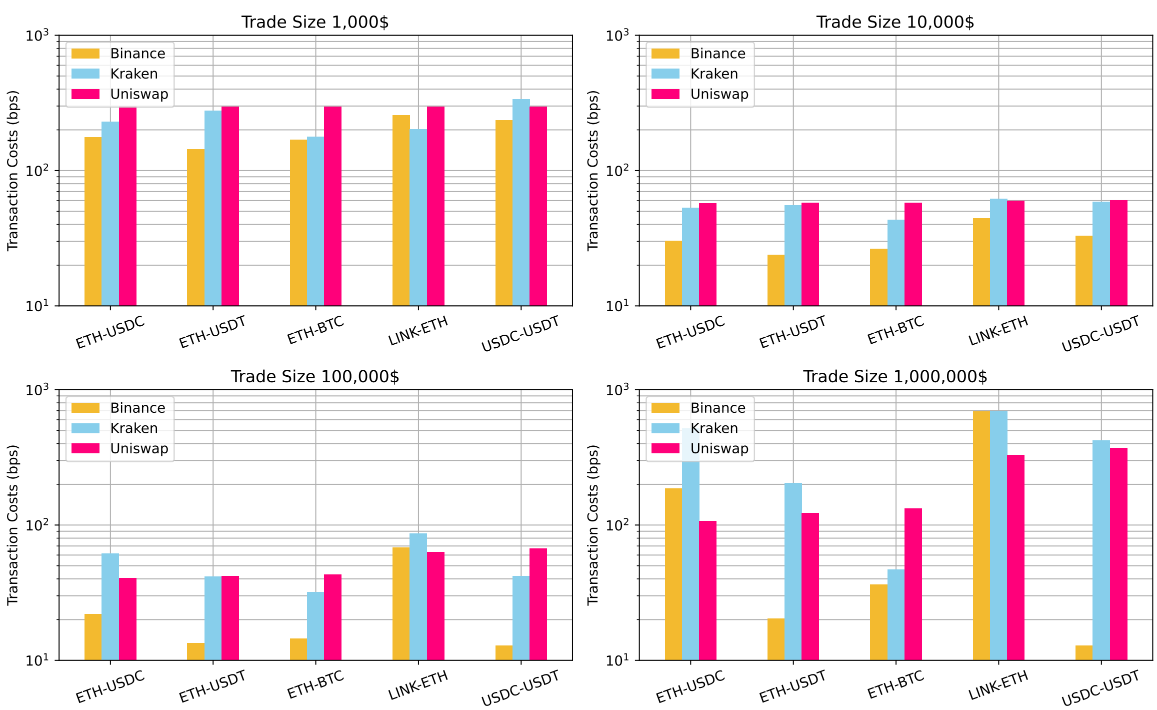

Figure 3. Transaction Costs. The figure presents transaction costs, computed as in equation (5) for the LOB-based Binance and Kraken and on equation (6) for the AMM-based Uniswap v2. These are computed at the hourly frequency for the five pairs in our sample and for different trade sizes ($10^3, 10^4, 10^5$, and 10^6 US dollars), then averaged from January, 2021 to December, 2021. As discussed in Section IV, the displayed transaction costs include B/A spreads, exchange fees (30 basis points for Uniswap, 10 basis points for Binance, and 26 basis points for Kraken), and settlement fees (gas fees for Uniswap, and deposit and withdraw fees for the two CEXs). The vertical axis uses log-scale and is reported in basis points.

