On The Quality Of Cryptocurrency Markets

Centralized Versus Decentralized Exchanges

Andrea Barbon, Angelo Ranaldo*

First Draft: July 2021. This Draft: December 2021

Abstract

Despite the growing adoption of decentralized exchanges, not much is yet known about their market quality. To shed light on this issue, we compare decentralized blockchain-based venues (DEX) to centralized crypto exchanges (CEX) by assessing two key aspects of market quality: price efficiency and market liquidity. Using a novel and comprehensive data set, we find that overall CEX provide better market quality but DEX become competitive for transactions exceeding $ 100,000. Further, the main determinant of the lower price-efficiency of DEX is the high gas price stemming from proof-of-work blockchains. We propose and empirically validate a stylized theory of DEX liquidity provision, which links trading volumes, protocol fees, and liquidity in equilibrium. Our model identifies quantitative conditions for DEX to overtake CEX 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. The authors would like to thank Charles Milliet (charles.milliet@student.unisg.ch) for his contribution to the work as research assistant.

I. Introduction

In modern financial markets the majority of asset classes, including equity securities and cryptocurrencies, are traded on centralized exchanges (CEX). 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 (DEX) have recently emerged as an alternative and innovative market structure for crypto assets. These venues are based on smart-contract implementations of Automated Market Making (AMM) and attracted an increasing amount of attention and trading volume. One of the questions arising from this development concerns the market quality offered by these new market systems compared to CEX.

We address this issue by analyzing cryptocurrency trading and assessing the market quality of CEX and DEX for a comprehensive set of cryptocurrencies. Specifically, we examine two key aspects of market quality: price efficiency and market liquidity. We find that CEX generally enjoy higher liquidity and tighter transaction costs. However, DEX become competitive in terms of transaction costs for large-sized trades, that is, for amounts above $100,000$\$. By analyzing the violation of the law of one price for triplets of exchange pairs, we also find that DEX enjoy more efficient pricing. We identify the high level of gas fees, that is, the cost of recording transactions on the blockchain and acting as fixed cost, as the main cause of price inefficiency of DEX markets. To draw the possible paths to efficiency, we provide a simple equilibrium model capturing the main characteristics of DEX markets and the risk-return trade-off endured by liquidity providers. Using a unique and representative data set, we test the main empirical predictions and quantify the necessary conditions to improve DEX market quality in order to compete with 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 interesting for at least two reasons. First, it is a new market design which could potentially be applied to more traditional financial securities, given their advantages with respect to CEX. For instance, the fact that DEX relies on AMM rather that 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. Also, the custody of assets is fully kept by 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 DEX is desirable.

Our empirical work leverages on a unique dataset, comprising (i) LOB snapshots for the most liquid centralized crypto exchanges (Binance, Kraken, Coinbase), (ii) liquidity pools levels and transactions for the most prominent DEX (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 re-construct quoted prices and spreads for a selected set of exchange pairs of cryptocurrencies, at the minute frequency. We proceed in three steps. First, analyze market liquidity by computing effective transaction costs for each pair, for different sizes, that an end-user has to bear. For CEX our proxy 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 DEX 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 arise: (i) in general CEX offer better liquidity and lower transaction costs, and (ii) DEX become competitive for expensive transactions (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. Using the above-defined proxy for transaction costs we compute arbitrage bounds, i.e. regions on which the profits from a triangular trade would be lower than effective transaction costs of executing it. We employ the size of those bounds as an inverse proxy for price inefficiency. Our empirical analysis of exchange triplets documents that, once accounting for transaction costs, no-arbitrage conditions are less restrictive for DEX and result into larger deviations from the theoretically efficient price.

In the third part of the paper we outline a simple theoretical model that consistently capture the trade-off faced by DEX liquidity providers (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. And 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 the pool size. On the other side of the trade-off, LPs face the risk of incurring in the so-called *impermanent loss*, the AMM analogue to adverse selection is LOB markets. Hence the prediction of our model, stating that the equilibrium level of assets staked in a liquidity is proportional to the expected volume and inversely proportional to the ex-ante impermanent loss volatility. Based on the insights from our model, we make quantitative predictions on the possible future evolution of liquidity and price efficiency of DEX, conditional on expected levels of trading volume. Our analysis suggests that DEX are likely to catch up with and potentially overtake CEX shortly.

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 CEX. Second, we propose a simple model of liquidity provision in DEX and show that it can explain the vast majority of the cross-sectional and time-series variation in observed liquidity levels. Third, using evidences based on both theory and empirics, we argue that DEX 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 for future efficiency levels. Finally, Section VII concludes.

II. AMM Markets

A. High-level description of AMM Markets

Most exchanges in the current financial markets use a central Limit Order Book (LOB) system that often required a central institution that keeps a record of available buy and sell orders. On this type of exchanges, the market price is determined by the most recently matched buy and sell order. Unlike order-book-based exchanges, Automated Market Makers (AMMs) are based on an algorithm that determines automatically transaction and market prices based on the liquidity made available by market participants.

Implementing a 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 conception has a limited storage capacity which is very needed in order-based exchanges. Crypto exchanges, such as Binance, Coinbase, or Kraken, that still provide a LOB mechanism have to operate it off-chain and thus are centralized entities. This comes at the expense of the benefits offered by decentralized networks. Unlike centralized exchanges (CEX), 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 so-called liquidity providers (LP) 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 also allow users to directly swap assets without having to interact with a counter-party or third party. To incentivize users to provide funds to the liquidity pools, LPs are, in turn, compensated by a small fee charged on each transaction. Nonetheless, providing liquidity is not riskless. Price divergence between the time of provision and withdrawal leads to an Impermanent Loss (IL) (also called divergence loss). This arises from the fact that the liquidity provider 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, without accounting for the revenues from transaction fees.

The importance of decentralized exchanges (DEX) has kept rising since their inception. As of early May 2021, there were more than 23 billion US Dollars deposited in liquidity pools across Uniswap, Sushiswap and Pancakeswap combined. The following Section will discuss the advantages and inconveniences that DEXs can offer compared to centralized order-book-based exchanges.

B. Salient Features of DEX

Decentralized exchanges based on AMM provide their users with a fundamentally different experience relative to standard centralized exchanges based on limit order book (LOB). Below, we discuss a number of relevant advantages and drawbacks of DEX.

First of all, contrary to CEX, the custody of assets is fully kept by the user as there is no third party required to execute the trade. This feature implies that users can take full advantage of the censorship-resistant and trust-less 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 that are commonly associated with the deposit and withdrawal of assets in centralized exchanges.

Second, 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 nor in developing complex algorithms. On the contrary, in LOB-based exchanges, liquidity providers are usually highly specialized and entry costs are significant, both in terms of sophistication and capital. Market makers need fast computers and state-of-the-art algorithms to update their quotes as quickly as possible and avoiding being picked-off by high-frequency traders (Foucault et al., 2017).

Third, platform fees charged to each transaction are distributed to liquidity providers proportional to their share (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 liquidity providers.

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 the approval procedures of CEX may require significant time. Moreover, DEX may allow trading on tokens that are not available in centralized exchanges. On the one hand, this constitutes an advantage, enlarging the space of investment opportunities and improving diversification. On the other hand, this has the drawback to expose users to potentially malicious assets.

Fifth, since transactions are processed by smart contracts and directly recorded on the blockchain, users bear the cost of non-trivial gas fees, required to compensate miners. This fact implies that transactions are subject to an execution delay, whose duration depends on the speed of the underlying blockchain, the chosen gas price, and the level of congestion of the network.

