Decentralized Exchanges

Alfred Lehar*
Haskayne School of Business
University of Calgary

Christine A. Parlour†
Haas School of Business
UC Berkeley

January 12, 2022

Abstract

Keywords: Decentralized Exchanges, Blockchain, Decentralized Finance

*Corresponding author, email: alfred.lehar@haskayne.ucalgary.ca, Tel: (403) 220 4567. Alfred Lehar is grateful to the to the Canadian Securities Institute Research Foundation and the Fintech Dauphine Chair in partnership with Mazars and Crédit Agricole CIB for financial support. Christine Parlour is grateful for a Norwegian Finance Initiative Grant.

†email:parlour@berkeley.edu. We thank Andrea Barbon, Agostino Capponi, Dmitry Chebotarev, Arthur Gervais, Anand Goel, Terrence Hendershott, Jaisun Li, William Mann, Katya Malinova, Andreas Park, Daniel Rabetti, Gustavo Schwenkler, Fahad Saleh, Donghwa Shin, Andre Stenzel, and Haoxiang Zhu, for helpful comments, and seminar participants at the BIS, Insead, Penn State, Rochester, University of Copenhagen, University of Florida, 2021 Central Bank Conference on the Microstructure of Financial Markets, 2021 Central Bank Research Association conference, 6th RiskLab Bank of Finland, 2021 China International Conference in Finance, 6th Southern Denmark University Finance Workshop, 2021, NFA 2021, NBIM-Columbia Research Conference, UWA BlockChain and Crypto conference, FRA 2021, AFA 2022, Blockchain Technology Symposium, 2021 SNB-CIF Conference on Cryptoassets and Financial Innovation, 2021 FinteQC conference, 2021 Conference on Financial Innovation. Tokenomics 2021, MaCCI Mannheim 2021, 2021 SAFE Microstructure conference, 5th Shanghai-Edinburgh-London Fintech Conference 2021, 2021 Global AI Finance Research Conference, 2021 Coller School of Management Blockchain Research Institute Conference, and Protocol Labs

Decentralized Exchanges

Abstract

Uniswap is a system of smart contracts on the Ethereum blockchain and is one of the largest decentralized exchanges with a liquidity balance worth over 8 billion USD and daily trading volume of about 2 billion USD. It is a new model of liquidity provision, so called automated market making. We collect and analyze data on all 19 million Uniswap interactions from its inception to May 2021. For this new market, we characterize equilibrium liquidity pools and provide evidence that they are stable. We compare this automated market maker to a centralized limit order book (Binance) and document absence of long-lived arbitrage opportunities and show conditions under which the AMM dominates a limit order market.

1 Introduction

Uniswap is a decentralized exchange that launched in November 2018. To date, committed liquidity supply tops 9 billion USD across multiple cryptocurrencies. This liquidity facilitates transactions worth roughly 2 billion USD per day. One of the striking features of this successful exchange is that instead of a centralized limit order book, it uses a novel model of liquidity provision. In this paper, we provide a detailed empirical analysis of UniSwap and analyze the way in which "automated market making" provides liquidity, and what this new protocol informs us about centralized order limit order markets.

In an automated market maker (AMM) such as Uniswap, each asset pair comprises a distinct pool or market. Agents supply liquidity by adding the pair in proportion to the existing pool. Agents demand liquidity by adding one asset and removing the other. The relative proportion of the two traded assets determines the average price paid and is calculated according to a predetermined downward sloping, convex relationship. This is referred to as a bonding curve. The convexity implies that larger orders have a larger price impact. In addition, all liquidity demanders pay a proportional fee to the liquidity suppliers.

To analyze this new form of liquidity provision, we collected a novel, detailed data set of 43,349,198 interactions with the Uniswap smart contract. Our observations are at block frequency, and include the ordering sequence within each block. Using these data we can identify all flows into and out of 36,958 liquidity pools – that is changes in liquidity supply as well as all the token trades. Thus, we can trace how liquidity is both supplied and demanded for each set of asset pairs and the payoffs to participants.

To structure our analysis, we focus on two key differences between an AMM and a limit order market. First, in the AMM, the benefits and costs of supplying liquidity are mutualized: Liquidity suppliers are not in competition. In contrast, in the limit order book, strategic liquidity suppliers actively compete with each other – the costs and benefits of supplying liquidity are individual to each liquidity supplier. Second, in the AMM, price impact is deterministic. In particular, the transaction price is determined by the bonding curve and is perfectly predictable given the size of the liquidity pool and incoming order. By contrast, in the limit order market, liquidity suppliers choose to post at a price that maximizes their profits and may even withdraw liquidity strategically.

We first investigate the effect of these two key differences in liquidity provision for an asset whose value is volatile. Risk neutral liquidity suppliers, a liquidity demander and an arbitrageur all interact. Liquidity suppliers may be adversely selected as liquidity is posted before any potential asset innovation. We compare strategic liquidity provision in a stylized limit order market and an AMM.

In the stylized limit order book market, competing liquidity suppliers post prices to trade off adverse selection risk against profitable liquidity supply. Of course, if two liquidity suppliers are competing in the same market, each earns zero in expectation. Because of this, each liquidity supplier has an incentive to invest to find trading opportunities in which he does not have to compete. This captures the idea that liquidity is cheap if two suppliers are competing on price. However, if they compete on other dimensions such as speed, it may increase the cost of liquidity.

In the AMM, a liquidity supplier deposits a pair of asset in a pool in exchange for a liquidity token. On redemption, the liquidity token yields both assets but in their current proportion. The direct payoff to supplying liquidity is in the form of fees, and the indirect cost of supplying liquidity is the loss of committed capital if an arbitrageur has traded against the pool. In this case, if the relative price of one of the asset shifts, an arbitrageur would find it profitable to buy the underpriced asset. The arbitrageur effectively rebalances the liquidity suppliers' portfolio at disadvantageous terms. The equilibrium size of a pool balances the fee revenue against this "picking off" risk. Equilibrium is reached through a change in the size of a pool rather than a change in price because larger pools mechanically have a smaller price impact.