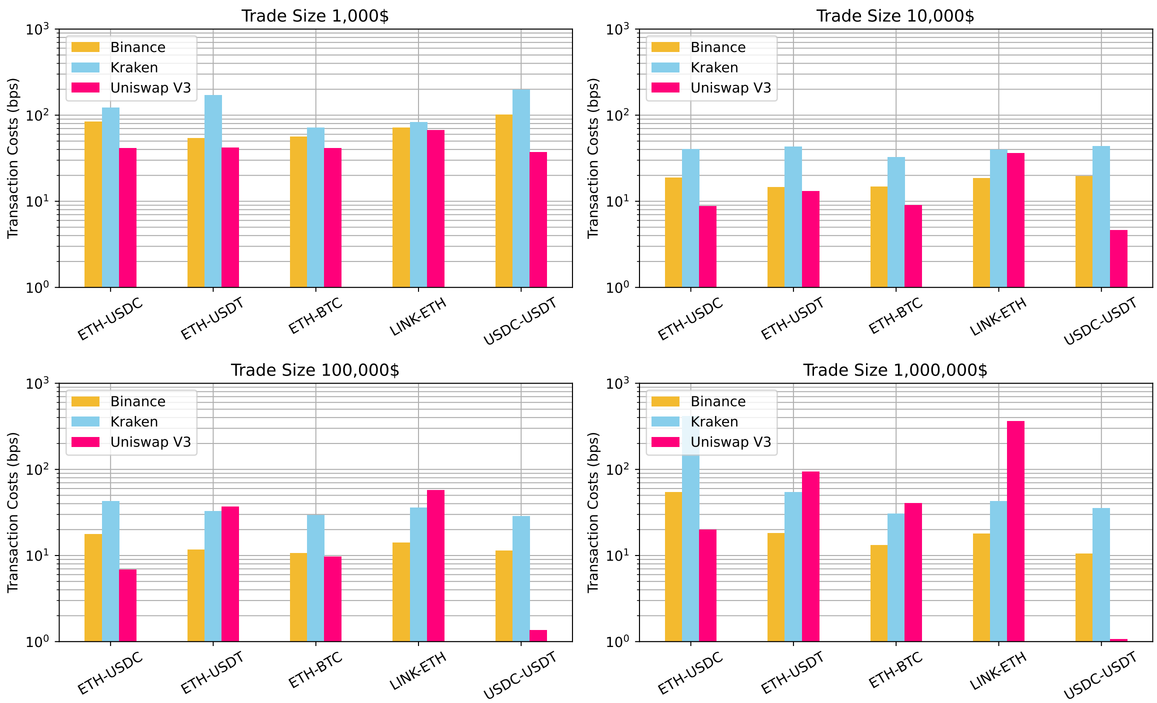

Figure 4. Transaction Costs – Uniswap v3. The figure presents transaction costs for the AMM-based Uniswap gas, and for the LOB-based Binance and Kraken. The figures are obtained in real-time from the official Uniswap v3 SDK and from the official APIs of Binance and Kraken, respectively. These are fetched every 30 minutes for the five pairs in our sample and for different trade sizes ($10^3, 10^4, 10^5$, and 10^6 USD), then averaged from July 22, 2022 to September 11, 2022. As discussed in Section IV, the displayed transaction costs include B/A spreads, exchange fees (30 basis points for Uniswap, 10 basis points for Binance, and 26 basis points for Kraken), and settlement fees (gas fees for Uniswap, and deposit and withdraw fees for the two CEXs). The vertical axis uses log-scale and is reported in basis points.

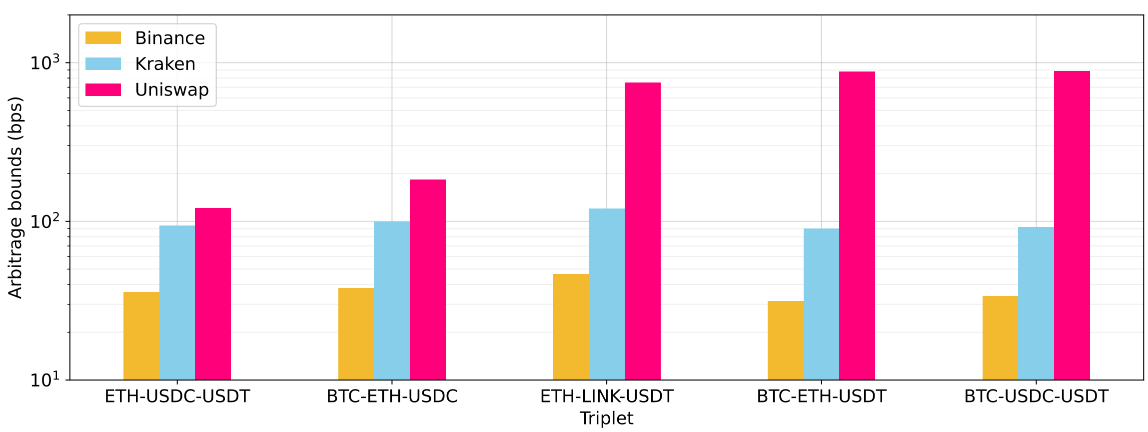

Figure 5. Price Inefficiency. The figure presents price inefficiency levels, proxied by the size of arbitrage bounds computed as in equation (14). These are based on equation (11) for the LOB-based Binance and Kraken and on (13) for the AMM-based Uniswap. They are estimated at the hourly frequency for the five triplets in our sample, then averaged from January, 2021 to December, 2021. The vertical axis uses log-scale and is reported in basis points.