C. Related Literature

We contribute to the 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 which 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 is 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 given periods such as the early hours of a day and 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), there have been a few more studies using 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 (CEX) and decentralized (DEX) crypto exchanges based on innovative blockchain-based venues called automated market making" (AMM). So far, only a few papers have studied AMM exchanges. On the theoretical side, Aoyagi and Ito (2021) study the conditions for the coexistence of such CEX and DEX exchanges and Evans (2020) outlines the pay-offs of liquidity providers 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 of Uniswap prices to closely refer to the reference market price. Capponi and Jia (2021) model the impact on utility for liquidity providers and traders of the curvature of the pricing function on Uniswap. The closest paper to our study is Lehar and Parlour (2021) comparing AMM and a limit order markets. They theorize the liquidity provision in AMM market venues and empirically investigate tokens listed on both Uniswap and Binance. Our contribution to the literature is threefold: 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 CEX and DEX crypto exchanges. By doing so, we quantify under which conditions and transaction size DEX exchanges become competitive compared to CEX ones. 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 so-called *constant product rule* which enables to algebraically determine the market price and transactions price based on the available reserves (Adams et al., 2021). Let's consider a liquidity pool that contains x tokens of

X and y tokens of Y. The 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 with f the protocol fees charged by the decentralized exchange ($f = 0.003$ 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 hold constant before and after any transactions. Hence, 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

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

Solving for Δy 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 it 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 decentralized exchanges are based on smart contracts deployed on blockchains, records of every single interaction with those contracts is available to the public. This rich data-set includes as primitives the creation of exchange pairs, the addition/removal of liquidity from liquidity providers, and swap transactions between two quoted tokens. Building on those, one can re-construct liquidity levels, quoted prices, transaction prices, and trading volume at the pair level at any point in time. We leverage on the API of TheGraph.com to obtain data for Uniswap from the Ethereum MainNet blockchain. We download data on liquidity pool reserves and volumes at hourly frequency for the pairs made of the 5 crypto-tokens subject of our analysis.

For centralized exchanges, on the contrary, data is proprietary. We obtain minute-frequency OHLCV data and full LOB snapshots from Kaiko for all pairs quoted in 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 significant 10-fold increase from around 100 million USD at the start of the sample (August 2020) to roughly 1 billion at the end of the sample (September 2021). A similar upward trend is 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, and both present a significant upward trend. Figure 3 displays trading volumes for both the AMM-based and LOB-based exchanges in our sample, averaged across the three exchanges of each category. The average volume of DEX 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 CEX and DEX in our sample. Even though the increase in DEX is significantly steeper, the wedge within the two categories is still around one order of magnitude at the end of our sample.

Table II presents the daily average trading volume in million USD over the period January-September 2021, which will be the focus of our market quality analysis, for the 8 pairs we consider. These pairs provide a representative sample, as they generate roughly one-third of the volume in 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 values (Foucault et al., 2013). As a proxy for market illiquidity we employ the effective transaction costs associated with a single trade, expressed in 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 with 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 Δx 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[1]. We finally define the transaction costs as the sum of the quoted half-spread and the percentage transaction fees f charged by the respective 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, and required to interact with a smart contract and to record the transaction on the relevant blockchain. The dollar value of those fees depends on the computational complexity of the smart-contract function being utilized, on the execution priority chosen by the trader, and on the prevailing gas price at the execution time. For our purpose we are interested in the gas fees required to execute a swap transaction, that is, invoking the `swapExactTokensForTokens` function of the relevant router contract[2]. 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

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

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 congestion of the network. Figure 4 plots the evolution over time of the gas fees

[1] We only consider sell orders for the sake of simplicity. The half-spread can in principle be quantitatively different for buy orders in a 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.

[2] 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 needs to be used. Nevertheless, the amount of gas required is not significantly different.

(in US Dollars) required to execute a swap transaction in Uniswap. For our empirical analysis, we assume that gas fees for a swap transaction on the Uniswap exchange amount to 32 US Dollars, which equals to the median of the effective gas fees paid by Uniswap users during our sample period. We compute TC_{XY} in AMM and LOB exchanges at the hourly frequency, for the 6 pairs in our sample and for different trade sizes $(10^3, 10^4, 10^5, 10^6)$ expressed in US dollars.

B. Results

Figure 5 displays log levels of transaction costs for the AMM-based exchange Uniswap and the LOB-based exchanges Binance and Kraken, for different trade sizes. The top-left panel shows that Uniswap is extremely expensive for small traded sizes, with transaction costs ranging between 300 and 400 bps for all considered pairs. This finding does not come as a surprise, since gas fees constitute a large percentage (3.2%) of the traded amount. By the same token, the top-right shows that LOB-based exchanges are superior to Uniswap also for mid-sized transactions of $10,000$\$. The situation depicted in the bottom-left panel, for a more significant trade size of $100,000$\$, is somewhat different.

While Binance proves the most convenient choice for 5 out of 6 pairs, Uniswap delivers lower transaction costs for the LINK-ETH pairs. More generally, for this amount, trading costs in Uniswap are not too far from those of centralized exchanges for the 4 pairs involving Ethereum. The bottom-right panel reinforces the finding, showing that Uniswap is competitive with its centralized counterparts for those pairs. 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.

V. Price Efficiency

Finite liquidity and transaction fees constitute frictions limiting arbitrage forces, allowing deviations from efficient prices to persist. 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 θ, 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 θ 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 the 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 we will discuss later. 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,

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

Notice that equation (7) is simply the sum of the three spreads – associated with each transactions of the triangular trade – appropriately *discounted*, that is, adjusted to account for the fact that input amounts of the 2nd and 3rd 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

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

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 limit order book, and the transaction fees charged by the exchange. The total quoted spread and the total fees charged, as a percentage of the initial amount Δx, are defined as in (7) and (8), respectively. Thus, the execution cost of a triangular trade of size Δx 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 Δx is infinitesimal. This choice is based on the assumption that, in absence of fixed transaction costs, arbitrageurs are willing to perform arbitrage trades also with infinitely small dollar amounts, thus minimizing their price impact. Hence, the lower and upper arbitrage bounds for θ are given by

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

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 Δx, are defined as in (7) and (8), respectively. The total gas fees is just the gas fee for a single swap multiplied by a 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 Δx, while the latter increases with Δx. 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 θ 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 comprised between the above defined arbitrage bounds can be thought 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 can 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 θ at the hourly frequency, and find that 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 CEX and DEX for the triple BTC-ETH-USDT are provided in Figures **??**, **??**, and **??**.

D. *Results*

We estimate arbitrage bounds at the hour-frequency for the 5 triplets in our sample, then take the average over the period from January 2021 to September 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 the 5 exchange triplets in our sample, 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 around 200 bps, while for the less liquid ones goes above 10%. These estimates are much higher than those for centralized exchanges, which are lower than 100 bps for almost all the considered triplets. In particular, Binance dominated in terms of price efficiency, with bounds ranging from 30 to 50 basis points.

The first reason for such a significant discrepancy between AMM- and LOB-based exchanges relies on transaction fees. The protocol fees of 30 basis points charged by Uniswap are higher than those charged by centralized exchanges (10 bps for Binance and for 26 bps for Kraken). As triangular arbitrages require three transactions, these wedges become a more significant determinant the net profitability of the profitability of the trade.

The second (and quantitatively most important) cause of such low level of price efficiency enjoyed by Uniswap is related to high level of gas fees, required to compensate miners on the Ethereum blockchain. To compensate 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 compensate for the cost of expensive hardware and a significant energy consumption, thus requiring high gas prices in equilibrium.

On the contrary, no fixed costs are charged on centralized exchanges, since transactions are recorded in their internal databases, rather than on the blockchain. This allows arbitrageurs to exploit triangular arbitrage opportunities trading infinitesimally small amounts. In fact, arbitrage bounds are only slightly larger than transaction fees multiplied by a factor 3, suggesting that trading costs arising from quoted spreads are not as relevant.

VI. Paths to Efficiency and Future Scenarios

A. Equilibrium Liquidity

We model a representative liquidity provider (LP) with risk aversion γ and mean-volatility utility function, who faces the problem to provide the optimal quantity to the exchange pair X/Y. 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 and ΔP denote the total traded volume (in units of X) and the gross percentage change in the quoted price P_{XY}, respectively, between $t = 1$ and $t = 2$. As mentioned before, any price change leads to an impermanent loss (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. Therefore, LPs are compensated by the transactions fees. The resulting impermanent loss is given by (see Appendix A for a mathematical derivation)

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

Let $\mathrm{E}[V]$ denote the expected volume, and $\sigma(IL)$ the expected standard deviation of IL, both known at time $t = 0$. In this setting, the LP faces a risk-return trade-off between the losses arising from IL and the expected profits $f\,\mathrm{E}[V]$ arising from pocketing transaction fees.

