# On The Quality Of Cryptocurrency Markets

## Centralized Versus Decentralized Exchanges

Andrea Barbon, Angelo Ranaldo*

First Draft: July 2021. This Draft: April 2022

**Abstract**

Despite the growing adoption of decentralized exchanges, little is known about their market quality. Using a novel and comprehensive dataset, we compare decentralized blockchain-based venues (DEXs) to centralized crypto exchanges (CEXs) by assessing two key aspects of market quality: price efficiency and market liquidity. We find that CEXs provide better market quality overall and identify the main friction dampening DEX efficiency as the high gas price stemming from proof-of-work blockchains. We propose and empirically validate a stylized model of DEX liquidity provision, linking trading volume, protocol fees, and liquidity in equilibrium. Our theory identifies the quantitative conditions needed for DEXs to overtake CEXs in the future.

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 and participants in the SNB Technology and Finance Webinars and the World Federation of Exchanges Webinars 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, equity securities, cryptocurrencies, and many other asset classes are traded on centralized exchanges (CEXs). This dominant market structure often relies on an electronic limit order book (LOB), matching end-user orders in a fairly transparent, efficient, and centralized way. Recently however, fueled by the wave of innovation brought about by the advent of blockchain technology, decentralized exchanges (DEXs) have emerged as an alternative and innovative market structure for crypto assets. These venues, based on smart-contract implementations of automated market makers (AMM), have been attracting an increasing amount of 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 analyzing cryptocurrency trading and assessing the market quality of CEXs and DEXs, focusing on market liquidity and price efficiency.[2] We find that CEXs generally enjoy higher liquidity and tighter transaction costs. By analyzing the violation of the law of one price for triplets of exchange pairs, we also find that DEXs enjoy more efficient pricing. We identify the high level of gas fees – that is, the cost of recording transactions on the blockchain that depends on the dollar price of the native token and level of network congestion – as the main cause of price inefficiency of DEX markets. To draw possible paths to efficiency, we provide a simple equilibrium

[1]Throughout the paper, we use the acronym CEX to indicate an LOB-based centralized exchange and DEX to indicate an AMM-based decentralized one. However, the concepts of AMM and LOB are not identical to those of CEX and DEX. While CEXs depend on a proprietary IT platform (server, databases, account system, security, etc.), DEXs rely only on smart contracts and blockchain technology. The distinction between LOB and AMM, meanwhile, relates to the institutional arrangements regarding liquidity provision. In LOB markets, bid and ask limit orders are submitted and updated by market makers and recorded in the order book, while in AMM markets liquidity is posted to pools by market participants, and transaction prices are determined by a mathematical framework – a detailed explanation of the AMM system is provided in Section II. Even though it is possible to envision an LOB-based DEX, an effective implementation of an *on-chain* order book is currently unfeasible due to the limited speed of transaction and non-trivial gas costs. 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.

[2]Market quality is a broad concept that includes price efficiency, liquidity and fairness in the sense that each market participant has an equal chance of obtaining a market price that reflects the fundamental value of the financial security. In this study, we focus on the first two aspects and only briefly discuss fairness.

model capturing the main characteristics of DEX markets and the risk-return trade-off endured by LPs. Using a unique and highly representative dataset, we test the main empirical predictions and quantify the necessary conditions to improve DEX market quality in order to compete with a CEX equivalent.

While academic research on the topic is almost non-existent, assessing the present quality of decentralized AMM markets and their future potential 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 to improve the frictions of traditional markets. For instance, the fact that DEX relies on AMM rather than LOB implies that anyone, no matter who and what degree of sophistication she has, can offer liquidity to the exchange in a completely passive fashion by means of liquidity pools. In addition, the custody of assets remains fully with the user, thus ensuring the highest level of security and censorship resistance. Second, the political discussion has centered on the need to regulate cryptocurrency markets to protect their users and ensure financial stability. To properly address these issues, a thorough analysis of the quality of DEXs is desirable.

Our empirical work leverages a unique and very granular dataset that is comprised of three elements: (i) LOB high-frequency snapshots for the most liquid centralized crypto exchanges (Binance, Kraken, Coinbase), (ii) liquidity pools levels and transactions for the most prominent DEXs (Uniswap, Pancakeswap, Sushiswap), (iii) historical gas prices for the Ethereum blockchain and the Binance Smart Chain. This rich set of information allows us to accurately reconstruct quoted prices and spreads for a selected set of exchange pairs of cryptocurrencies at a per-minute frequency. We proceed in three steps. First, we analyze market liquidity by computing effective transaction costs for each pair and for different sizes that an end user has to bear. For CEXs, our measure for transaction costs is defined as the volume-weighted realized half-spread based on the available limit orders (implementation shortfall), plus the percentage transaction fees charged by the exchange. Similarly, for DEXs, we consider the sum of the realized half-spread (based on the available liquidity in the pools), the percentage transaction fees charged by the protocol, and the gas fees paid to miners operating the relevant blockchain. Two main findings emerged: (i) in general, CEXs feature better market quality in the sense that they offers higher liquidity and lower transaction costs, and (ii) DEXs become competitive for expensive transactions; we quantify this threshold as a transaction volume of more than 100,000$.

Next, we study price efficiency by examining triangular price deviations; that is, the difference between the quoted prices of a triplet of exchange pairs and those implied by no-arbitrage relations. For example, one could buy USD Coin (USDC) against Ethereum (ETH) directly or do so indirectly by first selling ETH for Tether (USDT) and then selling the latter to obtain USDC. 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, once accounting for transaction costs, no-arbitrage conditions are less restrictive for DEXs and result in larger 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 postulate 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 pool owned, 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; it is also called divergence loss), the AMM analogue to adverse selection cost in LOB markets. Hence, the prediction of our model states that the bid-ask spread depends on the equilibrium level of assets staked in a liquidity pool that in turn is proportional to the transaction fees earned by the LP and the expected volume and inversely proportional to the exante IL volatility. We test the empirical predictions with our data and find strong corroboration. Furthermore, based on the insights from our model and two planned institutional changes in the DEX market (i.e., the introduction of Ethereum 2.0 and lower protocol fees), we make quantitative predictions on the possible future evolution of liquidity and price efficiency of DEXs, conditional on expected levels of trading volume. Our analysis suggests that the market quality of DEXs will likely soon catch up with and potentially overtake that of CEXs.

Our contribution to the literature is at least three-fold. First, we provide a systematic analysis of liquidity and price efficiency in decentralized markets based on the AMM paradigm, highlighting the main reasons for the current dominance of CEXs. Specifically, we are the first to shed light on and analyze the crucial role of gas costs. Second, we propose a simple model of liquidity provision in DEX with intuitive economic predictions; that is, the spread depends on the LP's liquidity provision, which in turn is determined by the expected transaction fees, trading volume, and IL. These predictions find strong empirical support and can explain the vast majority of the cross-sectional and time-series variation in observed liquidity levels. Third, the theoretical and empirical analysis we apply to the anticipated institutional improvements indicates that DEXs are not destined to be shelved. On the contrary, DEX systems are likely to become 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, while 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 VIII concludes.

# II. AMM Markets

## *A. High-level description of AMM Markets*

Most exchanges in contemporary financial markets use a central LOB system that often requires a central institution to maintain a record of available and executed buy and sell orders. On this type of exchange, the market price is determined by the most recently matched buy 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 is very costly and slow due to the time-consuming mining process and gas fees paid to miners. Furthermore, blockchain technology by its very conception has a limited storage capacity, a resource that is sorely needed in order-based exchanges. Crypto exchanges such as Binance, Coinbase, or Kraken that still 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 the reserve of a smart contract or 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 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 IL. This arises from the fact that the LP receives more of the least valuable asset and less of the most valuable 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 icnrease since their inception. As of early May 2021, there were more than 23 billion U.S. dollars 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 involved with DEXs compared to order-book-based CEXs.

### *B. Salient Features of 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 the relevant advantages and drawbacks of DEXs.

First of all, contrary to CEXs, the custody of assets remains fully with the user, as no third party is required to execute the trade. This feature implies that users can take full advantage of the censorship-resistant and trustless nature of crypto assets based on blockchain technology (Pagnotta and Buraschi, 2018). It also neutralizes the risk of malicious agents (hackers) attacking the exchange and stealing assets, as the exchange does not possess the assets of its customers. Consequently, it allows users to save on the fees commonly associated with the deposit and withdrawal of assets in CEXs.

Second, DEX users can provide liquidity to the exchange in a completely passive fashion. Hence, liquidity provision is accessible to 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 significant in terms of both sophistication and capital. 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).

Third, 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.

Fourth, users can quote any pair of ERC-20 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.

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

### *C. Related Literature*

We contribute to the nascent but growing literature on cryptocurrencies by providing a comprehensive analysis of their market quality. Concerning price efficiency, 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 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 compute 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 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 influence 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) outlines the pay-offs of LPs on AMM exchanges. 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 closely follow the reference market price. Capponi and Jia (2021) model the impact on utility for LPs and traders of the curvature of the pricing function on Uniswap.

The closest paper to our study is Lehar and Parlour (2021), who compare AMM and a limit-order market. Our work differs from theirs in three main aspects. First, we propose a different theoretical model that postulates distinct and intuitive theoretical predictions outlining how liquidity provision arises from transaction fees, expected transaction volume, and IL. Second, while they only consider Binance as a representative CEX market, we analyze data from three CEX systems (Binance, Kraken, and Coinbase). Third, and most importantly, a crucial aspect of our analysis is the time series variation of gas fees, which those authors do not consider. While they find the liquidity provision in Uniswap to be more stable than in Binance, we find that CEX and DEX transaction costs are characterized by the same degree of predictability. Moreover, by properly accounting for oscillations in the gas price, transaction costs on Uniswap turn out to be only partially predictable and stable. We also find a significant co-movement of the effective transaction costs on Binance and the price of gas on the Ethereum blockchain, which is suggestive evidence of gas prices as a proxy for the aggregate demand of immediacy. Finally, 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.

Our contribution to the literature is three-fold: First, we provide a systematic analysis of price efficiency by studying the triangular no-arbitrage conditions based on a unique and very comprehensive set of cryptocurrencies. We determine the arbitrage boundary conditions and test their violations. Second, we investigate market liquidity and transaction costs considering all relevant features and trading costs of centralized and decentralized crypto exchanges. By doing so, we quantify the conditions and transaction size under which DEXs become competitive with CEXs. Third, we theorize equilibrium conditions including trading volume and protocol fees to efficiently provide liquidity on AMM exchanges.

### *D. Mathematical Foundations of AMM Markets*

This AMMs, such as Uniswap, use the *constant product rule*, which enables an algebraic determination of market price and transaction price based on the available reserves (Adams et al., 2021). Consider a liquidity pool that contains $x$ tokens of $X$ and $y$ tokens of $Y$. The combined amount of both tokens in the pool determines the current market price $P_{xy}$ or $P_{yx}$, 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 protocol fees charged by the DEX ($f = 0.003$ for Uniswap), 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 $\Delta y$ is mathematically determined through the following equation

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

Solving for $\Delta y$, one obtains that the output 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

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

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

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

# III. Data and Summary Statistics

Because DEXs are based on smart contracts deployed on blockchains, records of every single interaction with those contracts is available to the public. This rich dataset includes as primitives the creation of exchange pairs, the addition/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 point in time. We leverage the application programming interface of TheGraph.com to obtain data for Uniswap from the Ethereum MainNet blockchain. We download data on liquidity pool reserves and volumes at an hourly frequency for the pairs made of the five crypto-tokens that are 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 and Kraken.