Because of the different market designs, the gains from trade are split differently across the market participants. In the limit order market liquidity suppliers retain all the price impact revenue from supplying liquidity, whereas in the AMM, arbitrageurs obtain this benefit and liquidity suppliers only earn the fees. This affects agents' incentives to supply liquidity. Further, in the limit order market, the fact that there can be both competition that decreases transaction costs and competition that increases transaction costs means that it does not always dominate the AMM for the liquidity trader. Indeed, for assets that have lower volatility (and hence adverse selection) the AMM can be more effective (i.e., is cheaper) at providing liquidity. This observation depends on the fundamental parameters of the tokens traded.

Our extensive data allow us to make four robust empirical observations. First, we document an equilibrium pool size, pools are larger the lower the volatility of the token price, and the higher the level of liquidity trade. Second, we demonstrate that the AMM provides stable liquidity in comparison to a limit order market during an extreme market event. Third, we compare prices and volume for tokens listed on both Uniswap and Binance and find that prices are close. Pricing error is smaller when trading volume is evenly distributed between exchanges, when token price volatility is small, trading volume in general is high, transaction costs on the Ethereum blockchain are low, and when price impact is low. Consistent with our model, we find that price impact on Uniswap is small with low volatility, while price impact on Binance is higher and exhibits a high volatility. Of course, we observe this difference because of the equilibrium choice of pool size and trading venue. Thus, there are assets for which a liquidity trader would prefer an automated market maker.

Although time priority (or first in, first out queuing) is standard practice in most equity markets, there has been little research to determine the optimality of this precedence rule: Glosten (1998) is a rare exception. In this paper, he shows that pro rata rationing changes the marginal payoff to liquidity provision and hence the posted volume. Haynes and Onur (2020), using a natural experiment from the Treasury Futures market, find that under pro rata rationing, while order sizes and profitability of later submissions are higher, price efficiency is lower.

There is a large literature on liquidity provision in limit order markets. Since Glosten (1994), the efficiency of the limit order book in supplying liquidity has been widely accepted. Most modern markets operate as a form of an open electronic limit order book. We note that in contrast to equity markets, option markets have long used pro rata rationing. First, to allocate marketable orders across market makers and second to allocate early exercise assignments across writers. Hao, Kalay, and Mayhew (2009) document such rationing in options markets.

More recently, the rise of high frequency traders has generated research into competition that does not lead to cheaper liquidity. Biais, Foucault, and Moinas (2015) present a nuanced view of the effect of speed on market competition as it generates both positive and negative externalities. Empirically, Brogaard, Hendershott, and Riordan (2019) examine the limit orders submitted by HFT on a Canadian exchange. They document high limit order submission and cancellation (95% of the message traffic), which is consistent with strategic liquidity provision.

A few papers have analyzed the theoretical properties of constant function market makers. In a general framework, Angeris and Chitra (2020) show how this class of mechanisms can reflect "true" prices. They also provide a bound on the minimum value of assets held by such an automated system. These two concepts are related because of the increasing price impact faced by a potential arbitrageur. Further, Angeris, Kao, Chiang, and Noyes (2019) presents a more specific analysis of Uniswap, while Park (2021) points out blockchain related costs inherent in AMM's. Aoyagi (2020) characterizes the effect of information asymmetry on these types of markets and shows that the equilibrium liquidity supply size is stable. More recently, Ito and Aoyagi (2021) present a theory model in which an informed trader chooses between a centralized 'market maker' market and an automated market maker. By contrast, we characterize the endogenous liquidity supply for each set of markets and we solve for the optimal venue choice of a liquidity trader.

Most closely related to our work is Capponi and Jia (2021). They present a model and test of an AMM with a focus on competition among arbitrageurs. This competition allows them to consider the joint determination of gas fees and pool size. By contrast, our focus is on the comparison of a limit order market with AMM as markets for liquidity. In particular, we view the gas fees as an important pre-commitment not to withdraw liquidity from the AMM. This allows us to focus on the tradeoff between liquidity fees and adverse selection. We also note that our data are more extensive. A more recent working paper by Barbon and Ranaldo (2021) compares transaction costs and price efficiency on various decentralized exchanges and Binance for five different token pairs. They conclude that due to the fixed cost of gas fees, trading costs on the DEX are higher than the CEX. The data that they analyze comprise very rich data from Binance (a commercial vendor provides high frequency snapshots of the limit order book and transactions). By contrast, their data on the DEX is generated by low frequency (hourly) downloads. Our paper instead investigates block-by-block (15 second interval) data for the universe of UniSwap liquidity pools. Their data only permits them to use the median gas fee over the entire sample, which as we discuss, overestimates the cost of transacting.

1.1 Detailed Description of Constant Product, Automated Market Making

In this subsection we describe the market making mechanics on UniSwap V2 for readers unfamiliar with this protocol. We present an additional numerical example in Appendix B. Our model is detailed in Section 2 below.

Providing Liquidity: Each swap pool comprises a pair of cryptocurrencies. Most frequently, as we document below, one of the currencies is Eth, the native cryptocurrency on the Ethereum Blockchain. We will typically use Eth as the numeraire, and refer to the other generic coin as the 'token.' An agent wishing to provide liquidity to their preferred pool deposits both Eth and the token into the pool. The deposit ratio of Eth to token is determined by the existing ratio in the pool, which implicitly defines the Eth price of the token.

An agent who makes such a deposit receives a proportional amount of a liquidity token. This third token is specific to the pool and represents an individual liquidity provider's share of the total liquidity pool. As the pool trades with users, the value of the liquidity pool may rise or fall. Liquidity providers can redeem their liquidity tokens at any time and get their share of the current liquidity pool paid out in equal value of Eth and tokens. Changes in the composition of the pool from the time a liquidity token is minted to when it is cashed in, potentially constitute adverse selection. However, providing liquidity is potentially profitable because each trade faces a fee of 30bps which is redeposited into the pool.

Consummating Trade: Suppose a trader wishes to buy the token. In this case, he will deposit Eth into the pool, and withdraw the token. The amount that he has to deposit or withdraw depends on the bonding curve which is illustrated in Figure 1. Before the trade, there are E_0 Eth and T_0 tokens. The ratio of Eth to tokens is the implied price quoted by the pool. Someone who is interested in selling an arbitrarily small amount of the Token, would pay or receive E_0. To trade a larger quantity, consider someone who wishes to sell some of the Token. This would mean that the trader deposits some amount $T_1 - T_0$ of the token into the pool. In return, he would receive $E_0 - E_1$. Thus, the amount of Eth in the pool drops.