The expected return $\mathrm{E}[R]$ from providing an infinitesimal amount of extra liquidity is equal to

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

Hence, assuming ΔP is a Gaussian $\mathcal{N}(0, \sigma^2)$, it is optimal for the LP to provide additional liquidity if and only if

$$x < \frac{f\,\mathrm{E}[V]}{2\gamma\sigma(IL)}\,. \qquad (14)$$

Under rationality assumption, the equilibrium level of liquidity is therefore

$$x = \frac{f\,\mathrm{E}[V]}{2\gamma\sigma(IL)} \tag{15}$$

and we should observe the above relationship linking liquidity, expected volume and IL risk in the data.

We use daily liquidity data to test this hypothesis, proxying for $\mathrm{E}[V]$ with the rolling average of daily traded volume and for $\sigma(IL)$ with the rolling standard deviation of the daily impermanent loss, estimated over the previous two weeks. We regress daily log values of empirically observed liquidity on the ones predicted with (15), for 100 exchange pairs over the period from May 2020 to April 2021. Results are reported in Table V and figure 10, showing a remarkable fit with $R^2 = 72.94\%$.

B. *Future Scenarios*

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

$$\mathrm{E}[V] = \frac{2\,\gamma\,\sigma(IL)x}{f} \;\propto \frac{x}{f} \tag{16}$$

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 protocol), or a combination of both. 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 study potential future scenarios for AMM exchanges, based on assumptions on the time-series dynamics of trading volume. In particular, we focus on scenarios for which the expected increase in trading volume on DEX from 3 to 600 times with respect to the trading volume at the end of our sample. Following our model, a given increase ΔV

in volume gives rise to scenarios with a decrease in fees f, an increase of liquidity x, or a combination thereof. Moreover, we will include three possible scenarios for the dollar value of gas fees g, which is exogenous to the other parameters as it is determined by technological evolution.

C. Results

Table VI presents predicted levels of transaction costs for each scenario. More precisely, hypothetical transactions costs for a $10,000$\$ transaction executed through Uniswap are reported, expressed in basis points. The current situation is represented by the last row, with fees equal to 30 bps, gas cost of 32\$, and unitary liquidity multiplier. The third row from the bottom, assuming a reduction of gas fees to 0.1 USD, shows that transaction costs are roughly halved for most of the pairs. This result quantifies the fact that the currently high level of gas fees represent a significant friction for DEX efficiency. As shown in Figure 12, however, such a reduction in transaction costs would lead to a situation in which Uniswap is still dominated by CEX in terms of trading costs. Our second Scenario (B) is depicted in 13 and assumes, on top of low gas fees, a 6-fold increase in volume leading to a 6-fold reduction in protocol fees to 5 basis points. Under these assumptions Uniswap would be highly competitive with centralized exchanges, offering significantly lower transaction costs with respect to Binance and Kraken for the majority of pairs. 13 presents Scenario C, assuming a more sizeable 30-fold increase in trading volume, resulting into to a 3-fold reduction in protocol fees to 10 basis points 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 suggests that the most important friction undermining DEX arises from high levels of protocol fees, rather than from 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 liquidity providers. Our conclusion applies to the trade size we consider in this analysis ($10,000$\$), while it would likely differ for larger transactions. For those, an increase of the available liquidity could provide more benefits to traders with respect to a reduction in the protocol fees.

Moving to price efficiency, 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 triplets, 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, gas cost of 32$, and trivial liquidity multiplier. The third row from the bottom, assuming a reduction of gas fees to 0.1 USD, shows that price efficiency is significantly increased for all the triplets. Lower gas fees for USDC-USDT-ETH and USDC-BTC-ETH result into a 30% increase in efficiency, while the benefits for the other triplets are higher 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 costs – on the profitability of potential triangular arbitrages. These results highlight the fact that high level of gas fees represent a significant friction for price efficiency on DEX, and 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 CEX. The second Scenario (B) is pictured in 17 and assumes, on top of low gas fees, a 6-fold increase in volume leading to a 6-fold reduction in protocol fees to 5 basis points. 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. 17 presents Scenario C, assuming a more sizeable 30-fold increase in trading volume, resulting into to a 3-fold reduction in protocol fees to 10 basis points and a 10-fold increase in pool liquidity. Under these assumptions, benefits for Uniswap price efficiency are larger for less liquid triplets, while are reduced for the most liquid ones. This suggests that price efficiency on DEX 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. Conclusion

Our analysis of the market quality of cryptocurrency exchanges highlights a number of interesting conclusions on the weaknesses and the future potentials of decentralized venues based on AMM. First of all, the data shows that Uniswap and the other prevalent DEX experienced a steep rise in adoption and trading volume during the last year, accompanied by a significant increase in the available liquidity. Second, we provide evidence showing that DEX are, at present, still not competitive with the largest CEX regarding transaction costs and price efficiency. These two facts can be reconciled by the observation that DEX provide a number of advantages with respect to DEX, in particular 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 venues instead of centralized ones.

We highlight multiple factors determining the degree of market quality of DEX, 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 decentralized exchanges. 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 as they 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 DEX clarifies the risk-return trade-off faced by liquidity providers, based on impermanent loss and expected profits from trading fees. As the former only depends 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 liquidity providers in equilibrium. We analyze a number of future scenarios based on different assumptions for future levels of trading volume routed through DEX, concluding that those could soon become as efficient as CEX under relatively modest increases in volume, provided that gas costs will be reduced 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, DEX based on AMM could soon offer a competitive alternative to centralized venues. 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. *Available at 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.

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

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. 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. *Nuttawat, Bitcoin Liquidity (May 31, 2019)*, 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. *Available at SSRN 3142022*, 2018.

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

Appendix

A. LP Returns and Impermanent Loss

Let's assume a 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 unit of y is

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

At time $t = 1$, the value of her position in 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 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 impermanent loss 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 ρ, notice that market price and effecitve 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 ρ_X

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

As the transaction price can be expressed using the price impact and market price, we rewrite the output amount from the 1st trade as follow

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

Using the previous expression, we can express ρ_Y in function of Δx

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

Similarly for ρ_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 cumlative price impact and the total spread is given by $\rho = \rho_X \rho_Y \rho_Z$ and $S_{XYZ} = 1 - \rho$. Using the previouly 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 size for the 2nd and 3rd trade decrease due to the price impact and previous plaftorm fees charged.

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

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

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 1000 most liquid exchange pairs in our sample. We report the total value of liquidity in US Dollars, 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 US Dollars over the period January-September 2021 for each pair in our sample. The percentage of the aggregate volume represented by these pairs on each exchanges is reported below. For pairs involving USDC on Kraken, we report the volume for the corresponding pair based on USD, since volumes on USDC-based pairs are close to zero.

Pair	Binance	Kraken	Uniswap
BTC-USDC	119.36	257.39	2.17
BTC-USDT	3384.25	17.92	0.04
BTC-ETH	538.05	33.43	36.16
ETH-USDT	2223.72	11.39	114.58
LINK-ETH	9.65	0.78	14.1
USDC-ETH	60.39	202.87	120.9
LINK-USDT	233.28	1.13	0.0
USDC-USDT	156.27	0.0	5.61
Fraction of Total Volume	24.71%	30.45%	29.15%

Table III. Impermanent Loss. The table shows the summary statistics of the daily Impermanent Loss (IL) for 100 pairs traded on Uniswap over the period May 2020 to April 2021. Summary statistics for the IL are available at pair-day, pair and day level.

Level	N	Mean	Std	1%	10%	50%	90%	99%
Pair-Day	18183	-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. Impermanent Loss by Type. The table shows the summary statistics of the daily Impermanent Loss (IL) for 100 pairs traded on Uniswap over the period May 2020 to April 2021. IL displayed are divided into three categories. Category stable-stable refers to a pair of two cryptocurrency that are not pegged to a fiat currency, e.g. ETH-BTC. Category stable-crypto refers to a pair of one stablecoin and one classic cryptocurrency that is not pegged to a fiat currency, e.g. USDT-BTC. Category stable-stable refers to a pair of two stablecoins, e.g. USDT-DAI.