## *A. Summary Statistics*

Figure 1 displays daily volumes for the AMM-based exchanges in our sample (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 USD on August 2020 to roughly 1 billion at the end of the sample (December 2021). Similar upward trends are displayed for Sushiswap and Pancakeswap, which were deployed in September 2020. Figure 2 reports trading volumes for the LOB-based exchanges in our sample (Binance, Kraken, and Coinbase). Binance is the dominant exchange in terms of volumes across the entire sample, rising from roughly 4 billion to 23 billion USD.

Volumes on Coinbase and Kraken are comparable; both present a significant upward trend. Figure 3 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 in millions of USD over the January–December 2021 period, which is the focus of our market quality analysis, for the pairs we consider. These pairs provide a representative sample, as they generate roughly one third of the volume on each exchange.

## 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., 2013). 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.

### *A. Variables Definition*

For LOB markets, we observe the full depth of ask and bid quotes present in the order book at any point in time, so that the quoted half-spread associated with a market order can be computed directly using the volume-weighted average price (VWAP).

More specifically, we define the transaction price $T_{XY}$ for a sell order of size $\Delta x$ as

$$T_{XY}(\Delta x) = \frac{\sum_{i=1}^{N} v_i p_i}{\Delta x} \qquad \text{such that} \qquad \sum_{i=1}^{N} v_i = \Delta x \ ,$$

where $v_i$ and $p_i$ represent the volume and the price of each filled limit order $i$. The quoted half-spread for a sell order is thus given by

$$S_{XY}(\Delta x) = \frac{P_{XY} - T_{XY}(\Delta x)}{P_{XY}} \ , \tag{3}$$

where $P_{XY}$ is the quoted mid-price.[3] Finally, we define the transaction costs as the sum of the quoted half-spread and the percentage transaction fees $f$ charged by the exchange:

$$TC_{XY}(\Delta x) = S_{XY}(\Delta x) + f, \tag{4}$$

For AMM exchanges we also need to account for gas fees paid directly to miners that are required to interact with a smart contract and record the transaction on the relevant blockchain. The dollar value of those fees depends on the computational complexity of the smart-contract function being used, 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 fees required to execute a swap transaction; that is, invoking the `swapExactTokensForTokens` function of the relevant router contract.[4] We denote the gas fees as $g$, expressed in units of the traded token $X$. The transaction costs on AMM markets are computed as the sum of the quoted half-spread $S$ defined in (2), the percentage protocol fee $f$, and the gas fee $g$ as a fraction of the trade size

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

[3]For the sake of simplicity, we consider only sell orders. The half-spread can in principle be quantitatively different for buy orders in an LOB-based market if the available liquidity is asymmetric around the mid-price. Nevertheless, re-running the analysis using buy orders does not have a significant impact on our analysis.

[4]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.

The dollar value of gas prices on the Ethereum blockchain exhibit strong time series variation, depending on both the dollar price of the native token (ETH) and the level of network congestion. Figure 4 plots the evolution over time of the gas fees (in USD) required to execute a swap transaction on Uniswap. For our empirical analysis, 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.[5] We then approximate the gas cost of a swap during each hour of our sample period, multiplying $\Gamma$ by the average gas price paid across all blocks verified during that hour. We compute $TC_{XY}$ in AMM and LOB exchanges at the hourly frequency for the six pairs in our sample and for different trade sizes $(10^3, 10^4, 10^5, 10^6)$, expressed in USD.

### *B. Results*

Figure 5 displays log levels of transaction costs for different trade sizes on the AMM-based exchange Uniswap and the LOB-based exchanges Binance and Kraken. The panel on the top left shows that Uniswap is extremely expensive for small-size trades, with transaction costs at roughly 300 bps for all considered pairs. This finding does not come as a surprise, since gas fees, on average, constitute a large percentage (3.2%) of the traded amount. The panel on the top right shows that LOB-based exchanges are also superior to Uniswap for mid-sized transactions of 10,000$. The situation depicted in the panel on the bottom left, for a more significant trade size of 100,000$, is somewhat different. While Binance proves the most convenient choice for five out of six pairs, Uniswap delivers lower transaction costs for the LINK-ETH pair. The panel on the bottom right reinforces the finding, showing that Uniswap is competitive with its centralized counterparts for the four pairs involving Ethereum. It is worth noticing that Binance offers the lowest average transactions costs for all trade sizes, especially for the pairs involving the stable coins USDT and USDC.

[5] We estimate 110,000 by collecting all swap transactions executed on the Uniswap v2 Router contract using the latest 1,000 blocks, and taking their average gas usage. The variation across pairs is minimal for the pairs we consider (± 10,000 gas units at most). This figure is significantly larger than the gas required by a simple *transfer* function on a ERC20 contract, which costs around 10,000 gas units.

# 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 relate it to liquidity levels. 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 $\theta$, defined as

$$\theta = P_{XY}\, P_{YZ}\, P_{ZX} - 1\,, \tag{6}$$

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 profit $\theta$ of a triangular trade has to be higher than the associated 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 $R(\Delta x) > 0$ 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, which are discussed below. The total quoted spread for a triangular trade on $X$, $Y$, and $Z$, is given by

$$S_{XYZ}(\Delta x) = 1 - \Big(1 - S_{XY}(\Delta x)\Big)\Big(1 - S_{YZ}(\Delta y)\Big)\Big(1 - S_{ZX}(\Delta z)\Big), \tag{7}$$

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 (7) 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 $\Delta x$ as a result of the spreads of the previous transactions. The total fees charged, as a percentage of the initial amount $\Delta x$, are

$$F_{XYZ}(\Delta x) = f\Big(1 + (1 - S_{XY}(\Delta x)) + (1 - S_{XY}(\Delta x) \cdot (1 - S_{YZ}(\Delta y))\Big). \tag{8}$$

Given the execution cost $R(\Delta x)$, a triangular arbitrage is profitable if and only if

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

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

### *A. Arbitrage Bounds for LOB Markets*

Arbitrage bounds on LOB markets depend on the quoted spreads, defined in (3) and based on the available liquidity in the LOB, and the transaction fees charged by the exchange. The total quoted spread and the total fees charged, as a percentage of the initial amount $\Delta x$, are defined as in (7) and (8), respectively. Thus, the execution cost of a triangular trade of size $\Delta x$ is given by

$$R(\Delta x) \;=\; S_{XYZ}(\Delta x) \;+\; F_{XYZ}(\Delta x). \tag{9}$$

Note that the fact that we use the *best* quoted spread implies that the trade size $\Delta x$ is infinitesimal. This choice is based on the assumption that, in the absence of fixed transaction costs, arbitrageurs are also willing to perform arbitrage trades with infinitely small dollar amounts, thus minimizing their price impact. Hence, the lower and upper arbitrage bounds for $\theta$ are given by

$$.\theta^H, \theta^L = \pm\Big(S_{XYZ}(\Delta x) \;+\; F_{XYZ}(\Delta x)\Big). \tag{10}$$

*B. Arbitrage bounds for AMM markets*

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 protocol 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 $\Delta x$, are defined as in (7) and (8), respectively. The total gas fees are simply the gas fee for a single swap multiplied by the factor 3. Thus, the total execution cost can be described as (see Appendix B for more details)

$$R(\Delta x) \;=\; S_{XYZ}(\Delta x) \;+\; F_{XYZ}(\Delta x) \;+\; 3g/\Delta x. \tag{11}$$

As in models with entry costs, arbitrageurs face a trade-off between the cost of gas fees and the price impact. The former is reduced (in %) by increasing $\Delta x$, while the latter increases with $\Delta x$. Assuming rationality, they choose the optimal trade size $\Delta x^{\star}$ for which the cost $R(\Delta x)$ 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, $R(\Delta x^\star)$. Hence, the lower and upper arbitrage bounds for $\theta$ are given by

$$\theta^H, \theta^L = \pm\Big(S_{XYZ}(\Delta x^\star) \; + \; F_{XYZ}(\Delta x^\star) \; + \; 3g/\Delta x^\star\Big)\,. \tag{12}$$

*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

$$B = \frac{\theta^H - \theta^L}{2}\,. \tag{13}$$

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 the efficient levels. We construct bounds at the daily frequency for the six triplets in our sample, separately for each exchange. We then compare the bounds to the realized price deviations $\theta$ 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. Graphical representations of the resulting bounds for the triple USDC-USDT-ETH are provided in Figures 7, 8, and 9.

*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 (10) for the LOB-based Binance and Kraken and on (12) for the AMM-based Uniswap. Figure 6 presents the results, displaying the log-levels of price inefficiency for each triplet, as proxied by the size of arbitrage bounds defined in (13). It is evident that the AMM-based 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 is below 200 bps, while for the less liquid ones it rises above 1,000 bps. These estimates are much higher than those for CEXs, which are lower than 100 bps for almost all the considered triplets. Binance was particularly dominant in terms of price efficiency, with bounds ranging from 30 to 50 bps.

The first reason for such a significant discrepancy between AMM- and LOB-based exchanges is transaction fees. The protocol fees of 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 become a more significant determinant of the net profitability of the trade.

The second – and quantitatively most important – cause of such a low level of price efficiency enjoyed by Uniswap is related to the high level of the gas fees required to compensate miners on the Ethereum blockchain. To make up for such a significant fixed cost, triangular arbitrage on ETH-based AMM markets requires trading sizeable amounts in dollar terms. 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 zero-knowledge proof currently employed by the Ethereum network. Miners have to cover the costs of expensive hardware and significant energy consumption and thus require 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. In fact, arbitrage bounds are only slightly larger than transaction fees multiplied by the factor 3, suggesting that trading costs arising from quoted spreads are not as relevant.

# VI. Conditions for DEX Supremacy

In this section, we present a simple theoretical model of liquidity provision on AMM markets, highlighting the main economic trade-off faced by LPs. Solving the model gives rise to a rational linkage 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 CEX in terms of market quality.

## *A. Equilibrium Liquidity*

We model a representative LP who faces the problem of providing the optimal quantity to the exchange pair $X/Y$. We assume the LP is risk-neutral and that the market is perfectly competitive, as in Kyle (1985) and Glosten and Milgrom (1985). At time $t = 1$ the total liquidity in the pools is equal to $x$, and the LP can add or remove liquidity. At time $t > 1$, users start to trade on the pair until the trading stops at $t = 2$. Let the random variables $V$ denote the total traded volume (in units of $X$) and $\Delta P$ denote the gross percentage change in the quoted price $P_{XY}$, respectively, between $t = 1$ and $t = 2$. The profits and losses of the LP depend on two factors. On the one hand, LPs are compensated by pocketing transactions fees applied to traded amounts. On the other hand, a permanent price change leads to an IL for the LP. The IL arises from the fact that providing funds to a liquidity pool is less profitable relative to holding the tokens (Loesch et al., 2021). The resulting IL is given by (see Appendix A for a mathematical derivation)

$$IL = 2\frac{\sqrt{\Delta P}}{\Delta P + 1} - 1 \ < \ 0.$$

It is important to note that $IL$ measures the level of adverse selection faced by LPs, similar to that faced by market makers in LOB markets. In fact, $IL = 0$ if the order flow is uninformed and gives rise to only a temporary price impact ($\Delta P = 1$), while it increases in magnitude in presence of informed order flow, causing a permanent price change ($\Delta P \neq 1$).