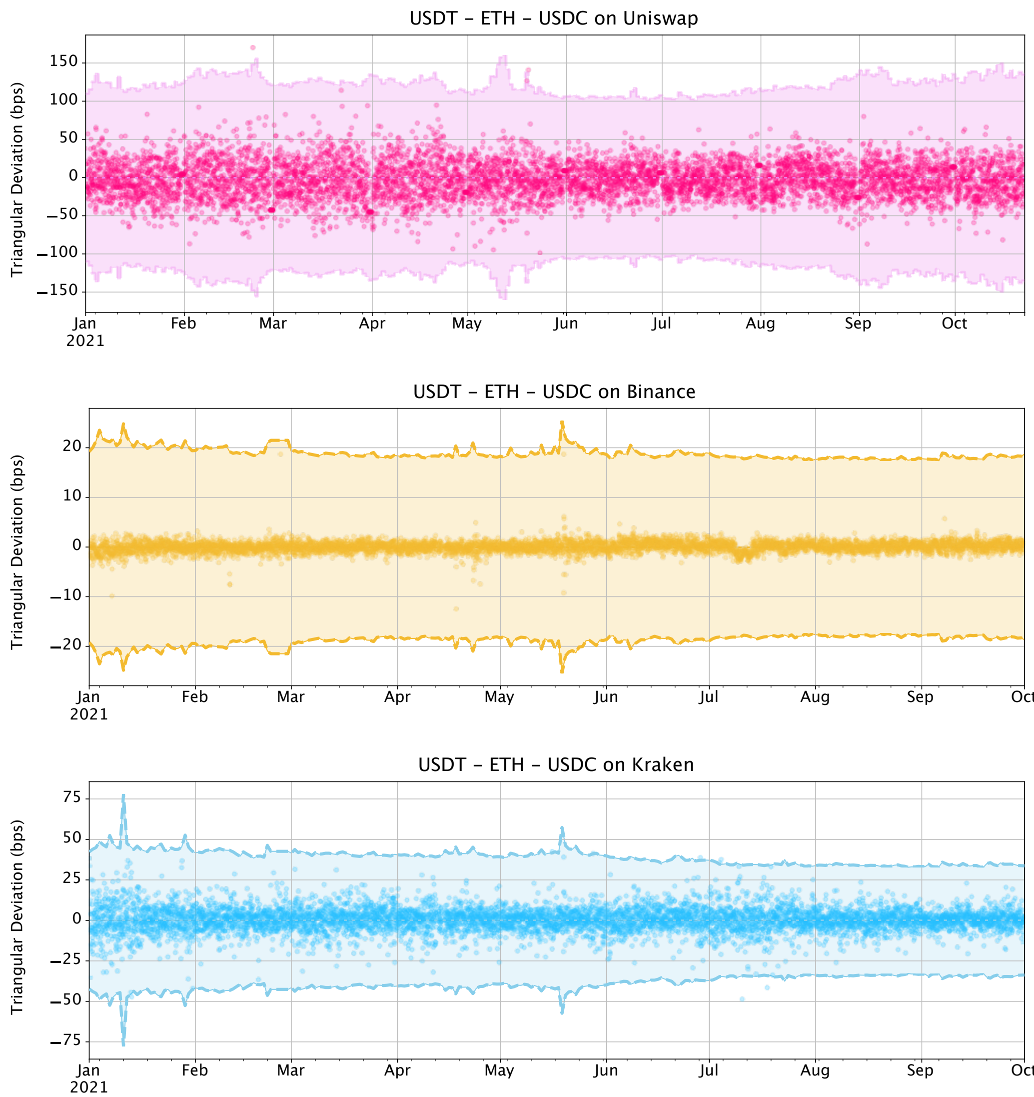

Figure 6. Arbitrage Bounds. The figure presents the estimated triangular arbitrage bounds $[\theta^L, \theta^H]$ and the realized deviations θ from the law of one price for the triplet USDC-USDT-ETH from January, 2021 to December, 2021. The top panel presents arbitrage bounds for the DEX Uniswap v2, computed as in equation (13) on a daily basis using the optimal trade size Δx^*. The mid panel presents arbitrage bounds for Binance, computed as in equation (11) on a daily basis using an infinitesimal trade size Δx, assuming some Binance users are able to trade without incurring the transaction fee (i.e., zero basis points). The bottom panel presents arbitrage bounds for Kraken, computed as in equation (11) on a daily basis using an infinitesimal trade size Δx and assuming a transaction fee of 10 basis points, which is provided to users whose 30-day trading volume is above 10,000 BTC (equivalent to roughly 300 million USD at the time of writing). For all exchanges, triangular price deviations θ based on equation (7) are computed at an hourly frequency.