Type	N	Mean	Std	1%	10%	50%	90%	99%
crypto-crypto	13950	-0.00174	0.01102	-0.02345	-0.00281	-0.00019	-0.0	-0.0
stable-crypto	2789	-0.00188	0.02041	-0.01887	-0.00171	-0.00011	-0.0	-0.0
stable-stable	1444	-0.00001	0.00011	-0.00008	-0.00001	-0.00000	-0.0	-0.0

Table V. 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 (15). Both the dependent and independent variables are computed at the pair-day level for 100 exchange pairs over the period from May 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)	0.73***	0.36***	0.72***	0.26***
	(21.36)	(7.53)	(21.78)	(6.43)
Constant	4.32***			
	(9.02)			
Observations	16,947	16,947	16,938	16,938
R-squared	0.73	0.94	0.75	0.96
Pair Fixed Effects	-	Yes	-	Yes
Date Fixed Effects	-	-	Yes	Yes
Ses Clustered By	Pair-Date	Pair-Date	Pair-Date	Pair-Date

Table VI. Uniswap Hypothetical Transaction Costs. The table displays hypothetical transactions costs for a 10,000$ transaction executed through Uniswap, expressed in basis points. They are computed as in (5), first at the hour-frequency and then averaged over the period from January 2021 to September 2021, for different levels of gas fees (*Gas*, in US dollars), platform fees (*Fees*, in basis points), and liquidity multiplier (*Liq*). Each row represents a potential future scenario requiring an increase in trading volume, as predicted by our model, equal to ΔV.

Fees	Liq	Gas	ΔV	BTC ETH	BTC USDC	ETH USDT	LINK ETH	USDC ETH	USDC USDT
5	100	0.10	600	5.11	5.64	5.11	5.13	5.11	5.13
5	100	10.00	600	15.01	15.54	15.01	15.03	15.01	15.03
5	100	36.00	600	41.01	41.54	41.01	41.03	41.01	41.03
5	10	0.10	60	5.20	10.46	5.19	5.40	5.18	5.45
5	10	10.00	60	15.10	20.36	15.09	15.30	15.08	15.35
5	10	36.00	60	41.10	46.36	41.09	41.30	41.08	41.35
5	1	0.10	6	6.13	58.39	6.03	8.14	5.88	8.57
5	1	10.00	6	16.03	68.29	15.93	18.04	15.78	18.47
5	1	36.00	6	42.03	94.29	41.93	44.04	41.78	44.47
10	100	0.10	300	10.11	10.64	10.11	10.13	10.11	10.13
10	100	10.00	300	20.01	20.54	20.01	20.03	20.01	20.03
10	100	36.00	300	46.01	46.54	46.01	46.03	46.01	46.03
10	10	0.10	30	10.20	15.46	10.19	10.40	10.18	10.45
10	10	10.00	30	20.10	25.36	20.09	20.30	20.08	20.35
10	10	36.00	30	46.10	51.36	46.09	46.30	46.08	46.35
10	1	0.10	3	11.13	63.36	11.03	13.14	10.88	13.57
10	1	10.00	3	21.03	73.26	20.93	23.04	20.78	23.47
10	1	36.00	3	47.03	99.26	46.93	49.04	46.78	49.47
30	100	0.10	100	30.11	30.64	30.11	30.13	30.11	30.13
30	100	10.00	100	40.01	40.54	40.01	40.03	40.01	40.03
30	100	36.00	100	66.01	66.54	66.01	66.03	66.01	66.03
30	10	0.10	10	30.20	35.45	30.19	30.40	30.18	30.45
30	10	10.00	10	40.10	45.35	40.09	40.30	40.08	40.35
30	10	36.00	10	66.10	71.35	66.09	66.30	66.08	66.35
30	1	0.10	1	31.13	83.26	31.03	33.13	30.87	33.56
30	1	10.00	1	41.03	93.16	40.93	43.03	40.77	43.46
30	1	36.00	1	67.03	119.16	66.93	69.03	66.77	69.46

Table VII. 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 (12), first at the hour-frequency and then averaged over the period from January 2021 to September 2021, for different levels of gas fees (*Gas*, in US dollars), platform fees (*Fees*, in basis points), and liquidity multiplier (*Liq*). Each row represents a potential future scenario requiring an increase in trading volume, as predicted by our model, equal to ΔV.

Fees	Liq	Gas	ΔV	USDC USDT ETH	USDC BTC ETH	USDT BTC ETH	LINK USDT ETH	USDC USDT BTC
5	100	0.10	600	15.36	15.51	21.60	20.47	21.64
5	100	10.00	600	18.57	20.14	80.93	69.67	81.32
5	100	36.00	600	21.80	24.75	139.85	118.60	140.62
5	10	0.10	60	15.87	16.62	35.87	32.30	35.99
5	10	10.00	60	23.63	31.25	222.53	187.34	223.70
5	10	36.00	60	31.43	45.83	406.37	340.59	408.59
5	1	0.10	6	17.61	20.14	80.93	69.67	81.30
5	1	10.00	6	41.16	66.37	661.65	554.37	665.32
5	1	36.00	6	64.65	112.36	1217.95	1024.40	1224.82
10	100	0.10	300	30.32	30.48	36.57	35.44	36.61
10	100	10.00	300	33.61	35.11	95.85	84.61	96.27
10	100	36.00	300	36.85	39.72	154.74	133.50	155.47
10	10	0.10	30	30.85	31.60	50.84	47.27	50.95
10	10	10.00	30	38.75	46.22	237.35	202.19	238.53
10	10	36.00	30	46.42	60.79	421.05	355.33	423.28
10	1	0.10	3	32.59	35.11	95.85	84.61	96.23
10	1	10.00	3	56.13	81.31	676.14	568.95	679.82
10	1	36.00	3	79.57	127.26	1232.01	1038.61	1238.93
30	100	0.10	100	90.08	90.24	96.31	95.19	96.35
30	100	10.00	100	93.30	94.85	155.42	144.20	155.80
30	100	36.00	100	96.90	99.45	214.12	192.95	214.91
30	10	0.10	10	90.61	91.35	110.53	106.97	110.65
30	10	10.00	10	98.51	105.93	296.49	261.43	297.67
30	10	36.00	10	106.21	120.46	479.64	414.11	481.87
30	1	0.10	1	92.34	94.85	155.42	144.20	155.79
30	1	10.00	1	115.81	140.91	733.95	627.08	737.68
30	1	36.00	1	139.19	186.73	1288.13	1095.32	1295.21

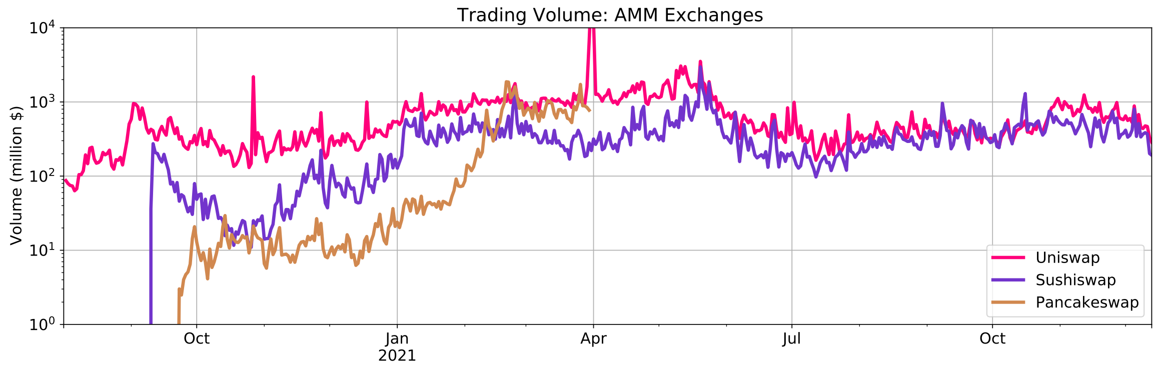

Figure 1. DEX Volume. The figure presents traded volumes for the AMM-based exchanges in our sample, namely, Uniswap, PancakeSwap, and SushiSwap, for the period starting in August 2020 and ending in November 2021. The displayed traded volume is the summation of the volume for all the trading pairs listed on each exchange. The vertical axis uses log-scale and reported in million US Dollars.