Let $\mathrm{E}[V]$ denote the expected volume and $\mathrm{E}[IL]$ the expected $IL$, both known at time $t = 0$. The net expected percentage return $\mathrm{E}[R]$ from providing an amount $x$ of liquidity is equal to

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

. The assumption of perfect competition results in zero expected returns for the LP, hence the equilibrium level of liquidity is

$$x = -\frac{f\,\mathrm{E}[V]}{\mathrm{E}[IL]}, \tag{14}$$

and we should observe the above relationship for which the provided liquidity increases with the expected trading volume and (percentage) protocol fees remunerating the LP, while it decreases with the IL risk. The equilibrium condition (14) has a clear economic interpretation that is conceptually related to standard microstructure models featuring market makers. First, the level of liquidity $x$ provided by the LP determines the quoted spread available to traders, as in (2). Second, as noted above, the expected IL can be thought as a proxy for the level of adverse selection risk faced by the LP. Thus (14) simply says that spreads are increasing in the level of adverse selection; in other words, LPs require a compensation for the losses caused by informed trading.

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. We regress daily log values of empirically observed liquidity on the ones predicted with (14), for 100 exchange pairs over the period from May 2020 to March 2022. Results are reported in Table IV and Figure 10, showing a highly significant positive correlation between predicted and observed liquidity levels, with a remarkable $R^2$ coefficient equal to 87.59%. The results are robust to the inclusion of pair and time fixed effects.

### *B. Supremacy Conditions*

Re-arranging equation 14, we can link trading volume to liquidity and protocol fees as

$$\mathrm{E}[V] = \frac{-\,\mathrm{E}[IL]x}{f} \propto \frac{x}{f}. \tag{15}$$

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 protocol), or a combination of the two. This makes intuitive sense; since higher trading volume corresponds to more fees proceedings 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 an increase in $x$ (reducing the expected returns per addition units of liquidity provided). We can thus use the above 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 from 3 to 600 times with respect to the trading volume consistent with the growth rate as during the last period of our sample. Following our model, a given increase $\Delta V$ 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.

### *C. Results*

Table VI presents predicted levels of transaction costs for each parameters combination. More precisely, hypothetical transactions costs for a 10,000$ transaction executed through Uniswap are reported, expressed in bps. The current situation is represented by the last row, with fees equal to 30 bps and unitary gas and liquidity multipliers. The third row from the bottom, which assumes a reduction of gas fees by a factor 500, shows that transaction costs are roughly halved for most of the pairs. This assumption seems reasonable in the context of Ethereum 2.0 adoption.[6] This result quantifies the fact that the currently high level of gas fees represents a significant friction for DEX efficiency. As shown in Figure 12, such a reduction in transaction costs would lead to a situation in which Uniswap is strictly dominated by Binance, but only partially by Kraken. Our second scenario (B) is depicted in Figure 13 and assumes, in addition to low gas fees, a six-fold increase in volume leading to a six-fold reduction in protocol fees to 5 bps, which is a viable assumption in the context of Uniswap v3.[7] Under these assumptions, Uniswap would be highly competitive with CEXs, offering significantly lower transaction costs with respect to Binance and Kraken for the majority of pairs. Figure 14 presents scenario C, assuming a more sizeable 30-fold increase in trading volume, resulting in a three-fold reduction in protocol fees to 10 bps and a 10-fold increase in pool liquidity. The results show that Uniswap would offer roughly the same level of transaction costs as Binance for most of the pairs. This corroborates one of our previous main findings; that is, the most important friction undermining DEXs arises from high levels of protocol fees rather than low levels of liquidity. All in all, given an increase in trading volume, it would thus be preferable to reduce fees (as in scenario B) versus attracting more 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 protocol fees.

[6]The introduction of Ethereum 2.0, with ZK Rollups and data sharding implemented, is expected to allow for around 10 million transaction per seconds, 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 network activity, we thus get to an effective reduction in the gas price by a factor of 500.

[7]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 three distinct liquidity pools for the same exchange pair, featuring 5, 30, and 100 bps. This new feature effectively allow LPs to exploit the trade-off highlighted in our model, lowering their required fees to attract or respond to an increase in trading volume.

Moving to price efficiency, the results are reported in Table VII, which presents predictions on the degree of price inefficiency for each scenario. More precisely, the table presents the size of arbitrage bounds for each exchange triplet, computed as in (13). These are based on (10) for the LOB-based Binance and Kraken and on (12) for the AMM-based Uniswap. The current situation is represented by the last row, with fees equal to 30 bps and trivial gas and liquidity multipliers. The third row from the bottom, 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 a 30% increase in efficiency, while the benefits for the other triplets are greater than 80%. This heterogeneous effect of gas fees depends on the diverse size of optimal triangular trades. Since the first two triplets enjoy higher liquidity, the optimal trade size is larger, reducing the impact of gas fees – a fixed cost – on the profitability of potential triangular arbitrages.

Overall, these results highlight the fact that high levels of gas fees represent a significant friction for price efficiency on DEXs and that this effect is more important for triplets involving low-liquidity pairs. As shown in Figure 16, such an improvement in price efficiency would lead to a situation in which Uniswap is still dominated by CEXs. The second scenario (B) is depicted in Figure 17 and assumes, on top of low gas fees, a six-fold increase in volume leading to a six-fold reduction in protocol 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. Figure 18 presents scenario C, assuming a more sizeable 30-fold increase in trading volume, resulting in a three-fold reduction in protocol 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 protocol fees and by low liquidity, depending on which friction is more pronounced for the pairs composing the triplets.

## VII. Transaction Cost Determinants

In equation (5), we showed that the DEX transaction cost is the sum of three components: the quoted half-spread, the percentage protocol 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 that arises is how the bid-ask spread and gas costs impact transaction costs (while the protocol fee is exogenously fixed). This question is particularly relevant for market participants who want to predict transaction costs. To shed more light on this issue, we analyze the degree of predictability of effective transaction costs on Binance and Uniswap, focusing on the marginal impact of gas prices on the latter. 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 $\rho$ of transaction costs associated with a 10,000 \$ trade, for every exchange-pair couple. Second, we run for each exchange and each pair the following time series regression at the hourly frequency:

$$\mathrm{TC}(t) = \alpha + \beta\ \mathrm{TC}(t-1) + \varepsilon(t)$$

. For Uniswap, we repeat the exercise and impose a time-invariant gas price equal to the sample average. Figure 19 plots the cross-pair average $\rho$ 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 is removed due to fluctuating gas prices, 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 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. We conclude our analysis by examining whether gas fees can also affect transaction costs in CEX markets. This is an interesting question since, as shown in equation (5), the CEX transaction cost is determined only by the quoted half-spread and the percentage transaction fees. On the other hand, gas levels may capture the aggregate demand for immediacy in the overall cryptocurrency market. For instance, in the presence of new fundamental information, market participants rush into trading and have incentives to pay higher gas prices to gain priority of execution. In LOB markets, new information is associated with widening bid/ask spreads and thus increased transaction costs (Glosten and Milgrom, 1985). If gas prices proxy for new information, they should therefore also exhibit positive time series correlations with CEX spreads and not only with DEX spreads. To test our hypothesis, we run an hourly frequency regression of transaction costs of Binance on hourly levels of gas prices. Table V shows that transaction costs in Binance are indeed positively related to Ethereum gas prices, thus providing suggestive evidence that the latter proxy for overall demand for immediacy. Future research can study the mechanisms that determine gas costs more thoroughly and how they determine the supply of and demand for liquidity.

## VIII. Conclusion

Our analysis of the market quality of cryptocurrency exchanges highlights a number of important conclusions on the weaknesses and future potential of decentralized venues based on automated market making. First of all, the data shows that Uniswap and the other prevalent DEXs experienced a steep rise in adoption and trading volume during the last year, accompanied by a significant increase in available liquidity. Second, we provide evidence showing that DEXs are not yet of sufficient quality to compete with the largest CEXs regarding transaction costs and price efficiency. These two facts can be reconciled by the observation that DEXs provide a number of advantages over CEXs, especially in terms of security, censorship resistance, and accessibility. It is thus reasonable to speculate that end users value these features and are willing to pay a premium by using decentralized rather than centralized venues.

We highlight multiple factors determining the degree of market quality of DEXs, including the amounts of capital staked in liquidity pools and the level of transaction fees charged by the platform. Nevertheless, high levels of gas fees required by proof-of-work blockchains constitute the most significant friction harming the market quality of DEXs. Price efficiency is particularly harmed since large amounts of capital are required to make arbitrage trades profitable despite the fixed cost associated with transactions. This is particularly relevant for triangular price deviations, which require three distinct transactions to be executed. While high gas prices are not the main determinant of transaction costs for small trades, those involving large amounts are less impacted by gas fees in percentage terms.

Our equilibrium model of liquidity provision in DEXs clarifies the risk-return tradeoff faced by LPs, based on IL 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 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 stacked liquidity and trading fees required by LPs in equilibrium. We analyze a number of future scenarios based on realistic assumptions for future levels of trading volume routed through DEXs, concluding that they could soon become as efficient as CEXs under relatively modest increases in volume, provided that gas costs decline thanks to new proof-of-stake blockchains. We argue that, given the positive outlook toward a future improvement in efficiency and the valued utility in terms of security and censorship resistance, DEXs based on AMM could soon offer a competitive alternative to CEXs. More broadly, the innovative market structure of decentralized crypto exchanges could be applied to other asset classes in the future.

# 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.

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

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

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, A. Roell, and A. Röell. *Market liquidity: theory, evidence, and policy*. Oxford University Press, 2013.

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

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.

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

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

A. S. Kyle. Continuous auctions and insider trading. *Econometrica: Journal of the Econometric Society*, 53(6):1315–1335, 1985.

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.

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.

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

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

# Appendix

## *A. LP Returns and IL*

Assume an LP owns a share $s$ of a liquidity pool containing tokens $x$ and $y$ at time $t = 0$ and the current price is $P_0 = \frac{y_0}{x_0}$. At time $t = 0$, the value of her position in a unit of $y$ is

$$s(x_0 P_0 + y_0) = 2sy_0$$

. At time $t = 1$, the value of her position in a unit of $y$ changes to $s(x_1 P_1 + y_1) = 2sy_1$. The liquidity provision return can be expressed as

$$R_{LP} = \frac{2sy_1}{2sy_0} = \frac{y_1}{y_0}$$

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

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

. Similarly, at $t = 1$, the value in a unit of the LP position is

$$y_1 = \sqrt{kP_1}$$

. Hence, the LP return depends solely on the price change between $t = 0$ and $t = 1$:

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

. The return $R_H$ from holding the tokens is

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

. The IL from providing liquidity instead of holding on the tokens is therefore

$$IL = \frac{R_L P}{R_H} - 1 = 2\frac{\sqrt{\Delta P}}{\Delta P + 1} - 1$$

. Thus, the maximal IL is 0 when there is no price change ($\Delta P = 1$); otherwise $IL \leq 0$.