If the seller was a liquidity trader, the post trade price in the pool $(\frac{E_1}{T_1})$ is now too low, and a potential arbitrageur would enter the market and trade in the opposite direction to return the ratio of Eth to tokens to equilibrium.

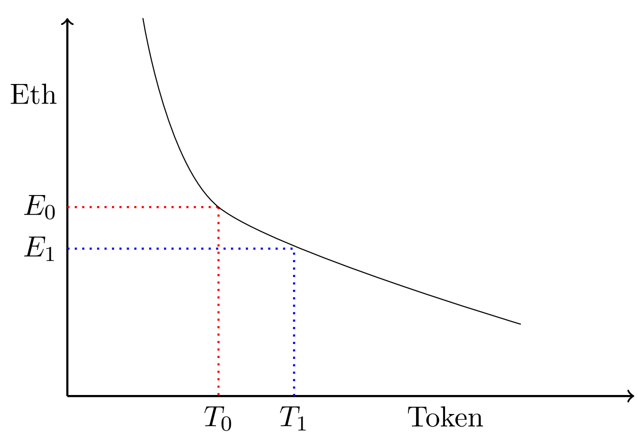

Figure 1. A bonding curve. From an initial amount of Eth and Tokens of E_0 and T_0 respectively, a trader deposits $T_1 - T_0$ tokens (sells) in exchange for $E_0 - E_1$ Eth. the price impact of this trade is determined by the bonding curve.

Specifically, if T_0 is the amount of tokens and E_0 the amount of Eth in the contract's liquidity pool, then the terms of trade are such that for any post trade quantities before any fee revenue T_1, E_1

$$k := T_1 \cdot E_1 = T_0 \cdot E_0. \tag{1}$$

In other words, the product of the Token and Eth quantities is always on the bonding curve. For each pool, the constant k, depends on the amount of liquidity that has been deposited in the pool up to this point. We note that if more liquidity is posted, the constant changes. This is the mechanism through which the market equilibrates.

Assessing Liquidity Fees: The previous clarifies the terms of trade absent the liquidity fee. Of course, remuneration is important for the liquidity providers. To see how the fee affects trades and prices, suppose that an agent wants to trade e Eth in exchange for tokens. The exchange collects a fee τ, which benefits liquidity holders.[1] Thus the effective amount of Eth that gets traded is $(1-\tau)e$. This leads to a post trade, but before fee revenue liquidity pool balance of $E' = E + (1-\tau)e$. Following the logic of the bonding curve (1), the post trade token balance must be

$$T' = \frac{T \cdot E}{E'} = \frac{T \cdot E}{E + (1-\tau)e}. \tag{2}$$

The smart contract which executes the trade accepts the e ETH and returns the difference between the pre and post trade token balances. Or, the amount of token t that the trader receives is given by

$$t = T - T' = \frac{(1-\tau)eT}{(1-\tau)e + E}. \tag{3}$$

Therefore, the terms of trade expressed in Eth/token is given by

$$p^{tot} = \frac{e}{t} = \frac{e}{T} + \frac{E}{(1-\tau)T}. \tag{4}$$

The terms of trade have a natural interpretation as a spread. Suppose that the fundamental value of the token denominated in Eth is p_0. If the pool is in equilibrium then $p_0 = \frac{E}{T}$. The liquidity fee generates what is essentially a tick size that is distinct from the volume-induced price impact that the trader pays when he moves long the bonding curve, then

$$\lim_{e \to 0} \frac{p^{tot}}{p_0} = \frac{ET}{ET(1-\tau)} = \frac{1}{1-\tau} \tag{5}$$

That is, when buying tokens, traders have to pay a fixed spread of $\frac{1}{1-\tau}p_0$. Similarly for token sales traders have to pay a fixed spread of $(1-\tau)p_0$.

Pool size: The price that a trader gets is determined by the bonding curve and the volume of posted liquidity. In particular, the price impact of a marginal increase in the order is $\partial p/\partial e = 1/T$. As the liquidity pool grows, the price impact of a fixed order size decreases. Thus, understanding the payoff to liquidity provision is an important determinant of AMM market quality.

Figure 2 presents an example of an 'orderbook' that an incoming trader might face. The blue line is for a small pool and the orange line for a larger pool. Because Uniswap has a unique

[1] Uniswap collects a fee of 30bps per trade.

mapping of trading quantity to price, the graph shows the exact amount that is traded at any price. The spread or fixed cost of trading is manifested in the interval around the mid-price of 10 for which no quantities can be bought. Clearly, larger orders face worse marginal prices as they do in a limit order book.

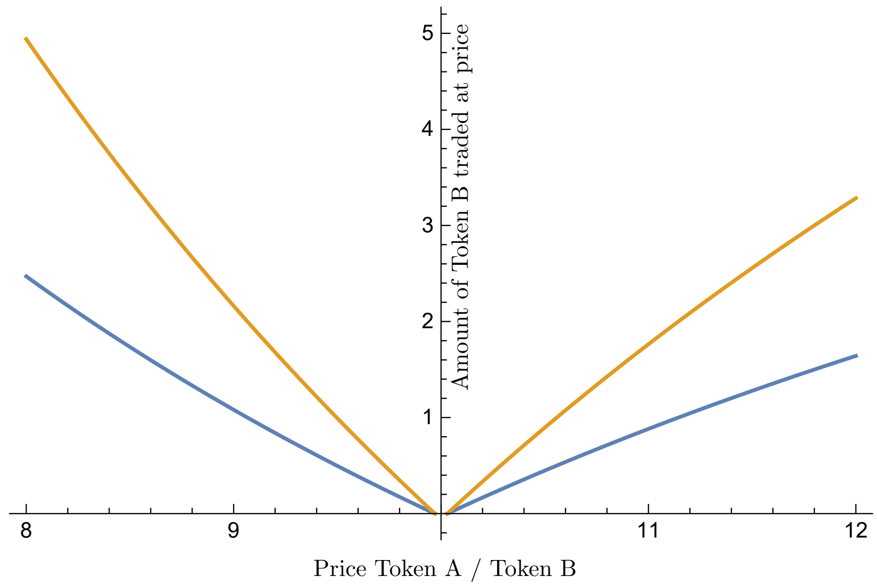

Figure 2. Uniswap orderbook depth The graph shows how many Token B could be bought or sold at a given price for a large (orange) and small (blue) liquidity pool, respectively. The parameters are: $\tau = 0.003$ and $T = 20, E = 200$ for the large pool, and $T = 10, E = 100$ for the small pool.