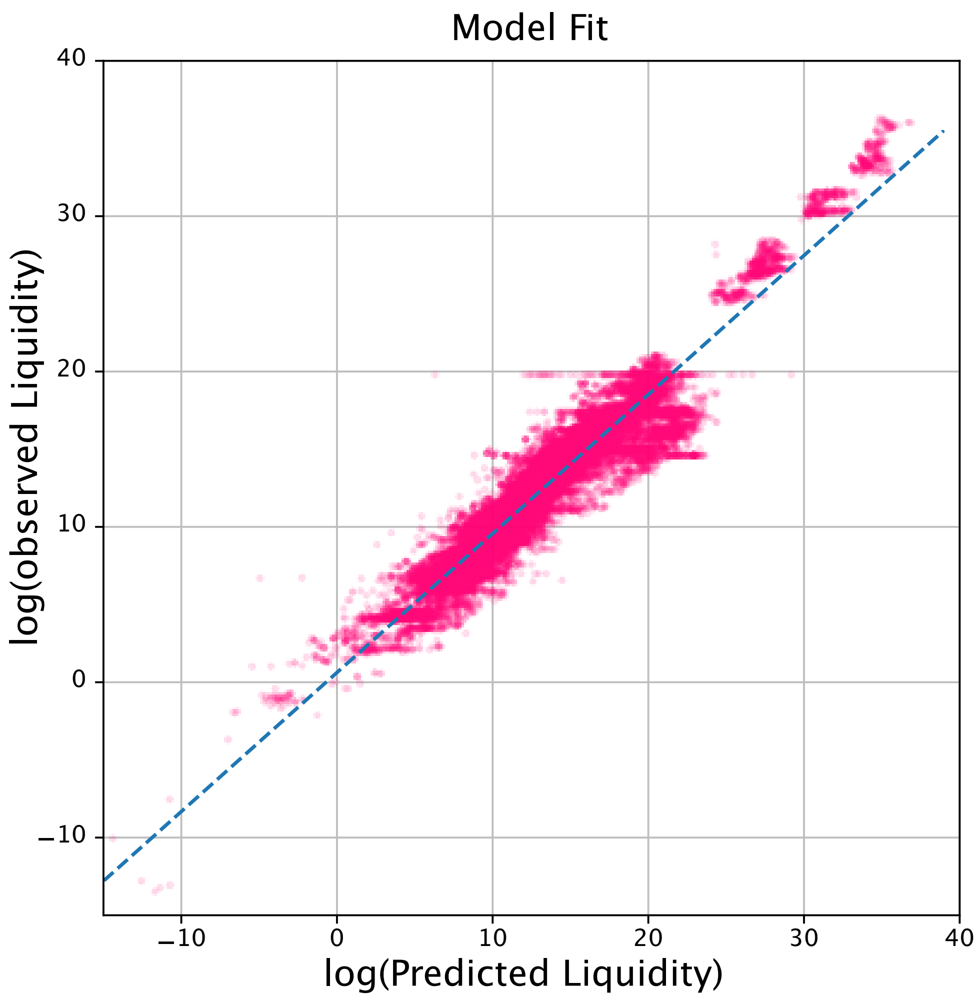

Figure 7. Model Fit. The figure presents a scatter plot of The observed levels of liquidity (y-axis) and those predicted by our model and computed as in equation (15) (x-axis), based on 42, 299 daily observations of 100 exchange pairs quoted in the AMM-based Uniswap from May, 2020 to March, 2022.