*B. Arbitrage bound on AMMs*

To derive the expression for the price impact $\rho$, notice that market price and effective transaction price are given by

$$P_{XY} = \frac{y}{x} \qquad \text{and} \qquad T_{XY} = \frac{x}{x + (1-f)\Delta x}$$

. We define the price impact $\rho_X$ as follows:

$$\rho_X = 1 - S_{XY} = 1 - \frac{\Delta x}{x + (1-f)\Delta x} = \frac{x}{x + (1-f)\Delta x}$$

As the transaction price can be expressed using the price impact and market price, we rewrite the output amount from the first trade as follows:

$$\Delta y = \frac{(1-f)\Delta x y}{x + (1-f)\Delta x} = (1-f)\rho_X P_{XY} \Delta x$$

Using the previous expression, we can express $\rho_Y$ as a function of $\Delta x$:

$$\rho_Y = \frac{y}{y + (1-f)\Delta y} = \frac{y}{y + (1-f)^2 \rho_X P_{XY} \Delta x}.$$

Similarly, for $\rho_Z$,

$$\rho_Z = \frac{z}{z + (1-f)\Delta z} = \frac{z}{z + (1-f)^3 \rho_X \rho_Y P_{XY} P_{YZ} \Delta x}.$$

The cumulative price impact and total spread are given by $\rho = \rho_X \rho_Y \rho_Z$ and $S_{XYZ} = 1 - \rho$. Using the previously defined price impact, we define the total platform fees charged as

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

Notice that the fact that the trade sizes for the second and third trades decrease due to the price impact and previous platform fees charged.

Therefore, we can represent the cost of an triangular trade $R(\Delta x)$ as follows:

$$R(\Delta x) = S_{XYZ}(\Delta x) + F_{XYZ}(\Delta x) + 3g/\Delta x) = 1 - \rho + f(1 + \rho_X(1-f) + \rho_X \rho_Y (1-f)^2) + 3g/\Delta x.$$

Based on the cost of the triangular trade, we can express the arbitrage bounds

$$\theta^H, \theta^L = \pm\Big(\rho - 1 - f(1 + \rho_X(1-f) + \rho_X \rho_Y (1-f)^2) - 3g/\Delta x\Big).$$

# Tables and Figures

**Table I. Liquidity pools.** The table reports summary statistics on the liquidity pools underlying the 1,000 most liquid exchange pairs in our sample. We report the total value of liquidity in USD, the number of swap transactions, and the time since the pool was initiated in days.

| | N | Mean | Std | 1% | 10% | 50% | 90% | 99% |
|---|---|---|---|---|---|---|---|---|
| Liquidity (Million USD) | 1000 | 6.71 | 41.52 | 0.01 | 0.03 | 0.99 | 8.32 | 114.54 |
| Transactions (thousands) | 1000 | 37.54 | 121.69 | 7.24 | 8.28 | 18.13 | 62.77 | 238.35 |
| Age (days) | 1000 | 193.05 | 96.03 | 2.33 | 47.04 | 204.09 | 333.69 | 347.26 |

**Table II. Trading volume per pair.** The table reports the daily average trading volume in million USD over the January–December 2021 period for each pair in our sample. The percentage of the aggregate volume represented by these pairs on each exchange is reported below. For pairs involving USDC on Kraken, we report the volume for the corresponding pair based on USD, since volumes in USDC-based pairs are close to zero.

| Pair | Binance | Kraken | Uniswap |
|---|---|---|---|
| ETH-USDC | 60.39 | 202.87 | 120.9 |
| ETH-USDT | 2223.72 | 11.39 | 114.58 |
| ETH-BTC | 538.05 | 33.43 | 36.16 |
| LINK-ETH | 9.65 | 0.78 | 14.1 |
| USDC-USDT | 156.27 | 0.0 | 5.61 |
| BTC-USDC | 119.36 | 257.39 | 2.17 |
| Fraction of Total Volume | 24.71% | 30.45% | 29.15% |

**Table III. Impermanent loss.** The table shows summary statistics of daily impermanent loss for 100 pairs traded on Uniswap over the period between May 2020 and March 2022, aggregated at different levels.

| Level | N | Mean | Std | 1% | 10% | 50% | 90% | 99% |
|---|---|---|---|---|---|---|---|---|
| Pair-Day | 42,299 | -0.00162 | 0.01254 | -0.02180 | -0.00239 | -0.00013 | -0.00000 | -0.00000 |
| Pair | 100 | -0.00210 | 0.00291 | -0.01383 | -0.00428 | -0.00135 | -0.00019 | -0.00000 |
| Day | 351 | -0.00201 | 0.00324 | -0.01752 | -0.00435 | -0.00107 | -0.00042 | -0.00012 |

**Table IV. Model fit.** The table reports results from a panel regression of observed liquidity levels in logs onto log liquidity levels predicted by our model and computed as in (14). Both the dependent and independent variables are computed at the pair-day level for 100 exchange pairs over the period from May 2020 to March 2022. 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 levels. Asterisks denote significance levels (***= 1%, **= 5%, *= 10%).

| | (1) | (2) | (3) | (4) |
|---|---|---|---|---|
| Dependent Variable | Log(Liquidity) | Log(Liquidity) | Log(Liquidity) | Log(Liquidity) |
| Log(Predicted Liquidity) | 0.88*** | 0.56*** | 0.89*** | 0.53*** |
| | (37.97) | (9.46) | (39.39) | (9.27) |
| Constant | 5.92*** | | | |
| | (21.71) | | | |
| Observations | 42,299 | 42,299 | 42,299 | 42,299 |
| R-squared | 0.88 | 0.46 | 0.89 | 0.41 |
| Pair Fixed Effects | - | Yes | - | Yes |
| Date Fixed Effects | - | - | Yes | Yes |
| Ses Clustered By | Pair-Date | Pair-Date | Pair-Date | Pair-Date |

**Table V. Binance transaction costs and gas prices.** The table reports results on hourly time series regression of transaction costs on Binance (averaged across pairs) on the average price of gas from Ethereum blocks validated during the same hour. We saturate the regression model with day fixed effects. T-stats are reported in parentheses, based on robust standard errors. Asterisks denote significance levels (***= 1%, **= 5%, *= 10%).

| | (1) | (2) |
|---|---|---|
| Dependent Variable | Binance TCs | Binance TCs |
| Constant | 12.38*** | 12.39*** |
| | (115.81) | (91.12) |
| Gas Price | 0.01*** | 0.01*** |
| | (6.39) | (5.177) |
| Observations | 6,802 | 6,802 |
| R-squared | 0.15 | 0.83 |
| Date Fixed Effects | - | Yes |

**Table VI. Uniswap hypothetical transaction costs.** The table displays hypothetical transactions costs for a 10,000$ transaction executed through Uniswap, expressed in bps. They are computed as in (5), first at the hourly frequency and then averaged over the period from January 2021 to December 2021 for different gas fee reduction factors (*Gas*), platform fees (*Fees*, in bps), and liquidity multipliers (*Liq*). Each row represents a potential future scenario requiring an increase in trading volume, as predicted by our model, equal to $\Delta V$.

| Fees | Liq | Gas | $\Delta V$ | BTC ETH | BTC USDC | ETH USDT | LINK ETH | USDC ETH | USDC USDT |
|---|---|---|---|---|---|---|---|---|---|
| 5 | 100 | 500 | 600 | 5.06 | 5.59 | 5.06 | 5.08 | 5.06 | 5.09 |
| 5 | 100 | 10 | 600 | 7.68 | 8.21 | 7.68 | 7.70 | 7.68 | 7.71 |
| 5 | 100 | 1 | 600 | 31.73 | 32.25 | 31.73 | 31.75 | 31.73 | 31.75 |
| 5 | 10 | 500 | 60 | 5.16 | 10.42 | 5.15 | 5.36 | 5.13 | 5.40 |
| 5 | 10 | 10 | 60 | 7.78 | 13.03 | 7.76 | 7.98 | 7.75 | 8.02 |
| 5 | 10 | 1 | 60 | 31.82 | 37.08 | 31.81 | 32.02 | 31.80 | 32.06 |
| 5 | 1 | 500 | 6 | 6.09 | 58.34 | 5.98 | 8.09 | 5.83 | 8.53 |
| 5 | 1 | 10 | 6 | 8.71 | 60.96 | 8.60 | 10.71 | 8.45 | 11.15 |
| 5 | 1 | 1 | 6 | 32.75 | 85.00 | 32.65 | 34.75 | 32.49 | 35.19 |
| 10 | 100 | 500 | 300 | 10.06 | 10.59 | 10.06 | 10.08 | 10.06 | 10.09 |
| 10 | 100 | 10 | 300 | 12.68 | 13.21 | 12.68 | 12.70 | 12.68 | 12.71 |
| 10 | 100 | 1 | 300 | 36.73 | 37.25 | 36.73 | 36.75 | 36.73 | 36.75 |
| 10 | 10 | 500 | 30 | 10.16 | 15.41 | 10.15 | 10.36 | 10.13 | 10.40 |
| 10 | 10 | 10 | 30 | 12.78 | 18.03 | 12.76 | 12.98 | 12.75 | 13.02 |
| 10 | 10 | 1 | 30 | 36.82 | 42.07 | 36.81 | 37.02 | 36.80 | 37.06 |
| 10 | 1 | 500 | 3 | 11.09 | 63.31 | 10.98 | 13.09 | 10.83 | 13.53 |
| 10 | 1 | 10 | 3 | 13.71 | 65.93 | 13.60 | 15.71 | 13.45 | 16.14 |
| 10 | 1 | 1 | 3 | 37.75 | 89.98 | 37.64 | 39.75 | 37.49 | 40.19 |
| 30 | 100 | 500 | 100 | 30.06 | 30.59 | 30.06 | 30.08 | 30.06 | 30.09 |
| 30 | 100 | 10 | 100 | 32.68 | 33.21 | 32.68 | 32.70 | 32.68 | 32.71 |
| 30 | 100 | 1 | 100 | 56.73 | 57.25 | 56.73 | 56.75 | 56.73 | 56.75 |
| 30 | 10 | 500 | 10 | 30.16 | 35.40 | 30.15 | 30.36 | 30.13 | 30.40 |
| 30 | 10 | 10 | 10 | 32.77 | 38.02 | 32.76 | 32.97 | 32.75 | 33.02 |
| 30 | 10 | 1 | 10 | 56.82 | 62.06 | 56.81 | 57.02 | 56.79 | 57.06 |
| 30 | 1 | 500 | 1 | 31.09 | 83.21 | 30.98 | 33.08 | 30.83 | 33.52 |
| 30 | 1 | 10 | 1 | 33.70 | 85.83 | 33.60 | 35.70 | 33.45 | 36.14 |
| 30 | 1 | 1 | 1 | 57.75 | 109.87 | 57.64 | 59.74 | 57.49 | 60.18 |