2 Framework

Consider a market for one asset, with current value p_0. With probability α there is an innovation and the asset is equally likely to jump up or down to $p_0 + \sigma$ or $p_0 - \sigma$ respectively, else the asset value remains p_0. A potential arbitrageur monitors the market and trades whenever profitable, otherwise a liquidity trader, who trades a fixed quantity q, arrives. The liquidity trader is equally likely to buy or sell, at any price $p \in [p_0 - \sigma, p_0 + \sigma]$. At the extreme prices a trading crowd stands ready to execute orders.

There are two rational, deep pocketed, liquidity suppliers who potentially enter the market before the passive trader and supply liquidity that optimally trades off the surplus they can extract from the liquidity trader against the possibility of being "picked off" by an arbitrageur. We focus on the case of two liquidity suppliers as it is the minimum required for competition. These rational liquidity suppliers can privately invest to increase the probability that they will be a monopolist liquidity supplier. The symmetric investment cost is $I(\gamma) = a\gamma^2$. If liquidity supplier i chooses γ_i and liquidity supplier j chooses γ_j then nature assigns i to be the sole liquidity supplier with probability $\gamma_i(1-\gamma_j)$. This private investment captures the idea that liquidity suppliers have an incentive to compete in many different ways, such as co-location or speed.

To simplify the exposition, in the text we describe the case where the informed trader buys, i.e., if there is an innovation the asset value jumps up. The case where the informed trader sells is symmetric. We characterize symmetric equilibria.

2.1 Limit order market

The sequence of events in the limit order market is as follows: First, each liquidity supplier chooses their level of private investment. Second, nature determines if there is one or two limit order submitters, and then each limit order submitter posts orders. Nature then determines the new asset value. If there is no relative value change, the liquidity trader arrives and trades against the best quote or randomizes if indifferent. If there was a relative value change, the arbitrageur trades if it is profitable.

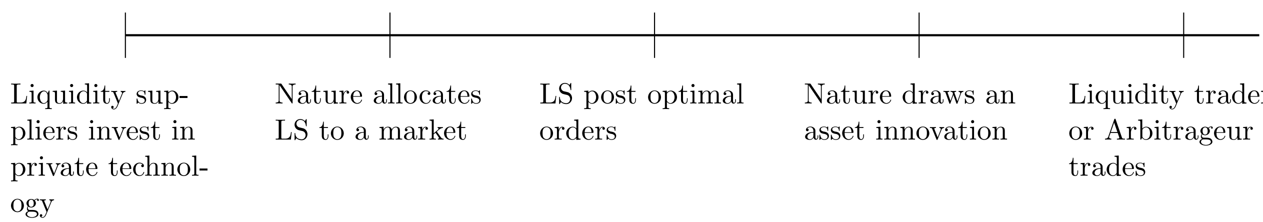

Figure 3. Sequence of Events in the Limit Order Market

From the above, it follows that liquidity suppliers may be alone or competing in a market. We characterize their optimal trading strategies in both cases, and then consider their private investment. The amount that the liquidity trader trades is fixed, q, and so this is also the amount that the liquidity suppliers post. Notice, that the informed trader will trade the maximum amount possible if it is profitable, i.e., $2q$.

If a liquidity supplier is alone in the market, then he will always post a sell price of $p_0+\sigma$. Posting at this high price completely mitigates adverse selection, and at the same time extracts maximal surplus from the passive trader.

Lemma 1 A monopolist liquidity supplier in the market, will post a sell price of $p_0+\sigma$, and a buy price of $p_0-\sigma$ to obtain an ex ante profit of $q(1-\alpha)\sigma$.

By contrast, if two competing liquidity suppliers are in the market then a liquidity supplier who charges the highest feasible price will always be undercut and lose out on the profitable trade against the passive trader. In this way, rivalrous liquidity provision will make them aggressively undercut. The symmetric equilibrium is in mixed strategies.

Lemma 2 If two competing liquidity suppliers are in the market offering orders to sell, then in the symmetric, mixed strategy equilibrium each will choose a distribution over prices $F^s(\cdot)$ over $[p^s_{\min}, p^s_{\max}]$, where

$$F^s(p) \quad = \quad \frac{(p-p_0)-\alpha\sigma}{(p-p_0)(1-\alpha)},$$

with $p^s_{\min} = p_0 + \alpha\sigma$ and $p^s_{\max} = p_0 + \sigma$.

A symmetric expression holds for competing liquidity buyers. Each competing liquidity supplier makes zero profits.

A sole liquidity supplier makes positive profits, while those in competition make zero profits. Increasing private investment makes it more likely that a liquidity supplier makes positive profits. The logic in support of competitive liquidity provision is that competition leads to lower prices. However, this is only true if suppliers are only competing on price. Given the complexity of electronic trading platforms, liquidity suppliers compete in may different ways. Such competition does not necessarily lead to lower prices and may even be inefficient. In the context of our framework, we capture this through liquidity suppliers' private investment to increase the likelihood that they are a monopolist.

Proposition 3 Each trader chooses private investment, $\gamma^* = \frac{(1-\alpha)\sigma\, q}{2a+\sigma\, q(1-\alpha)}$. The optimal private investment is increasing in σq, and decreasing in α and a.

The limit order traders find it profitable to trade with the noise traders, and σ captures the extent to which they can extract surplus from them.

2.2 Automated Market Maker

In the AMM or bonding curve market, liquidity suppliers choose a pool and commit quantities of both Eth and Tokens. The first thing to observe is that liquidity provision is not rivalrous and there is no incentive for private investment as there is in the limit order market. The second thing to observe is that the AMM requires committed capital.