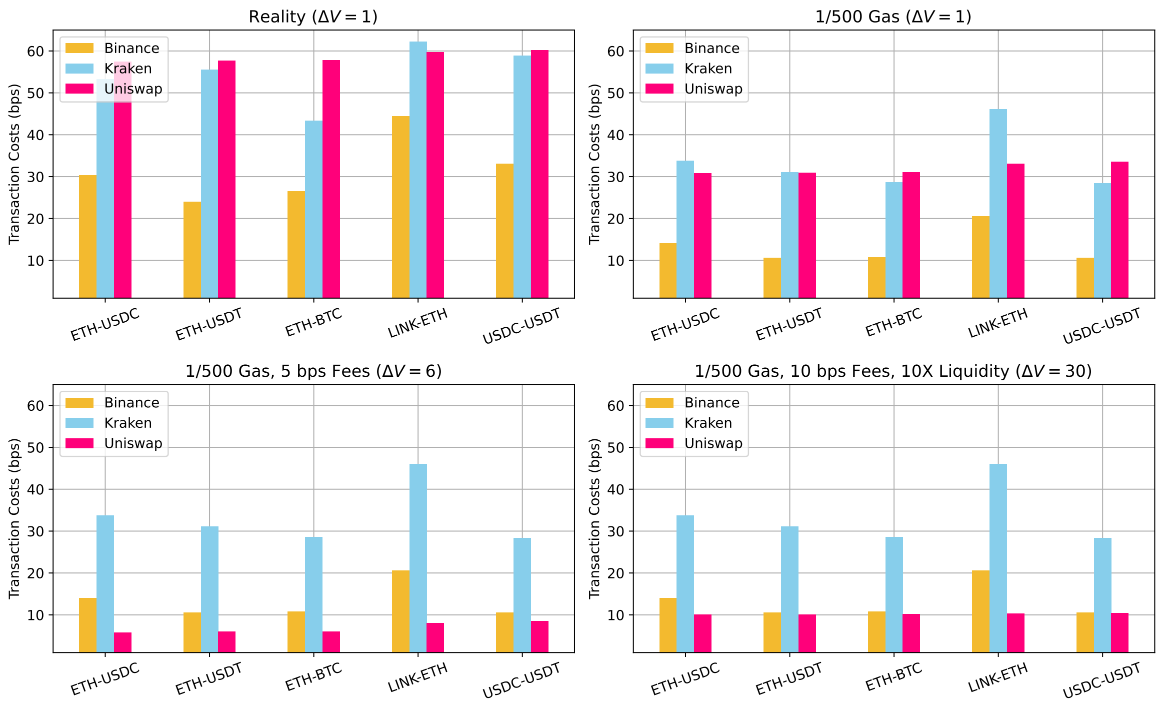

Figure 8. Transaction Costs – Equilibrium. The figure presents price inefficiency levels proxied by the size of arbitrage bounds and computed as in equation (14). These are based on equation (11) for the LOB-based Binance and Kraken and on (13) for the AMM-based Uniswap. They are estimated at an hourly frequency for the five triplets in our sample, then averaged from January, 2021 to December, 2021 and expressed in basis points. As discussed in Section IV, the displayed transaction costs include B/A spreads, exchange fees (30 basis points for Uniswap, 10 basis points for Binance, and 26 basis points for Kraken), and settlement fees (gas fees for Uniswap, and deposit and withdraw fees for the two CEXs). The top-left panel presents the current situation (as in Figure 5), while the other panels report hypothetical scenarios in which the gas cost to execute a swap transaction is reduced by a factor of 500, with different assumptions on the Uniswap exchange fees and trading volume being made. The assumed reduction in gas price could become possible before the end of 2022, when the proof-of-stake version of Ethereum (Ethereum 2.0) is expected to be deployed, and layer-2 solutions are expected to increase their market share.