**Table VII. Uniswap hypothetical price inefficiency.** The table displays hypothetical levels of price inefficiency on the Uniswap exchange, expressed in bps. They are estimated as in (12), first at the hourly frequency and then averaged over the period from January 2021 to December 2021, for different gas fee reduction factors (*Gas*), platform fees (*Fees*, in bps), and liquidity multipliers (*Liq*). Each row represents a potential future scenario requiring an increase in trading volume, as predicted by our model, equal to $\Delta V$.

| Fees | Liq | Gas | $\Delta V$ | USDC USDT ETH | USDC BTC ETH | USDT BTC ETH | LINK USDT ETH | USDC USDT BTC |
|---|---|---|---|---|---|---|---|---|
| 5 | 100 | 500 | 600 | 15.23 | 15.51 | 19.72 | 18.72 | 19.76 |
| 5 | 100 | 10 | 600 | 16.72 | 18.68 | 48.38 | 41.33 | 48.69 |
| 5 | 100 | 1 | 600 | 20.46 | 26.67 | 120.27 | 98.14 | 121.25 |
| 5 | 10 | 500 | 60 | 15.56 | 16.64 | 29.93 | 26.78 | 30.07 |
| 5 | 10 | 10 | 60 | 19.06 | 26.67 | 120.27 | 98.14 | 121.23 |
| 5 | 10 | 1 | 60 | 27.69 | 51.88 | 344.97 | 276.52 | 347.97 |
| 5 | 1 | 500 | 6 | 16.69 | 20.21 | 62.18 | 52.23 | 62.61 |
| 5 | 1 | 10 | 6 | 26.97 | 51.88 | 344.97 | 276.52 | 347.97 |
| 5 | 1 | 1 | 6 | 52.86 | 131.35 | 1029.11 | 827.75 | 1038.35 |
| 10 | 100 | 500 | 300 | 30.21 | 30.49 | 34.69 | 33.69 | 34.74 |
| 10 | 100 | 10 | 300 | 31.69 | 33.66 | 63.33 | 56.29 | 63.64 |
| 10 | 100 | 1 | 300 | 35.41 | 41.63 | 135.17 | 113.06 | 136.14 |
| 10 | 10 | 500 | 30 | 30.54 | 31.62 | 44.90 | 41.75 | 45.04 |
| 10 | 10 | 10 | 30 | 33.99 | 41.63 | 135.17 | 113.06 | 136.13 |
| 10 | 10 | 1 | 30 | 42.77 | 66.83 | 359.70 | 291.30 | 362.71 |
| 10 | 1 | 500 | 3 | 31.66 | 35.19 | 77.12 | 67.18 | 77.55 |
| 10 | 1 | 10 | 3 | 41.94 | 66.83 | 359.70 | 291.30 | 362.71 |
| 10 | 1 | 1 | 3 | 67.82 | 146.24 | 1043.31 | 842.11 | 1052.60 |
| 30 | 100 | 500 | 100 | 89.97 | 90.25 | 94.44 | 93.44 | 94.48 |
| 30 | 100 | 10 | 100 | 91.47 | 93.41 | 122.99 | 115.97 | 123.30 |
| 30 | 100 | 1 | 100 | 95.12 | 101.36 | 194.61 | 172.57 | 195.60 |
| 30 | 10 | 500 | 10 | 90.30 | 91.38 | 104.61 | 101.47 | 104.75 |
| 30 | 10 | 10 | 10 | 93.75 | 101.36 | 194.61 | 172.57 | 195.57 |
| 30 | 10 | 1 | 10 | 102.37 | 126.48 | 418.47 | 350.27 | 421.48 |
| 30 | 1 | 500 | 1 | 91.42 | 94.93 | 136.74 | 126.83 | 137.17 |
| 30 | 1 | 10 | 1 | 101.67 | 126.48 | 418.47 | 350.27 | 421.48 |
| 30 | 1 | 1 | 1 | 127.45 | 205.65 | 1100.00 | 899.42 | 1109.41 |

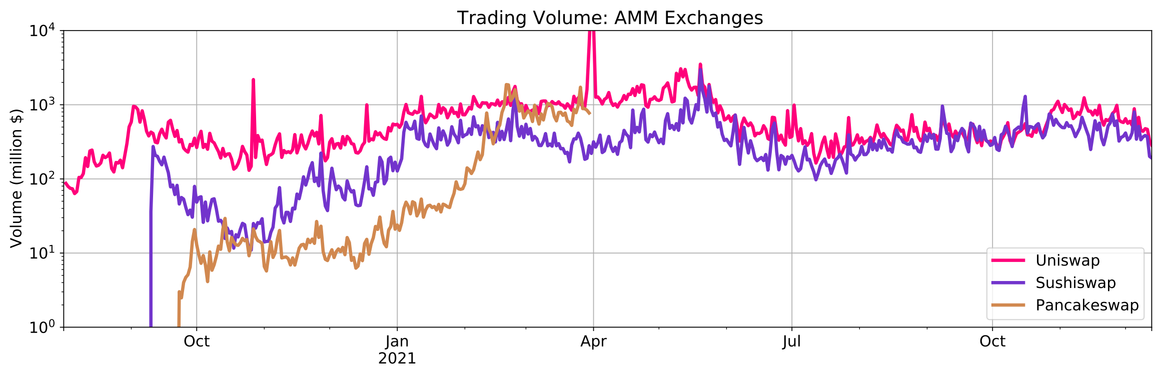


**Figure 1. DEX volume.** The figure presents traded volumes for the AMM-based exchanges in our sample; namely, Uniswap, PancakeSwap, and SushiSwap, for the period from August 2020 to December 2021. The displayed traded volumes are the sum of the volumes for all trading pairs listed on each exchange. The vertical axis uses log-scale and is reported in million USD.

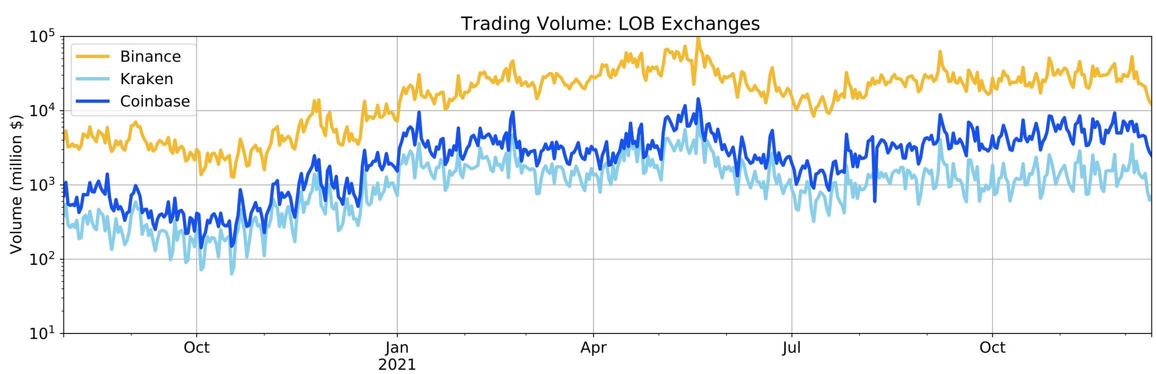


**Figure 2. CEX volume.** The figure presents traded volumes for the LOB-based exchanges in our sample; namely, Binance, Kraken, and Coinbase, for the period from August 2020 to December 2021. The displayed traded volumes are the sum of the volumes for all trading pairs listed on each exchange. The vertical axis uses log-scale and is reported in million USD.

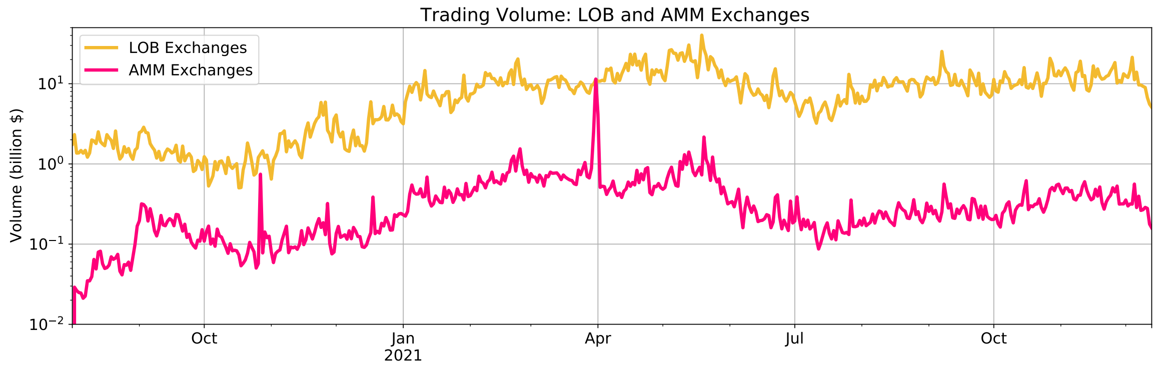


**Figure 3. CEX VS DEX volumes.** The figure presents traded volumes for both the AMM-based and LOB-based exchanges in our sample, averaged across the three exchanges in each category, for the period from August 2020 to December 2021. The displayed traded volumes are the sum of the volumes for all trading pairs listed on each exchange. The vertical axis uses log-scale and is reported in billion USD.

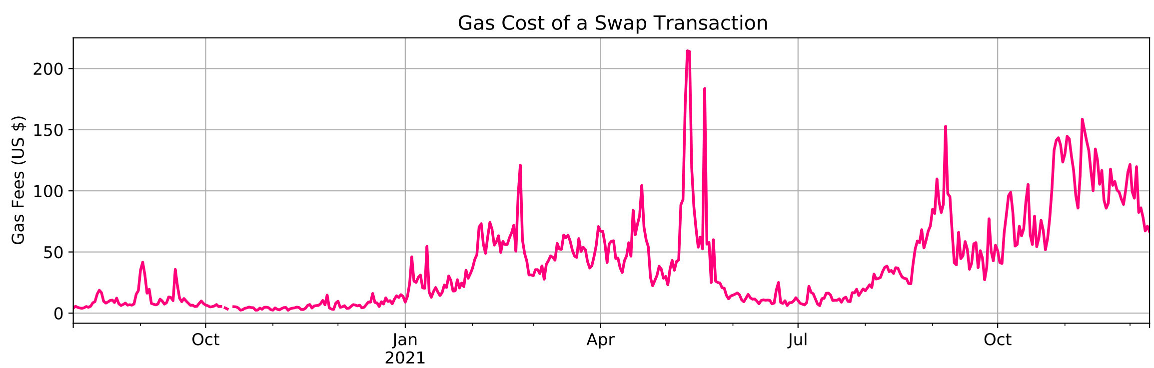


**Figure 4. Gas fees.** The figure presents the time series evolution of the gas costs of a swap transaction in our sample, in USD. This is computed at the hourly frequency, multiplying the units of gas required to execute a swap – roughly 110,000 – by the average gas price (in USD) associated with transactions in the blocks validated during each hour. 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.