Suppose that investors have each committed E_0 Eth and T_0 tokens. Two identities from the bonding curve will be useful: First, the ratio of Eth to tokens, in equilibrium, is the Eth price of tokens implied by the bonding curve market or

$$\frac{E_0}{T_0} = b_0. \tag{6}$$

We will start our analysis under the assumption that the price in the bonding curve market is equal to the equilibrium price, or $b_0 = p_0$. Second, any transactions must occur along the curve. Specifically, if an amount t tokens and e ether are traded, then

$$E_0 T_0 \quad = \quad k, \tag{7}$$

$$\text{and} \qquad (E_0+e)(T_0+t) \quad = \quad k \tag{8}$$

where k is the constant of the bonding curve.

In addition, as we indicated in subsection 1.1, any trade is assessed a proportional fee of τ. We simplify the algebra that follows by assuming that liquidity fees do not change the size of the pool, but are placed into a separate account. In reality, liquidity fees are paid into the pool and therefore change the bonding curve constant. The timing of events is show in Figure 4 below.

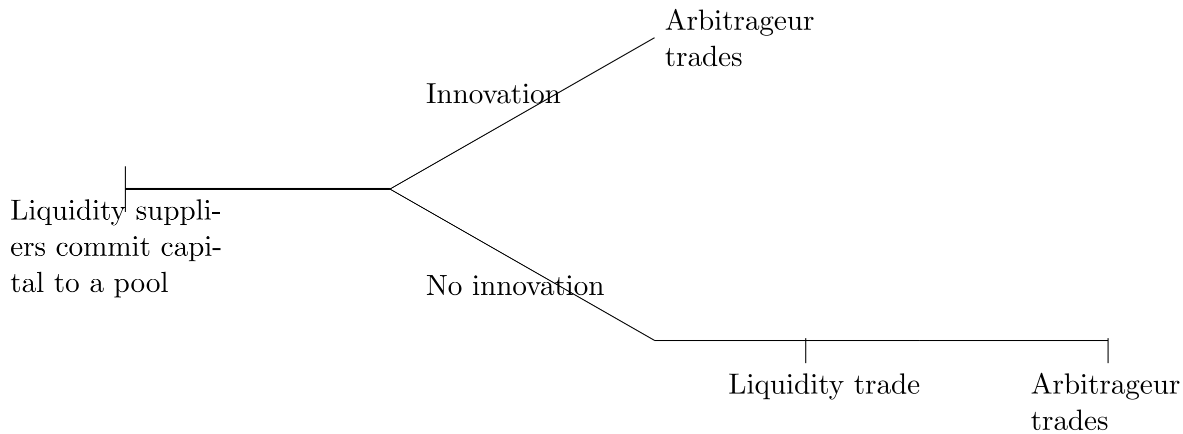

Figure 4. Sequence of Events in the Automated Market Maker

First consider the payoffs to liquidity provision if a liquidity trader arrives. If the liquidity trader buys tokens, they will remove q tokens, and will buy these with the numeraire good, Eth. Thus, they add in e_ℓ^b to the Eth pool. The specific amount of Eth they add is determined by the bonding curve, so

$$\begin{aligned}(E_0 + e_\ell^b)(T_0 - q) &= E_0 T_0 \\ e_\ell^b &= \frac{E_0 T_0}{T_0 - q} - E_0.\end{aligned}$$

Per the protocol, liquidity providers receive a fee for facilitating this transaction. We model this by assuming that a fee is deposited into a separate account. The fees paid by this noise trader are τe_ℓ^b.

Reversing this trade (so selling tokens to the pool) is profitable for the arbitrageur if $e_\ell^b(1-\tau) < p_0$. If τ is sufficiently small, the arbitrageur will add tokens and remove Eth. Under our maintained assumptions, the payoff to liquidity provision is twice the fee paid by the liquidity trader. Or, $2\tau e_\ell^b$. Symmetrically, if the liquidity trader is a seller, they will deposit q tokens and remove $e_\ell^s = E_0 - \frac{E_0 T_0}{T_0+q}$. The arbitrageur will buy tokens. These two transactions generate a fee revenue of $2\tau e_\ell^s$.

Our formulation makes use of two simplifications for tractability. First, in reality, the arbitrageur faces a different size pool than the liquidity trader as the liquidity fee has been paid into the pool. Second, the arbitrageur will trade slightly less than e_ℓ^b. The overall size of these is negligible and of the order τ^2.

Lemma 4 Suppose that the aggregate amount of Eth and Tokens in a liquidity pool are E_0 and T_0 respectively, and τ is sufficiently small. With probability $(1-\alpha)$, there is a liquidity event and the fee revenue for liquidity provision is:

$$2\tau p_0 q\left(\frac{T_0}{T_0^2-q^2}\right). \tag{9}$$

Now suppose that there was a positive innovation event so that an informed arbitrageur arrives. Since the pricing is deterministic she will trade an amount that maximizes her profit. She will buy t^b tokens and pay e_I^b for them. The Eth payment is pinned down by the bonding curve, which requires that $(T_0-t^b)(E_0+e_I^b)=E_0T_0$ which gives her a profit function of

$$\pi_I^b \;=\; (p_0+\sigma)t^b-(1+\tau)\left[\frac{E_0T_0}{T_0-t^b}-E_0\right]. \tag{10}$$

Given the convexity of the bonding curve, the optimal trading amount is determined by the first order condition,

$$(p_0+\sigma)-(1+\tau)\frac{E_0T_0}{(T_0-t^b)^2} \;=\; 0, \tag{11}$$

which implies

$$t^b \;=\; T_0-\sqrt{\frac{(1+\tau)E_0T_0}{p_0+\sigma}} \tag{12}$$

$$e_I^b \;=\; \sqrt{\frac{E_0T_0(p_0+\sigma)}{(1+\tau)}}-E_0. \tag{13}$$

We note that is optimal for the informed arbitrageur only if $t^b\geq 0$, or $\tau\leq\frac{\sigma}{p_0}$, so that the transaction cost is low relative to the innovation.