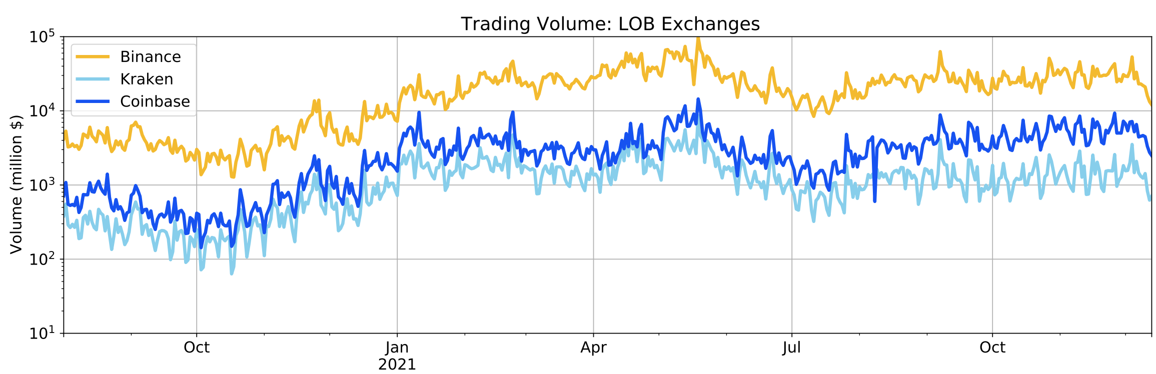

Figure 2. CEX Volume. The figure presents traded volumes for the LOB-based exchanges in our sample, namely, Binance, Kraken, and Coinbase, for the period starting in August 2020 and ending in November 2021. The displayed traded volume is the summation of the volume for all the trading pairs listed on each exchange. The vertical axis uses log-scale and it is reported in million US Dollars.

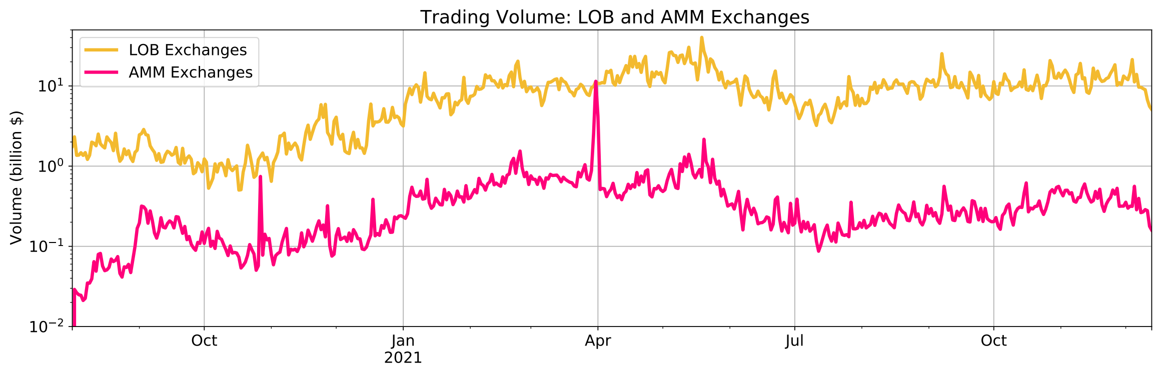

Figure 3. CEX VS DEX Volume. The figure presents traded volumes for both the AMM-based and LOB-based exchanges in our sample, averaged across the three exchanges of each category, for the period starting in August 2020 and ending in November 2021. The displayed traded volume is the summation of the volume for all the trading pairs listed on each exchange. The vertical axis uses log-scale and it is reported in billion US Dollars.

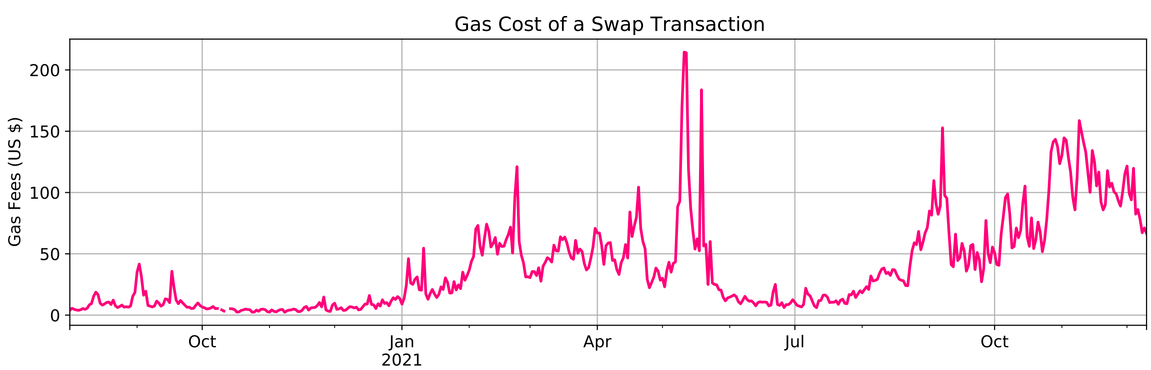

Figure 4. Gas Fees. The figure presents the time-series evolution of the gas cost of a swap transaction in our sample, in US dollars.

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 hour-frequency for the 6 pairs in our sample, for different trade sizes ($10^3, 10^4, 10^5$, and 10^6 US dollars), then averaged over the period from January 2021 to September 2021. The vertical axis is in log-scale and reported in basis points units.

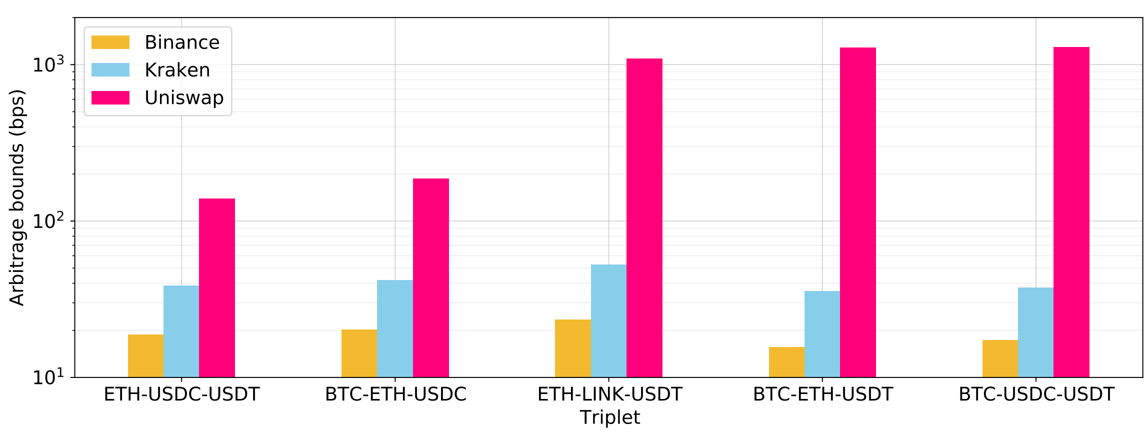

Figure 6. Price Inefficiency. The figure presents price inefficiency levels, proxied by the size of arbitrage bounds computed as in (13). These are based on (10) for the LOB-based Binance and Kraken, and on (12) for the AMM-based Uniswap. They are estimated at the hour-frequency for the 5 triples in our sample, then averaged over the period from January 2021 to September 2021. The vertical axis is in log-scale and reported in basis points units.