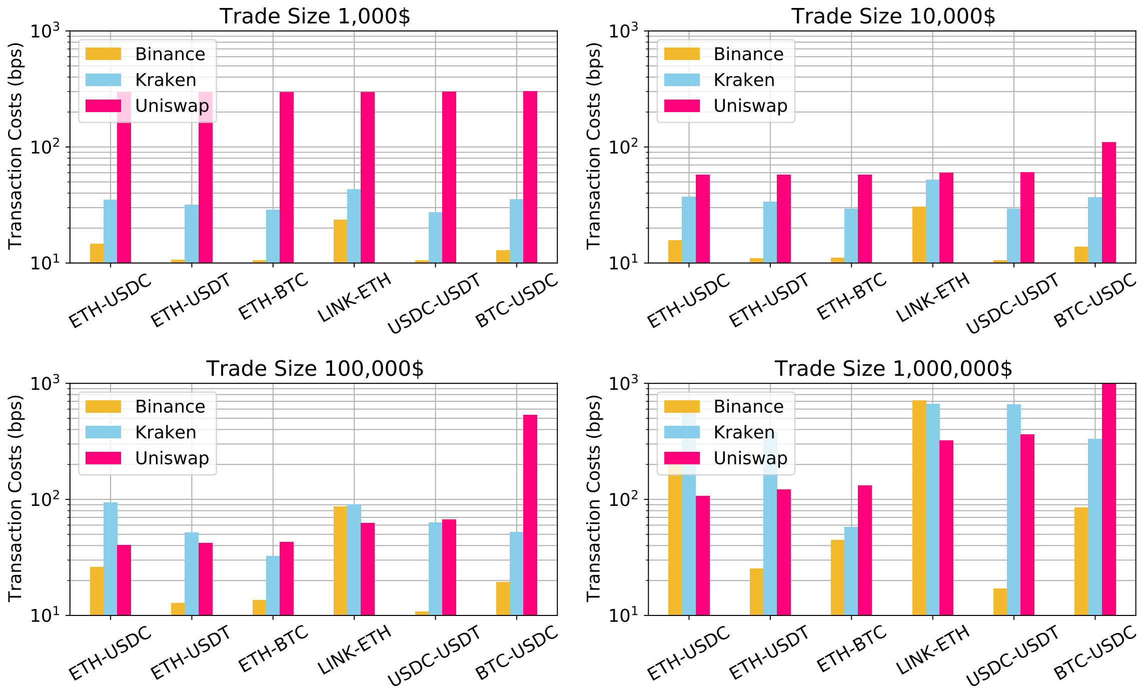


**Figure 5. Transaction costs.** The figure presents transaction costs, computed as in (4) for the LOB-based Binance and Kraken and on (5) for the AMM-based Uniswap. These are computed at the hourly frequency for the six pairs in our sample and for different trade sizes ($10^3, 10^4, 10^5$, and $10^6$ US dollars), then averaged over the period from January 2021 to December 2021. The vertical axis is in log-scale and reported in bps.

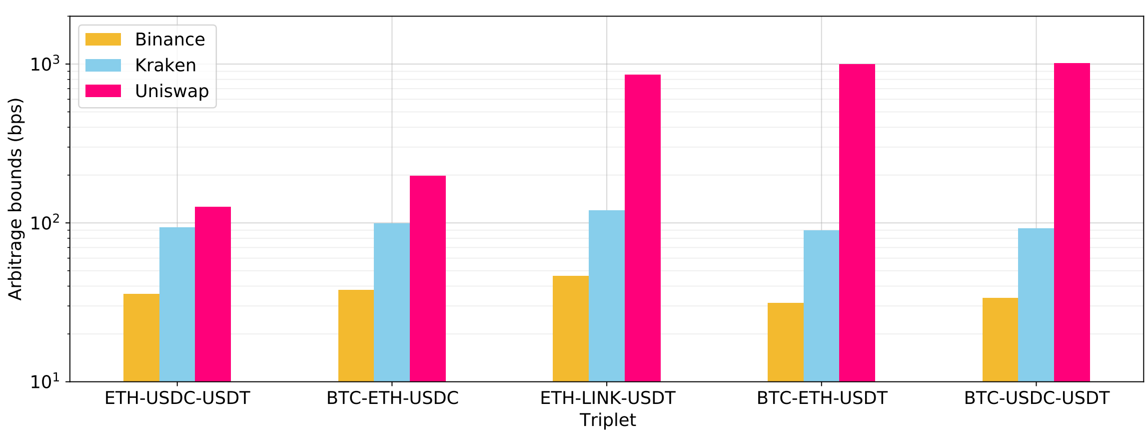


**Figure 6. Price inefficiency.** The figure presents price inefficiency levels, proxied by the size of arbitrage bounds computed as in (13). These are based on (10) for the LOB-based Binance and Kraken and on (12) for the AMM-based Uniswap. They are estimated at the hourly frequency for the five triplets in our sample, then averaged over the period from January 2021 to December 2021. The vertical axis is in log-scale and reported in bps.

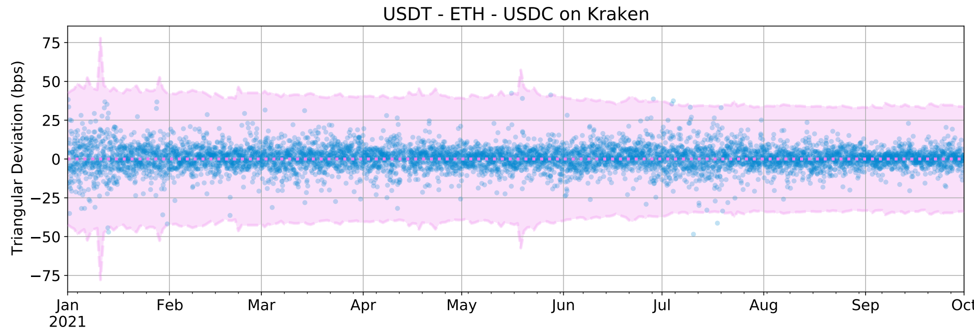


**Figure 7. Arbitrage bounds.** The figure presents the arbitrage bounds and $\theta$ the deviations from the law of one price for the triplet USDC-USDT-ETH on the LOB-based Kraken over the period from January 2021 to December 2021. Arbitrage bounds are computed as in (10) on a daily basis using an infinitesimal trade size $\Delta x$ and assuming a transaction fee of 10 bps, which is provided to users whose 30-day trading volume is above 1,0000 BTC (equivalent to around 300 million USD at the time of writing). Triangular price deviations $\theta$ based on (6) are computed at an hourly frequency.

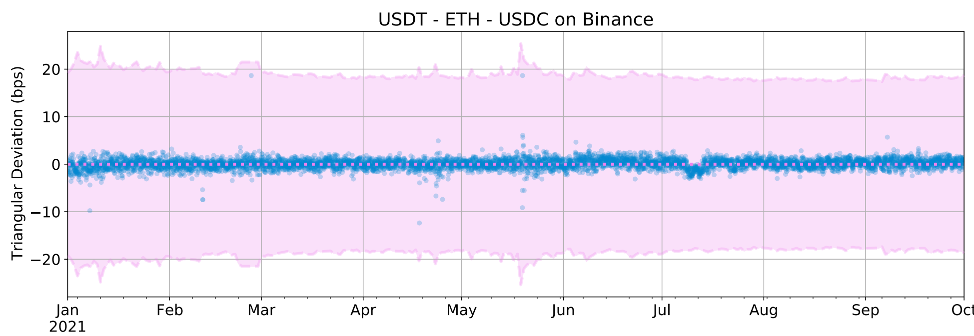


**Figure 8. Arbitrage bounds.** The figure presents the arbitrage bounds and $\theta$ the deviations from the law of one price for the triplet USDC-USDT-ETH on the LOB-based Binance over the period from January 2021 to December 2021. Arbitrage bounds are computed as in (10) on a daily basis using an infinitesimal trade size $\Delta x$, assuming some Binance users are able to trade without incurring the transaction fee (i.e., zero bps). Triangular price deviations $\theta$ based on (6) are computed at an hourly frequency.

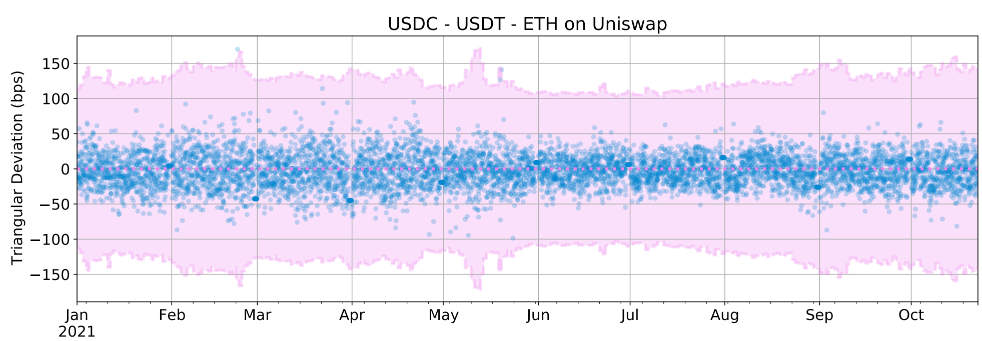


**Figure 9. Arbitrage bounds.** The figure presents the arbitrage bounds and $\theta$ the deviations from the law of one price for the triplet USDC-USDT-ETH on the AMM exchange Uniswap over the period from January 2021 to December 2021. Arbitrage bounds are computed as in (12) on a daily basis using an optimal trade size $\Delta x$. Triangular price deviations $\theta$ based on (6) are computed at an hourly frequency.