After the innovation, absent informed arbitrage trading, the Eth value of the total supplied capital would be $E_0+(p_0+\sigma)T_0$. Given the informed arbitrage trade, the Eth value of the supplied capital is

$$\begin{aligned}&E_0+e_I^b+(p_0+\sigma)(T_0-t^b)\\ =\;&\sqrt{E_0T_0(p_0+\sigma)}\left(\frac{2+\tau}{\sqrt{(1+\tau)}}\right).\end{aligned}$$

Therefore, the change in value of supplied capital for liquidity suppliers after an increase in the value of the asset is:

$$\left(\frac{2+\tau}{\sqrt{(1+\tau)}}\right)\left(\sqrt{T_0E_0(p_0+\sigma)}\right)-\left(E_0+(p_0+\sigma)T_0\right). \tag{14}$$

This change in value corresponds to "picking off" risk, in the sense that the informed arbitrageur rebalances the amount of Eth and Tokens to reflect the value in the wider market. In addition, however, the arbitrageur pays a liquidity fee. Consistent with the previous case, we assume that this is levied on the Eth total, for an amount equal to $\tau\left(\sqrt{\frac{E_0 T_0(p_0+\sigma)}{(1+\tau)}} - E_0\right)$.

Lemma 5 Suppose that the aggregate amount of Eth and Tokens provided are E_0 and T_0 respectively, and $\tau \leq \frac{\sigma}{p_0}$. With probability $\frac{\alpha}{2}$, the asset value jumps up and the payoff to liquidity provision is:

$$\underbrace{\left(\frac{2+\tau}{\sqrt{(1+\tau)}}\right) T_0 \left(\sqrt{p_0(p_0+\sigma)}\right) - (T_0 p_0 + (p_0+\sigma)T_0)}_{\text{Picking off}} + \underbrace{\tau\left(T_0\sqrt{\frac{p_0(p_0+\sigma)}{(1+\tau)}} - p_0 T_0\right)}_{\text{liquidity provision}}. \quad (15)$$

A symmetric expression holds for a jump down in the asset value.

Armed with Lemmas 4 and 5 we can determined the overall payoff to liquidity provision for the entire pool and thus the equilibrium size of the pool.

Proposition 6 Suppose that $0 < \tau < \frac{\sigma}{p_0}$, then the equilibrium supply of Tokens is given by

$$T_0 = q\left[\sqrt{1 + \frac{(1-\alpha)^2\tau^2 p_0^2}{\alpha^2\omega^2}} - \frac{(1-\alpha)\tau p_0}{\alpha\omega}\right] \quad (16)$$

Where $\omega = \sqrt{p_0(p_0+\sigma)(1+\tau)} + \sqrt{\frac{p_0(p_0-\sigma)}{1+\tau}} - 2p_0$.

The equilibrium pool size makes liquidity suppliers indifferent between committing capital to the AMM or to their outside option which, for simplicity, we normalize to zero. The parameter restrictions under which equation 16) characterizes pool size are intuitive. If the payoff to liquidity suppliers is too small, then pools are not sustainable. If the payoff to liquidity provision is too large, then arbitrageurs will not find it lucrative to trade on relative price information, and will also not trade to ensure that the price implied by the pool corresponds to the traded value of the token. Going forward, we assume that these conditions hold because our empirical focus is on the pools which we observe; further as we show in our empirical analysis, the prices implied by Uniswap consistently track traded prices on Binance which suggests that cross-market arbitrage is profitable.

3 Data and stylized facts

Decentralized exchanges (DEX) are smart contracts mostly deployed on the Ethereum blockchain. Users initiate trade as an Ethereum transaction that sends tokens to a smart contract which calls a function to perform the exchange. The smart contract then sends trade proceeds in the form of the appropriate tokens back. Since transactions on Ethereum are atomic, or in other words they either execute completely or fail, there is no settlement risk and users do not have to hand over custody of their digital assets to a third party. In keeping with other DeFi protocols, the source code for many DEXs is public and users can verify that the code is not fraudulent. This also allows them to perfectly predict how the smart contract will perform, and more specifically, the terms of trade.

Uniswap is open source with no owner or operator. It was launched in November 2018 at Devcon 4, and the first pool allowed swaps between ETH and the Maker token (MKR). In its first release, Uniswap V1 allowed exchange between any ERC20 tokens against Ether (ETH).[2] If no pool exists for a specific token pair, it can be freely created by invoking the Uniswap factory contract, and specifying the token for which a new pool should be created. The factory contract will then deploy a new pool for that specific token on the Ethereum blockchain, and enable subsequent trade.

Uniswap V2 was launched on May 18, 2020. The update allows direct trade of any ERC 20 token pairs and includes Tether (USDT). As we document below, most V2 pools trade tokens against wrapped Ether (WETH), which is a ERC 20 representation of Ether (ETH). In addition, V2 generates a moving average of past prices which other smart contracts can use as a reference price or "oracle." Many other smart contracts use Uniswap as a price feed, in the same way that traders in traditional financial markets use Bloomberg.[3] Because Uniswap has no designated operator, V1 pools cannot be deleted from the blockchain and exist in parallel to V2 pools.

Before we describe the data, it is important to note that in contrast to traditional exchanges, UniSwap liquidity pools are not certified. There are no listing requirements. A consequence of this is that some of the token pools are purposely misleading. For example, five different tokens in our sample share the same ticker symbol USDC. A naive user, who does not verify the relevant smart contract addresses, could be tricked into buying a worthless coin with the same ticker.[4] Overall trading activity in these fradulent pools is limited and will not affect our results. We provide detailed information on fake tokens in Appendix A.

To collect our data, we obtained a list of all Uniswap V1 and V2 liquidity pools from the original factory contract transactions. Our sample comprises 36,958 individual liquidity pools, consisting of 3,937 V1 pools and 33,021 V2 pools. We matched transactions into and out of these liquidity pools with block-by-block transactions on the Ethereum block chain. In total we have 47,204,920 transactions on Uniswap from its inception on November 2, 2018 until May 20, 2021. From the Ethereum blockchain we observe 1,084,581 liquidity injections into a pool, 582,063 withdrawals

[2] ERC 20 is the most commonly used standard for fungible tokens on Ethereum compatible blockchains. ERC 20 tokens must provide a standard interface through which other DeFi applications such as Uniswap can interact with the token contract.