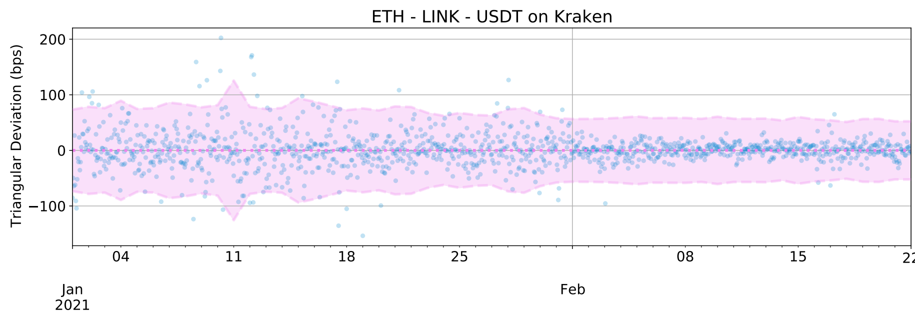

Figure 7. Arbitrage Bounds. The figure presents the arbitrage bounds and θ the deviations from the law of one price for the triple LINK-USDT-ETH on the LOB-based Kraken over the period from January 2021 to September 2021. Arbitage bounds are computed as in (10) on a daily basis using an infinitesimal trade size Δx assuming a transaction fee of 10bps, which is provided to users whose 30-day trading volume is above 10000 BTC (equivalent to around 300 million USD at the time of writing). Triangular price deviations θ based on (6) are computed at the hourly frequency.

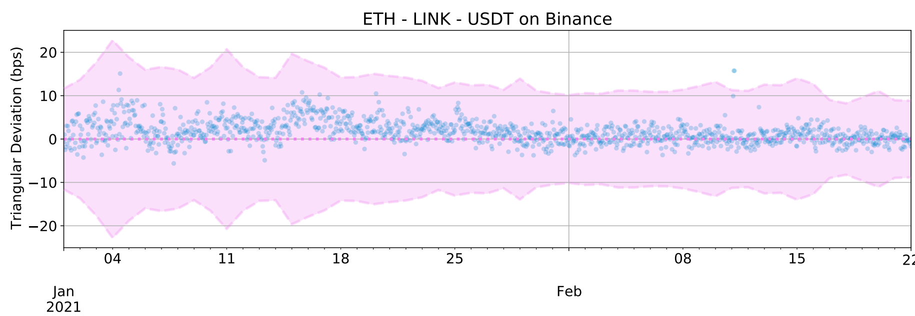

Figure 8. Arbitrage Bounds. The figure presents the arbitrage bounds and θ the deviations from the law of one price for the triple LINK-USDT-ETH on the LOB-based Binance over the period from January 2021 to September 2021. Arbitage bounds are computed as in (10) on a daily basis using an infinitesimal trade size Δx assuming some Binance users are able to trade without incurring the transaction fee (0bps). Triangular price deviations θ based on (6) are computed at the hourly frequency.

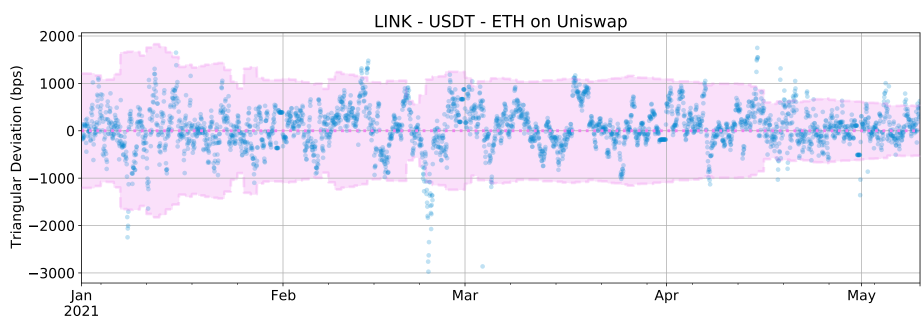

Figure 9. Arbitrage Bounds. The figure presents the arbitrage bounds and θ the deviations from the law of one price for the triple LINK-USDT-ETH on the AMM exchange Uniswap over the period from January 2021 to September 2021. Arbitage bounds are computed as in (12) on a daily basis using an optimal trade size Δx. Triangular price deviations θ based on (6) are computed at the hourly frequency.

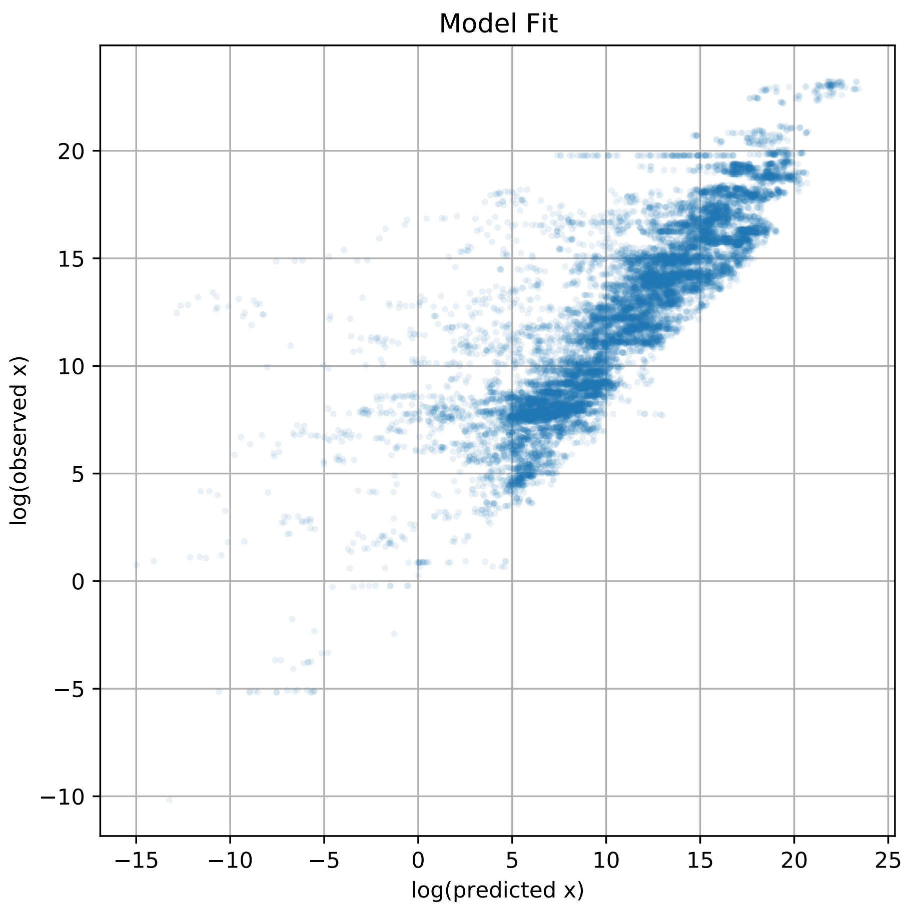

Figure 10. Model Fit. The figure presents a scatter plot of observed levels of liquidity (y-axis) and those predicted by our model and computed as in (15) (x-axis), based on daily observations of 100 exchange pairs quoted in the AMM-based Uniswap, over the period from May 2020 to April 2021.

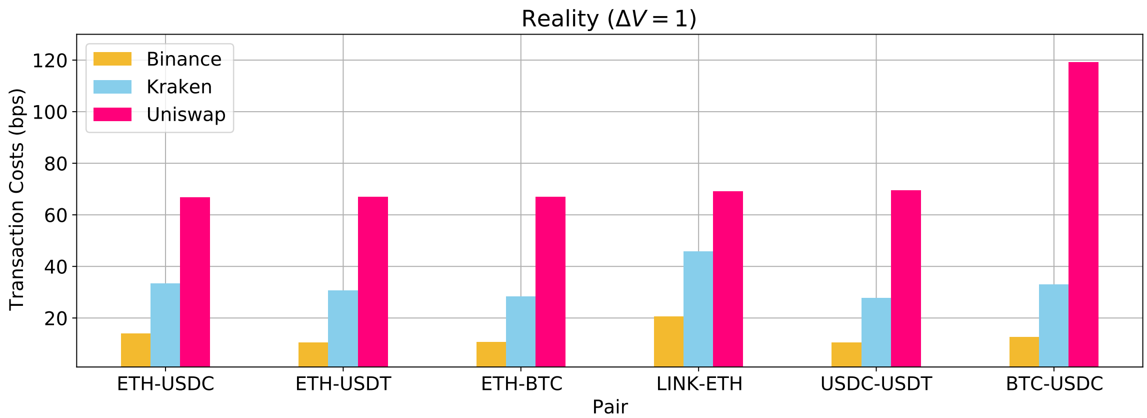

Figure 11. Transaction Costs – Reality. The figure presents transaction costs for a traded amount of 10,000\$, for the 6 pairs in our sample. These are computed as in (4) for the LOB-based Binance and Kraken, and on (5) for the AMM-based Uniswap.