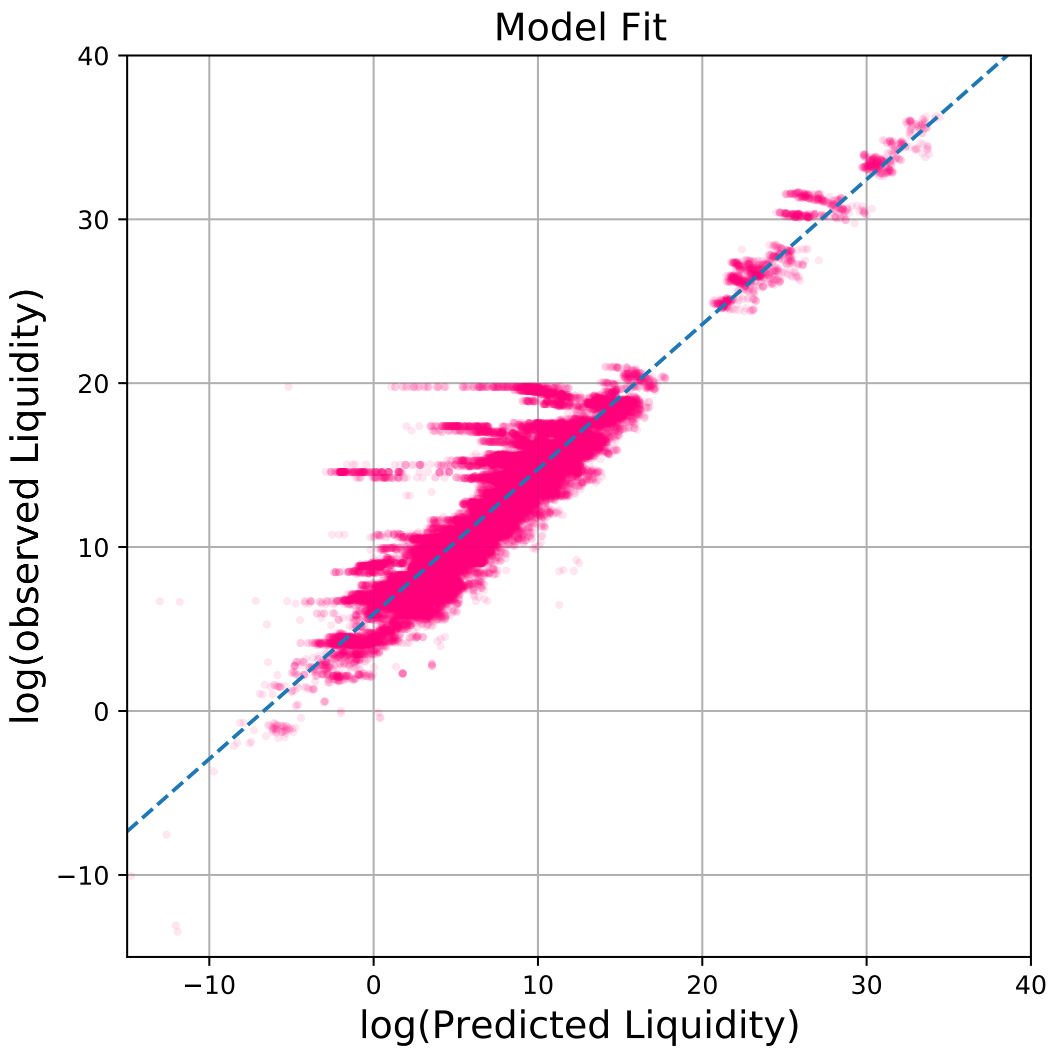


**Figure 10. Model fit.** The figure presents a scatter plot of observed levels of liquidity (y-axis) and those predicted by our model and computed as in (14) (x-axis), based on 42,299 daily observations of 100 exchange pairs quoted in the AMM-based Uniswap over the period from May 2020 to March 2022.

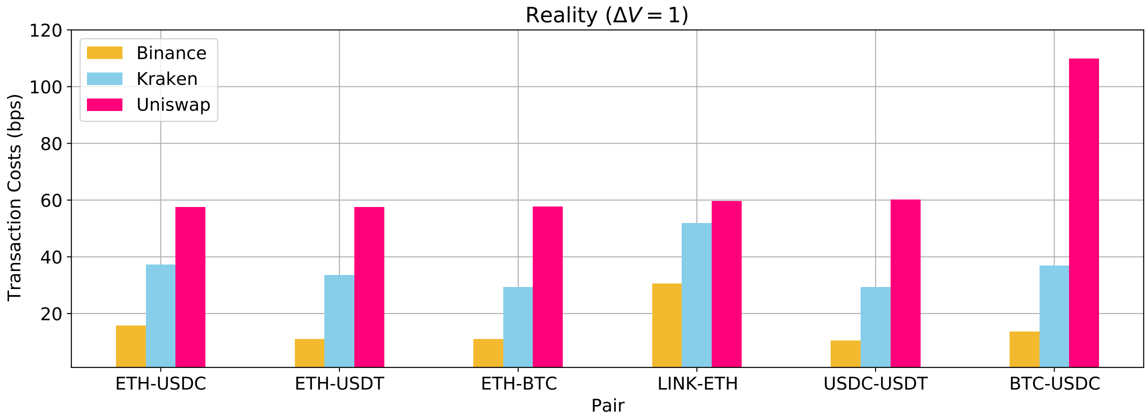


**Figure 11. Transaction costs – Reality.** The figure presents transaction costs for a traded amount of 10,000$, for the six pairs in our sample. These are computed as in (4) for the LOB-based Binance and Kraken and on (5) for the AMM-based Uniswap.

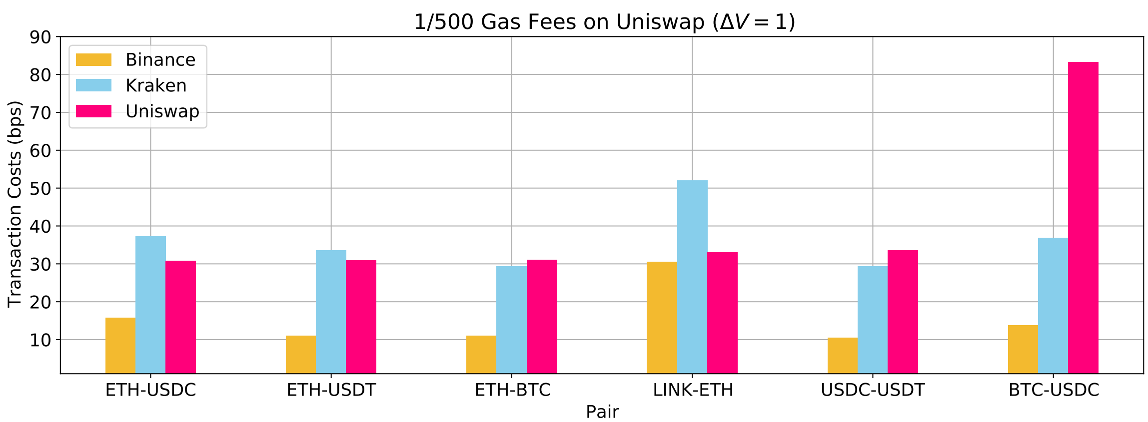


**Figure 12. Transaction costs – Scenario A.** The figure presents transaction costs for a traded amount of 10,000$, averaged over the period from January 2021 to December 2021, for the six pairs in our sample. These are computed as in (4) for the LOB-based Binance and Kraken and on (5) for the AMM-based Uniswap, assuming a hypothetical scenario in which the gas fee required to execute a swap transaction is reduced by a factor of 500. Such a scenario could materialize in 2022, when the proof-of-stake version of Ethereum (Ethereum 2.0) is expected to be deployed. The remaining parameters (liquidity and protocol fees) reflect the empirical values recorded over the period from January 2021 to December 2021.

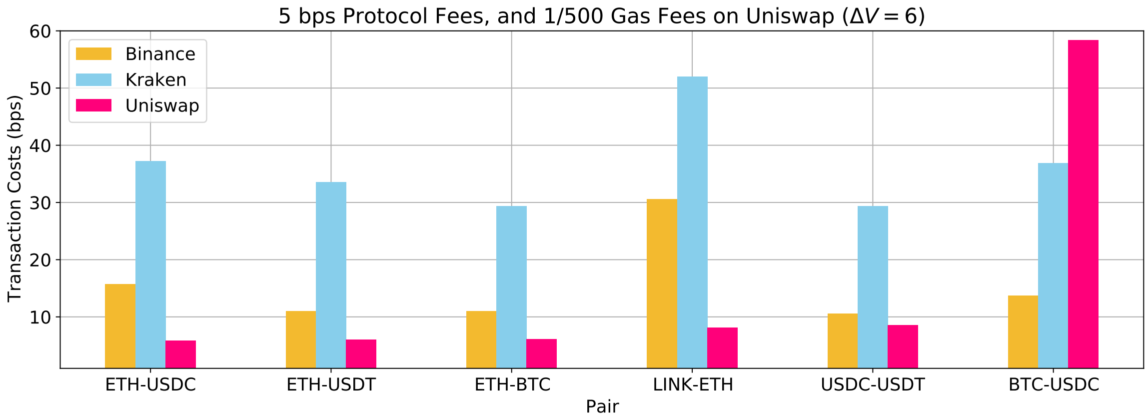


**Figure 13. Transaction costs – Scenario B.** The figure presents transaction costs for a traded amount of 10,000$, averaged over the period from January 2021 to December 2021, for the six pairs in our sample. These are computed as in (4) for the LOB-based Binance and Kraken and on (5) for the AMM-based Uniswap, assuming a hypothetical scenario in which the gas fee required to execute a swap transaction is reduced by a factor of 500 and the protocol fees of Uniswap are reduced to 5 bps. According to our equilibrium model, such a scenario could materialize if trading volume increases six-fold relative to the levels recorded in January and December 2021. Further, the low level of gas fees could become possible after the proof-of-stake version of Ethereum (Ethereum 2.0) is deployed, which is expected to occur in 2022. Protocol fees are assumed to be unchanged and reflect the empirical values recorded over the period from January 2021 to December 2021.

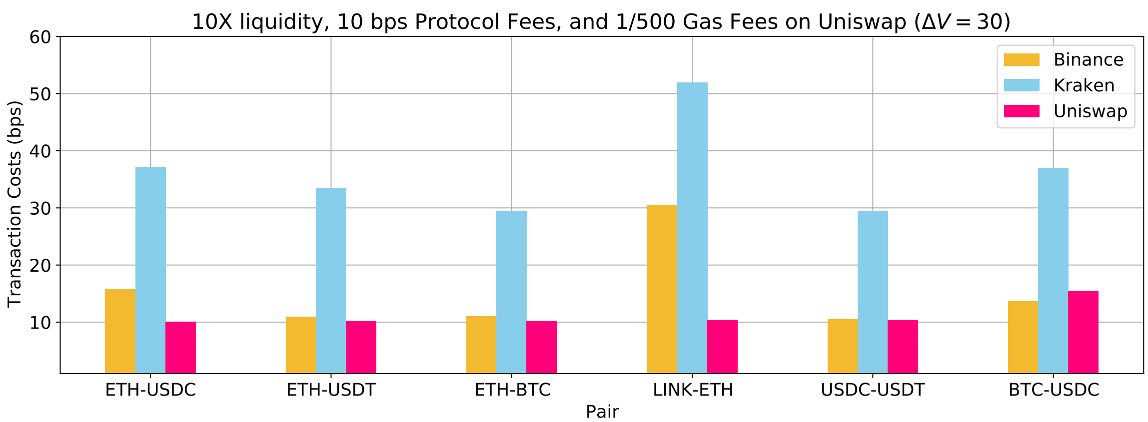


**Figure 14. Transaction costs – Scenario C.** The figure presents transaction costs for a traded amount of 10,000$, averaged over the period from January 2021 to December 2021, for the six pairs in our sample. These are computed as in (4) for the LOB-based Binance and Kraken and on (5) for the AMM-based Uniswap, assuming a hypothetical scenario in which the gas fee required to execute a swap transaction is reduced by a factor of 500, reserves staked in Uniswap's liquidity pools enjoy a 10-fold increase, and protocol fees are reduced to 10 bps. According to our equilibrium model, such a scenario could materialize if trading volume increases 30-fold relative to the levels recorded in January and December 2021. Further, the low level of gas fees could become possible after the proof-of-stake version of Ethereum (Ethereum 2.0) is deployed, which is expected to occur in 2022. Protocol fees are assumed to be unchanged, and reflect the empirical values recorded over the period from January 2021 to December 2021.