[3] Oracles provide information to other smart contracts. There have been cases of oracle manipulation: traders placed large orders on Uniswap V1 to affect the price used by other smart contracts to value loan collateral. Once prices revert to normal they default on the undercollateralized loan.

[4] The 'real' USDC stable coin resides under address 0xa0b86991c6218b36c1d19d4a2e9eb0ce3606eb48 and has over 4 million token transactions on the Ethereum blockchain. A token with the same ticker is 0x0xEFb9326678757522Ae4711d7fB5Cf321D6B664e6. Somebody created a Uniswap liquidity pool for this copycat token at the address 0x1bffb8a3fede9f83a3adc292ebf1716d40b220c1, which has a total of 10 trades and the size of the pool never exceeded 50 ETH.

of liquidity from a pool, and 45,481,500 trades of tokens.

An Ethereum transaction is a set of instructions that are processed as one unit within the same block. Apart from implicit limits on transaction sizes given by Ethereum block size, there is no limit on how many interactions with Uniswap liquidity pools can be done in one transaction. In our sample 79.7% of Ethereum transactions only have one interaction with one liquidity pool, another 18.4% have two interactions. 1615 transactions or 0.006% if the sample have 10 or more interactions with a liquidity pool. The most complex transaction in our sample has 60 interactions with 6 different Uniswap liquidity pools.[5] It is important to recognize that our analysis is on liquidity supply on both a limit order market and an AMM. Therefore, we do not consider market access fees – either on the limit order market (e.g., co-location) or on the AMM (e.g., gas fees). The latter are analyzed in Capponi and Jia (2021).

Transactions are finalized if they are incorporated onto the Ethereum blockchain, and anything that happens in the same transaction effectively happens at the same time. For this reason, flash swaps were introduced in Uniswap V2. This feature allows a user to borrow any amount up to the total liquidity available in a pool, only if the whole sum gets returned in the same Ethereum transaction. Because an Ethereum transaction is atomic, – i.e., it is either executed in its entirety or not at all – there is no credit risk for lenders as the loan is both originated and repaid in the same transaction.[6] In our data, 56,606 transactions combine liquidity additions or removals with swaps or flash swaps. To be consistent with websites like uniswap.info we include flash swaps in volume computations in this paper. Liquidity providers earn a fee identical to the one on regular token swaps that is based on the gross amount of the flash loan regardless whether the repayment is the same token that was borrowed or not. For liquidity providers flash swaps offer a risk free way to earn earn higher fees.

Most pools trade against WETH. To compute volume we take the WETH part of the trade and convert it to USD using Binance minute by minute data. Most of the remaining pools trade against a USD stablecoin. For those pools we convert the amount traded in the stablecoin to USD. For all remaining pools we search for all pools where one of the tokens trades against WETH and convert it using the prices from the pool with the highest volume.

Table 1 provides an overview of the 10 largest Uniswap V1 and V2 pools by total aggregate volume in ETH. V1 pools are smaller both in terms of volume as well as in number of trades, mostly because the introduction of V2 coincided with a huge boom in Decentralized Finance and caused most traders and liquidity providers to converge on the new protocol. The largest pool in terms of total volume traded is Tether (USDT) - Wrapped ETH with an aggregate volume (over all days) of over 26.5 billion USD. This pool also has the highest number of total trades in our sample – over 2.75 million trades. The "on average" largest pool in ETH is FEI - WETH

[5]see Ethereum transaction 0x2d732ab5aeb05eeb52eebb9a6086e77b15198fe61a827648b2e43a79fb1902ec. Uniswap V2 introduced router contracts that can perform complex transactions with one function call. Assume for example that a pool exists that trades tokens A and B, another pool trades tokens B and C, and there is no pool to swap A and C. The router contract can then be instructed to swap A and C by trading through token B. In our sample such a transaction would show up as two separate transactions, one for each of the two involved pools.

[6]Other protocols also offer flash loans, but Uniswap is unique because the borrowed amount can be repaid in any combination of pool tokens as long as the value repaid equals the value borrowed. Borrowers pay a fee on amounts borrowed.

Token 1		Token 2		Number Transactions	Volume (ETH)	Volume (USD)	Pool size (ETH)
		Panel A: Uniswap V2					
Wrapped Ether	WETH	Tether USD	USDT	7,516.2	83,445	72,383,925	211,915
USD Coin	USDC	Wrapped Ether	WETH	5,757.4	81,018	71,535,793	197,864
Dai Stablecoin	DAI	Wrapped Ether	WETH	3,008.9	46,683	36,897,989	162,671
Uniswap	UNI	Wrapped Ether	WETH	2,429.9	31,156	26,624,652	53,511
Wrapped BTC	WBTC	Wrapped Ether	WETH	957.9	29,277	23,932,848	284,151
Fei USD	FEI	Wrapped Ether	WETH	288.6	26,780	68,605,073	374,990
yearn.finance	YFI	Wrapped Ether	WETH	872.1	19,994	9,318,935	27,322
Tendies Token	TEND	Wrapped Ether	WETH	144.3	16,260	24,569,585	724
SushiToken	SUSHI	Wrapped Ether	WETH	894.5	14,860	6,750,425	77,097
Wrapped Ether	WETH	Truebit	TRU	3,680.3	14,171	43,746,104	1,647
		Panel B: Uniswap V1					
Ether	ETH	Dai Stablecoin	DAI	540.6	2,681	524,088	9,226
Ether	ETH	HEX	HEX	219.4	1,801	378,702	22,300
Ether	ETH	USD Coin	USDC	258.0	1,274	287,165	6,858
Ether	ETH	Maker	MKR	118.3	1,101	217,221	11,010
Ether	ETH	LoopringCoin V2	LRC	20.5	983	365,065	794
Ether	ETH	Sai Stablecoin v1.0	SAI	166.4	770	153,078	5,030
Ether	ETH	Synthetix Network Token	SNX	124.8	700	130,702	3,480
Ether	ETH	Synth sETH	sETH	44.1	576	110,465	26,579
Ether	ETH	UniBright	UBT	108.0	279	58,212	635
Ether	ETH	Pinakion	PNK	40.7	197	59,877	1,544

Table 1. Ten largest exchanges for Uniswap V1 and V2, respectively, sorted by volume. Number transactions is the daily average number of transactions, Volume (ETH) is the daily average volume in Ether, Volume (USD) is the daily average volume in USD, and Pool size (ETH) is the daily average pool size in Ether.

with an average size of around 375 thousand ETH.