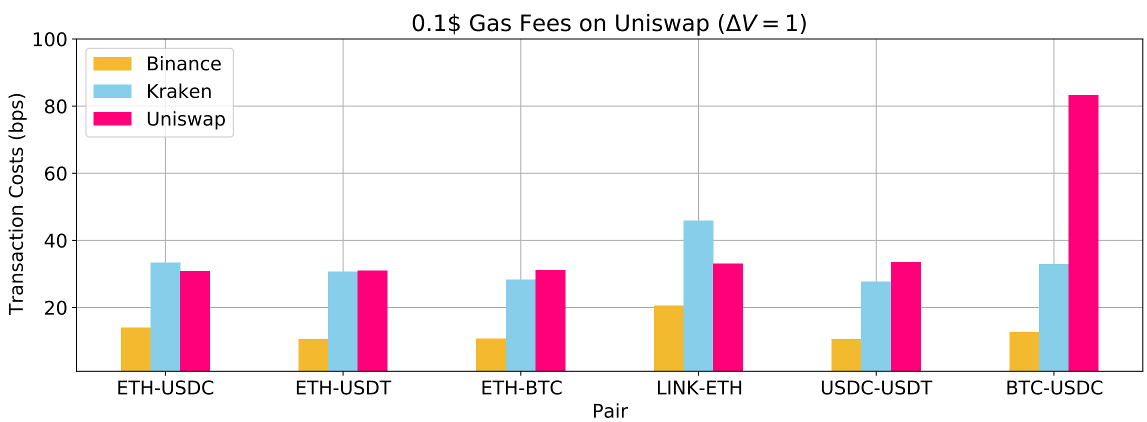

Figure 12. Transaction Costs – Scenario A. The figure presents transaction costs for a traded amount of 10,000\$, averaged over the period from January 2021 to September 2021, for the 6 pairs in our sample. These are computed as in (4) for the LOB-based Binance and Kraken, and on (5) for the AMM-based Uniswap, assuming a hypothetical scenario in which the gas fee required to execute a swap transaction is reduced to 0.1\$. Such a scenario could materialize in early 2022, when the proof-of-stake version of Ethereum (Ethereum 2.0) is expected to be deployed. Remaining parameters (liquidity and protocol fees) reflect the empirical values recorded over the period from January 2021 to September 2021.

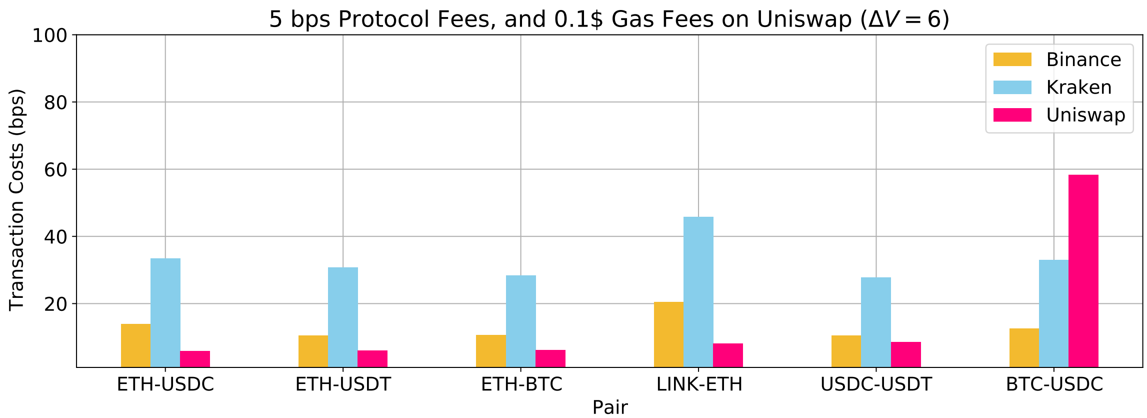

Figure 13. Transaction Costs – Scenario B. The figure presents transaction costs for a traded amount of 10,000$, averaged over the period from January 2021 to September 2021, for the 6 pairs in our sample. These are computed as in (4) for the LOB-based Binance and Kraken, and on (5) for the AMM-based Uniswap, assuming a hypothetical scenario in which the gas fee required to execute a swap transaction is reduced to 0.1$ 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 6-fold relative to levels recorded in January and September 2021. Further, the low level of gas fees could be made possible after the proof-of-stake version of Ethereum (Ethereum 2.0) is expected to be deployed, in early 2022. Protocol fees are assumed to be unchanged, and reflect the empirical values recorded over the period from January 2021 to September 2021.

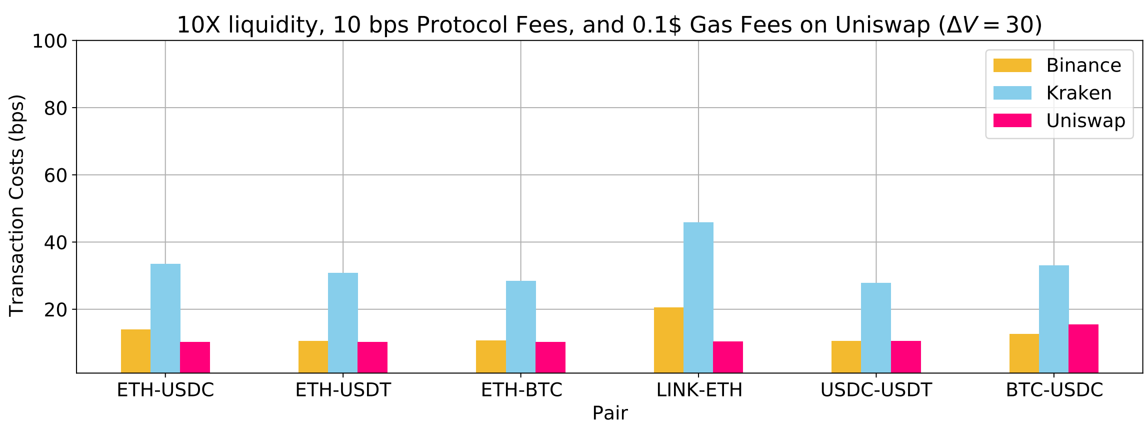

Figure 14. Transaction Costs – Scenario C. The figure presents transaction costs for a traded amount of 10,000$, averaged over the period from January 2021 to September 2021, for the 6 pairs in our sample. These are computed as in (4) for the LOB-based Binance and Kraken, and on (5) for the AMM-based Uniswap, assuming a hypothetical scenario in which the gas fee required to execute a swap transaction is reduced to 0.1$, reserves staked in Uniswap's liquidity pools enjoy a 10-fold increase, and the 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 levels recorded in January and September 2021. Further, the low level of gas fees could be made possible after the proof-of-stake version of Ethereum (Ethereum 2.0) is expected to be deployed, in early 2022. Protocol fees are assumed to be unchanged, and reflect the empirical values recorded over the period from January 2021 to September 2021.