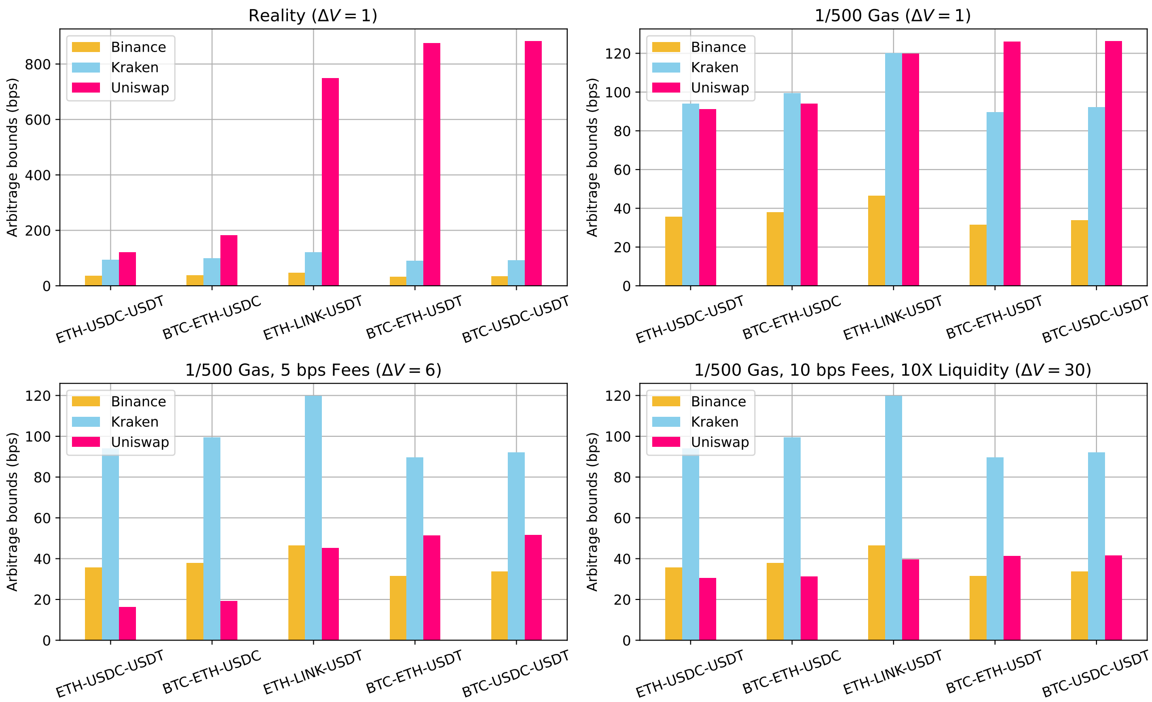

Figure 9. Price Inefficiency – Equilibrium. The figure presents price inefficiency levels proxied by the size of arbitrage bounds and computed as in equation (14). These are based on equation (11) for the LOB-based Binance and Kraken and on equation (13) for the AMM-based Uniswap. They are estimated at an hourly frequency for the five triplets in our sample, then averaged from January, 2021 to December, 2021 and expressed in basis points. The top-left panel presents the current situation (as in Figure 5), while the other panels report hypothetical scenarios in which the gas cost to execute a swap transaction is reduced by a factor of 500, with different assumptions on the Uniswap exchange fees and trading volume being made. The assumed reduction in gas price could become possible before the end of 2022, when the proof-of-stake version of Ethereum (Ethereum 2.0) is expected to be deployed, and layer-2 solutions are expected to increase their market share.

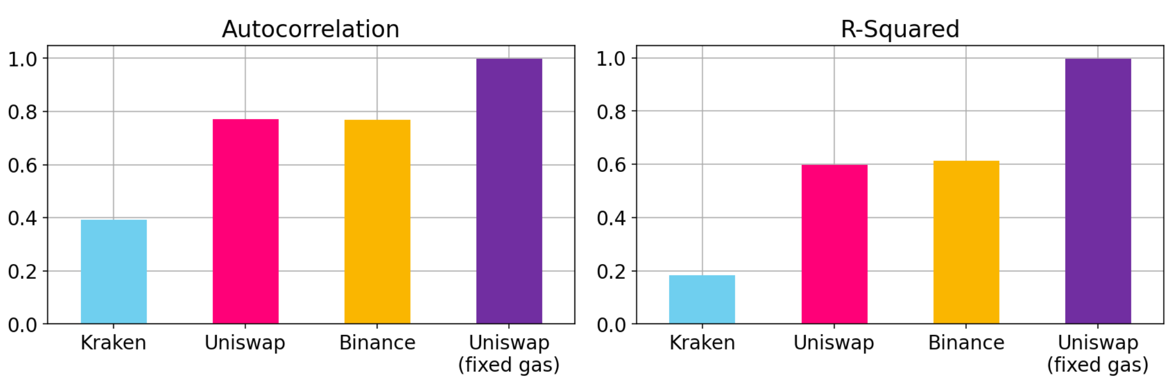

Figure 10. Predictability of Transaction Costs. The figure presents results on the degree of predictability of transaction costs on different exchanges. For each exchange-pair couple, considering hourly transaction costs for a 10,000$ transaction, we compute the auto-correlation coefficient ρ and run the time series regressions $\mathrm{TC}(t) = \alpha + \beta\ \mathrm{TC}(t-1) + \varepsilon(t)$. We then plot the average auto-correlation coefficient ρ (left panel) and the average R^2 from the time-series regressions (right panel).