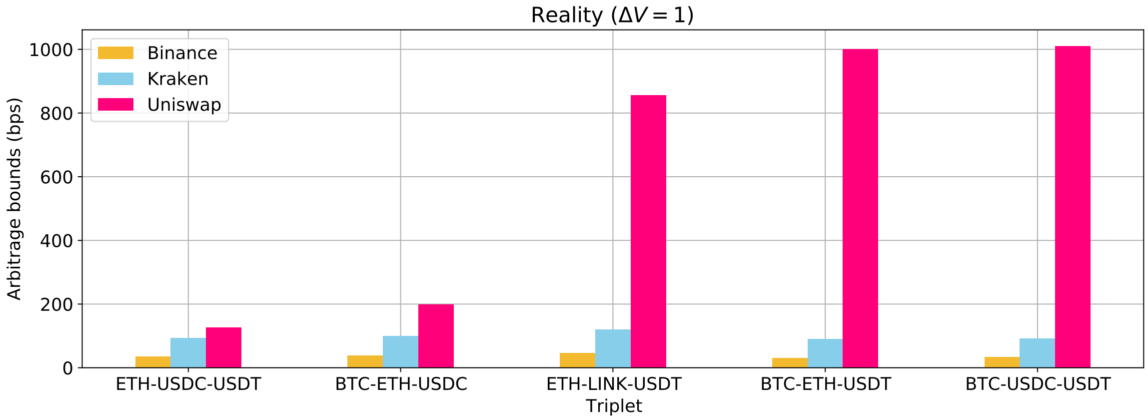


**Figure 15. Price inefficiency – Reality.** The figure presents price inefficiency levels proxied by the size of arbitrage bounds and computed as in (13). These are based on (10) for the LOB-based Binance and Kraken and on (12) for the AMM-based Uniswap. They are estimated at an hourly frequency for the five triplets in our sample, then averaged over the period from January 2021 to December 2021 and expressed in bps.

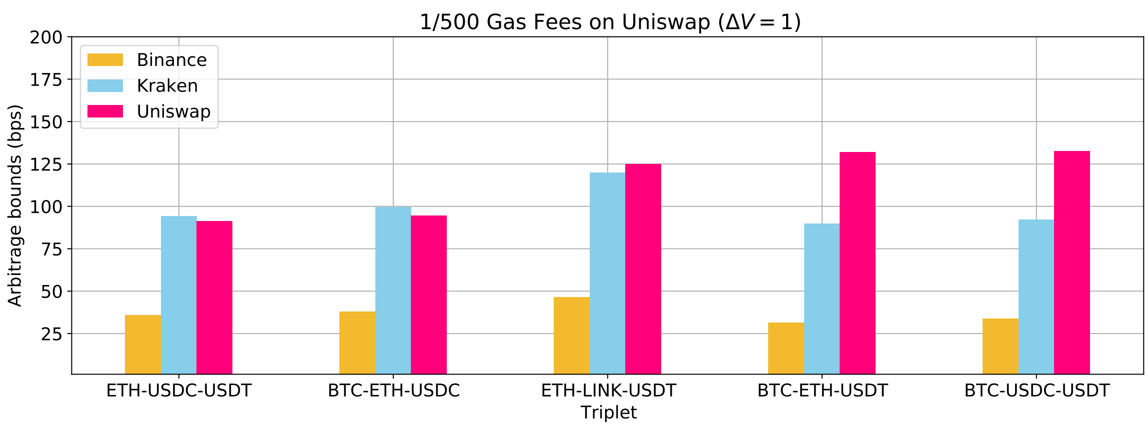


**Figure 16. Price inefficiency – Scenario A.** The figure presents price inefficiency levels averaged over the period from January 2021 to December 2021 for the five triplets in our sample. These are estimated as in (10) for the LOB-based Binance and Kraken and on (12) for the AMM-based Uniswap, assuming a hypothetical scenario in which the gas fee required to execute a swap transaction is reduced by a factor of 500. Such a scenario could materialize in 2022, when the proof-of-stake version of Ethereum (Ethereum 2.0) is expected to be deployed. The remaining parameters (liquidity and protocol fees) reflect the empirical values recorded over the period from January 2021 to December 2021.

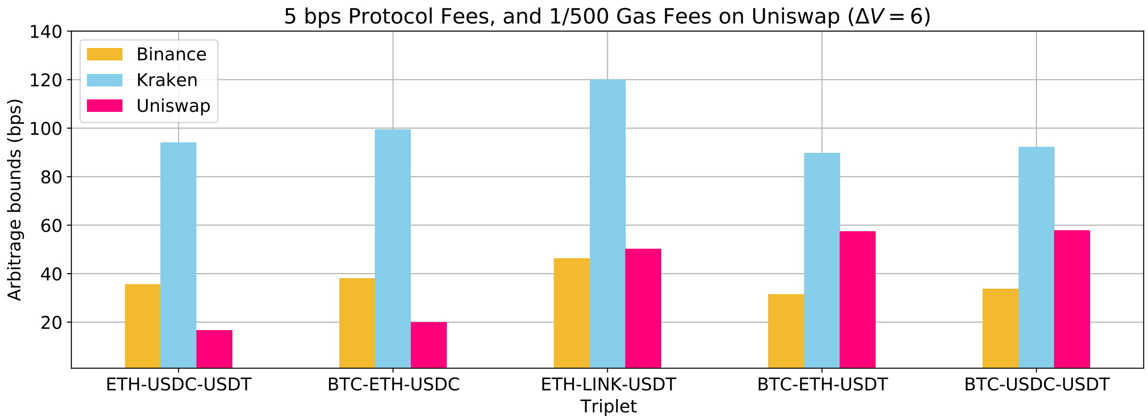


**Figure 17. Price Inefficiency – Scenario B.** The figure presents price inefficiency levels averaged over the period from January 2021 to December 2021 for the five triplets in our sample. These are estimated as in (10) for the LOB-based Binance and Kraken and on (12) for the AMM-based Uniswap, assuming a hypothetical scenario in which the gas fee required to execute a swap transaction is reduced by a factor of 500 and the protocol fees of Uniswap are reduced to 5 bps. According to our equilibrium model, such a scenario could materialize if trading volume increases six-fold relative to the levels recorded in January and December 2021. Further, the low level of gas fees could become possible in 2022, when the proof-of-stake version of Ethereum (Ethereum 2.0) is expected to be deployed. Protocol fees are assumed to be unchanged and reflect the empirical values recorded over the period from January 2021 to December 2021.

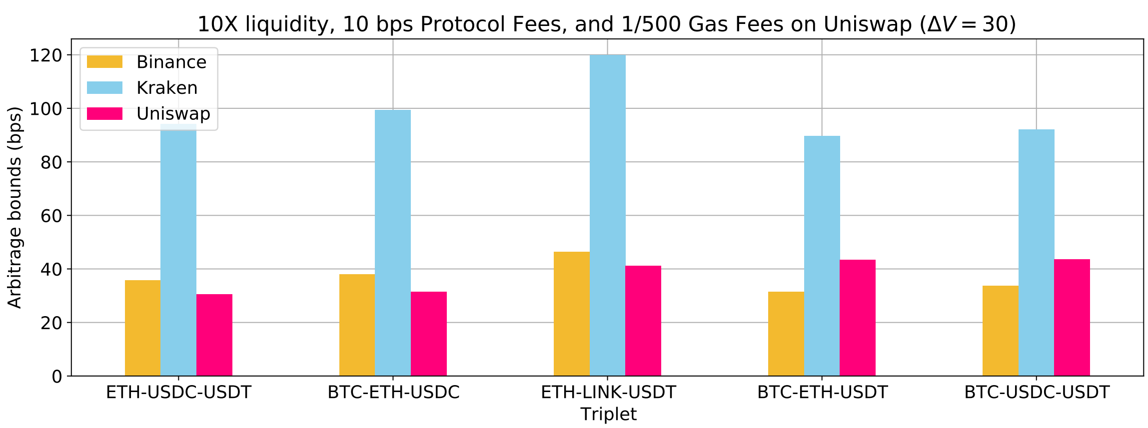


**Figure 18. Price inefficiency – Scenario C.** The figure presents price inefficiency levels averaged over the period from January 2021 to December 2021 for the five triplets in our sample. These are estimated as in (10) for the LOB-based Binance and Kraken and on (12) for the AMM-based Uniswap, assuming a hypothetical scenario in which the gas fee required to execute a swap transaction is reduced by a factor of 500, reserves staked in Uniswap's liquidity pools enjoy a 10-fold increase, and protocol fees are reduced to 10 bps. According to our equilibrium model, such a scenario could materialize if trading volume increases 30-fold relative to the levels recorded in January and December 2021. Further, the low level of gas fees could become possible after the proof-of-stake version of Ethereum (Ethereum 2.0) is deployed, which is expected to occur in 2022. Protocol fees are assumed to be unchanged and reflect the empirical values recorded over the period from January 2021 to December 2021.

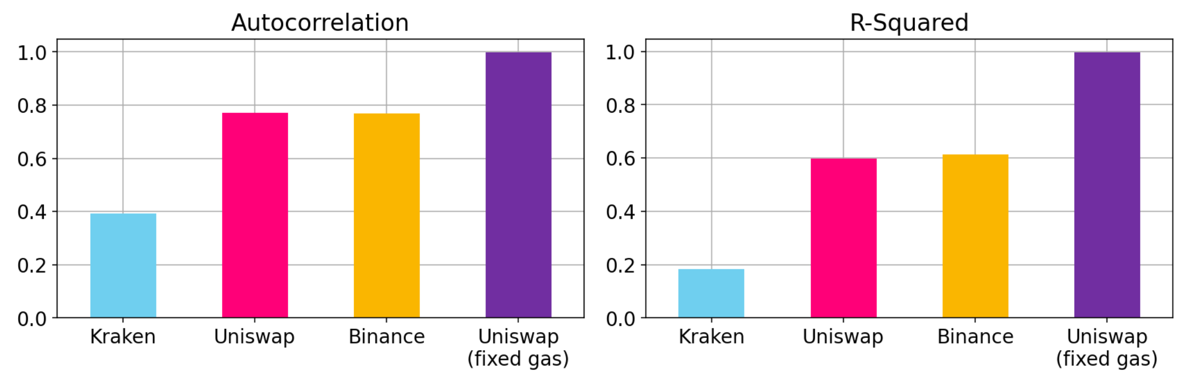


**Figure 19. Predictability of transaction costs.** The figure presents results on the degree of predictability of transaction costs on different exchanges. For each exchange-pair couple, considering hourly transaction costs for a 10,000$ transaction, we compute the auto-correlation coefficient $\rho$ 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 $\rho$ (left panel) and the average $R^2$ from the time-series regressions (right panel).