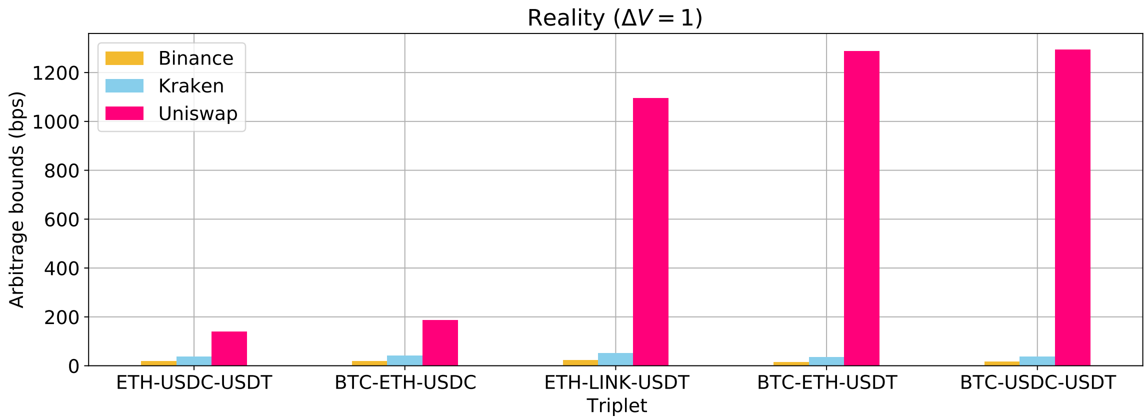

Figure 15. Price Inefficiency – Reality. The figure presents price inefficiency levels, proxied by the size of arbitrage bounds computed as in (13). These are based on (10) for the LOB-based Binance and Kraken, and on (12) for the AMM-based Uniswap. They are estimated at the hour-frequency for the 5 triples in our sample, then averaged over the period from January 2021 to September 2021, and expressed in basis points units.

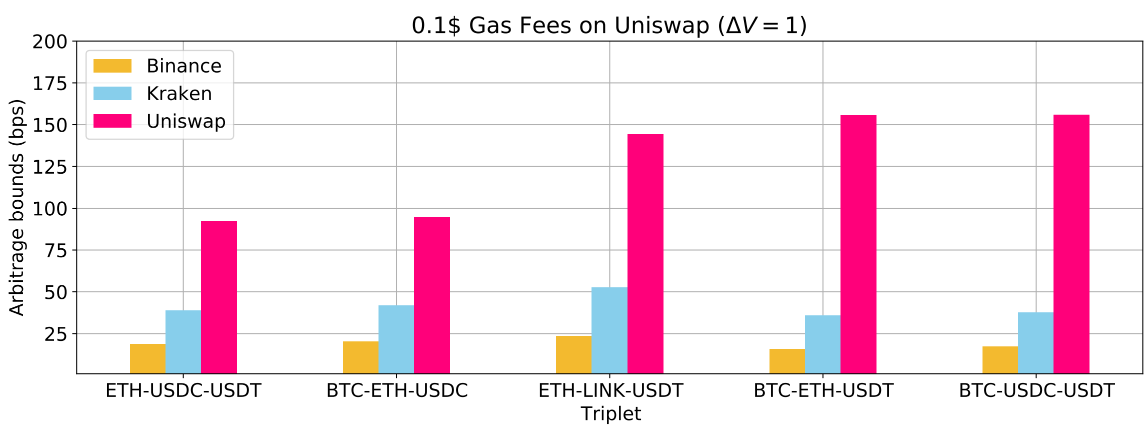

Figure 16. Price Inefficiency – Scenario A. The figure presents price inefficiency levels, averaged over the period from January 2021 to September 2021, for the 5 triples in our sample. These are estimated as in (10) for the LOB-based Binance and Kraken, and on (12) for the AMM-based Uniswap, assuming a hypothetical scenario in which the gas fee required to execute a swap transaction is reduced to 0.1$. Such a scenario could materialize in early 2022, when the proof-of-stake version of Ethereum (Ethereum 2.0) is expected to be deployed. Remaining parameters (liquidity and protocol fees) reflect the empirical values recorded over the period from January 2021 to September 2021.

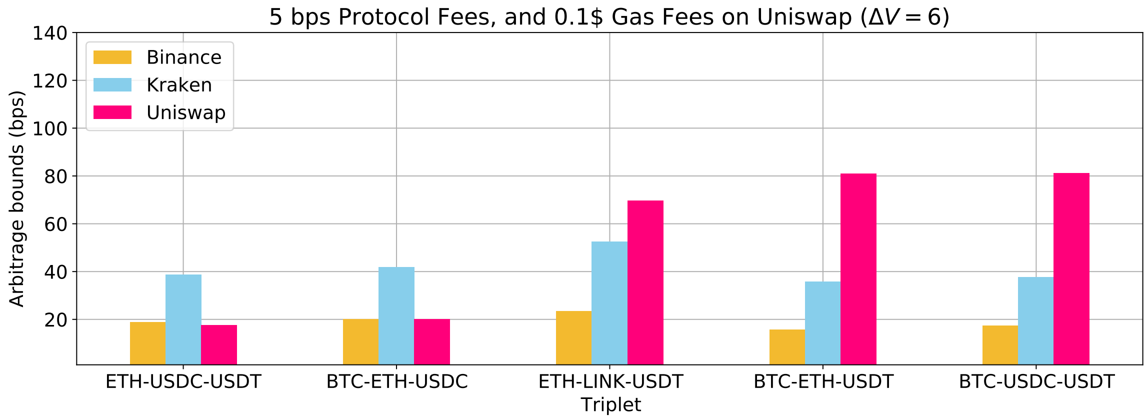

Figure 17. Price Inefficiency – Scenario B. The figure presents price inefficiency levels, averaged over the period from January 2021 to September 2021, for the 5 triples in our sample. These are estimated as in (10) for the LOB-based Binance and Kraken, and on (12) for the AMM-based Uniswap, assuming a hypothetical scenario in which the gas fee required to execute a swap transaction is reduced to 0.1$ 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 6-fold relative to levels recorded in January and September 2021. Further, the low level of gas fees could be made possible after the proof-of-stake version of Ethereum (Ethereum 2.0) is expected to be deployed, in early 2022. Protocol fees are assumed to be unchanged, and reflect the empirical values recorded over the period from January 2021 to September 2021.

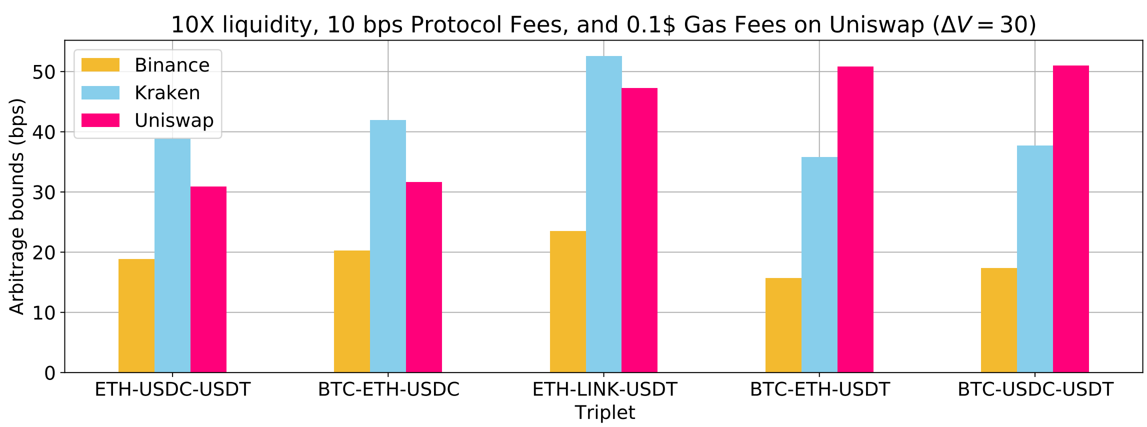

Figure 18. Price Inefficiency – Scenario C. The figure presents price inefficiency levels, averaged over the period from January 2021 to September 2021, for the 5 triples in our sample. These are estimated as in (10) for the LOB-based Binance and Kraken, and on (12) for the AMM-based Uniswap, assuming a hypothetical scenario in which the gas fee required to execute a swap transaction is reduced to 0.1$, reserves staked in Uniswap's liquidity pools enjoy a 10-fold increase, and the 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 levels recorded in January and September 2021. Further, the low level of gas fees could be made possible after the proof-of-stake version of Ethereum (Ethereum 2.0) is expected to be deployed, in early 2022. Protocol fees are assumed to be unchanged, and reflect the empirical values recorded over the period from January 2021 to September 2021.