While some of the pools are very active, many are not: 24,466 pools in our sample have fewer than 100 transactions. Figure 5 shows the number of trades by pool type. With the release of V2 trading activity in V1 declined. We also observe that for V2 pools trading against WETH (orange) dominates direct trading of other tokens (red).

Figure 6 illustrates the trading volume per day. Computing volume is not straightforward in bonding curve markets as attackers often deliberately push markets out of equilibrium. The highest volume in our sample is on March 31, 2021 with a volume of 18.33 billion USD. On that day, a trader moved 5.5 billion USD of a token back and forth between her own wallets. Another spike was on October 26 with a volume of over 5.5 million ETH or USD 2.1 billion, and is linked to an attack on Harvest Finance using a flash swap. A more detailed discussion of this incident and implications for Uniswap volume can be found in Appendix A. Use of Flash Swaps varies in our sample. Out of the 379 days when V2 was deployed, flash swaps occurred on 339 days. The median flash swap volume per day was 41,265 USD and the maximum was 17,1 billion USD on March 30, 2021.

The largest non-flash swap trade in our sample was on December 17, 2020 when a trader swapped 48,584,947.17 DAI for 342,252.89 WETH, worth about USD 220.4 million at the time as part of

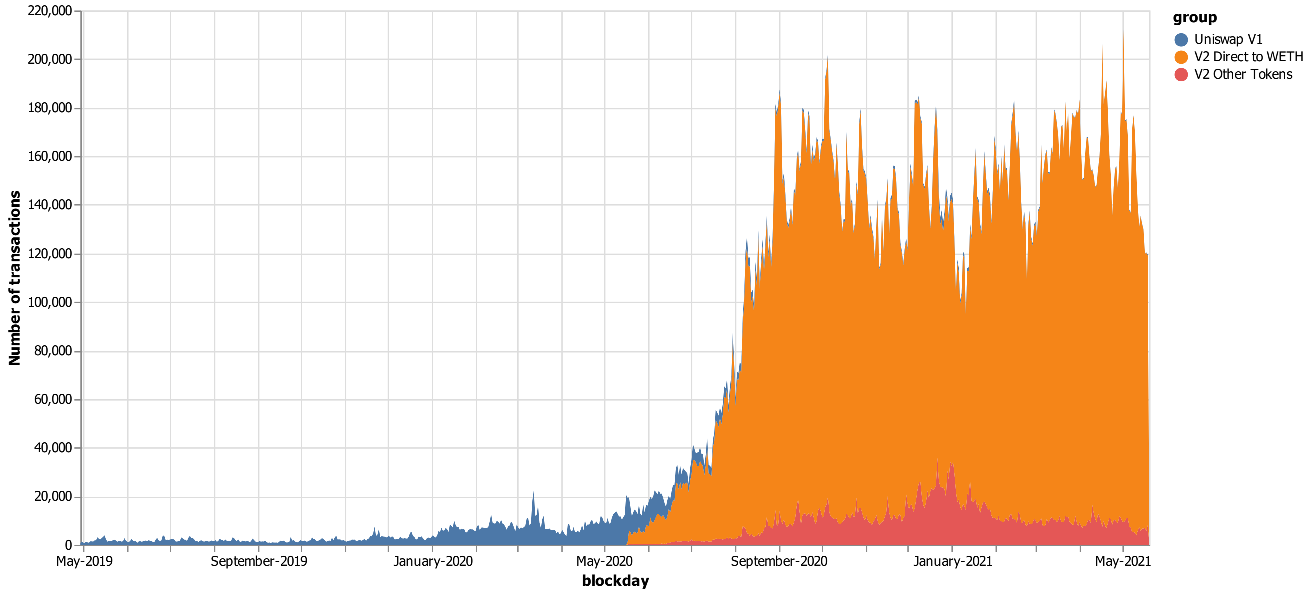

Figure 5. Number of transactions on Uniswap.

an attack on the platform Warp finance. Many large trades are part of an exploit that targets weaknesses in a platform's code. On June 18, 2020, a trader swapped 100,000.39 WETH (about USD 23.2 million USD at the time) for 1,695,998.19 UniBomb tokens as part of another exploit.[7] The median trade size in our sample is 845.21 USD. 33.5% of trades are below 0.5 ETH and 15.3% are below 100 USD. We provide details on our methodology in Appendix A.

Figure 7 presents the network of pools between all tokens that are part of the 50 largest pools by volume. The thickness of each edge corresponds to the trading volume between the tokens.. We can see that Wrapped Ether (WETH) takes a central position in the Uniswap network. For our whole sample of 36,958 tokens we find that 30,912 tokens, or 73.64%, trade directly against WETH. The second highest number of tokens, 1,538, trade against USDT. The highest volume and the most connections are between WETH and USD stable coins such as USDT, USDC, and DAI. 2,913 tokens or 7.88% of tokens are trading directly against these three stablecoins. We note in passing that the Uniswap network has a core-periphery structure similar to many other financial networks. 27,773 tokens or 89% of tokens trade only against one other token.

For our subsequent econometric analysis we purge pools that are very small or were only used for a few days. These pools may have been used for experiments, development, or exploits such as the fake tokens mentioned above. Specifically, we drop pools with less than 30 trading days and with less than 100 ETH average balance. The reduced sample has 37,902,764 observations

[7]See transaction 0x0x8492ce3b1ea8ec796471997731e557c057c2fb0a3ade7f9c0477450d53ad4791. Unibomb is deflationary token that burns 1% of the each transaction, thus increasing its value. Somebody seems to have borrowed 100.000 ETH from the lending platform dXdY and converted them to Unibomb. The transaction decreased token supply and the user could reconvert the Unibomb tokens to ETH with a slight gain in price, leaving a profit after the repayment of the loan.

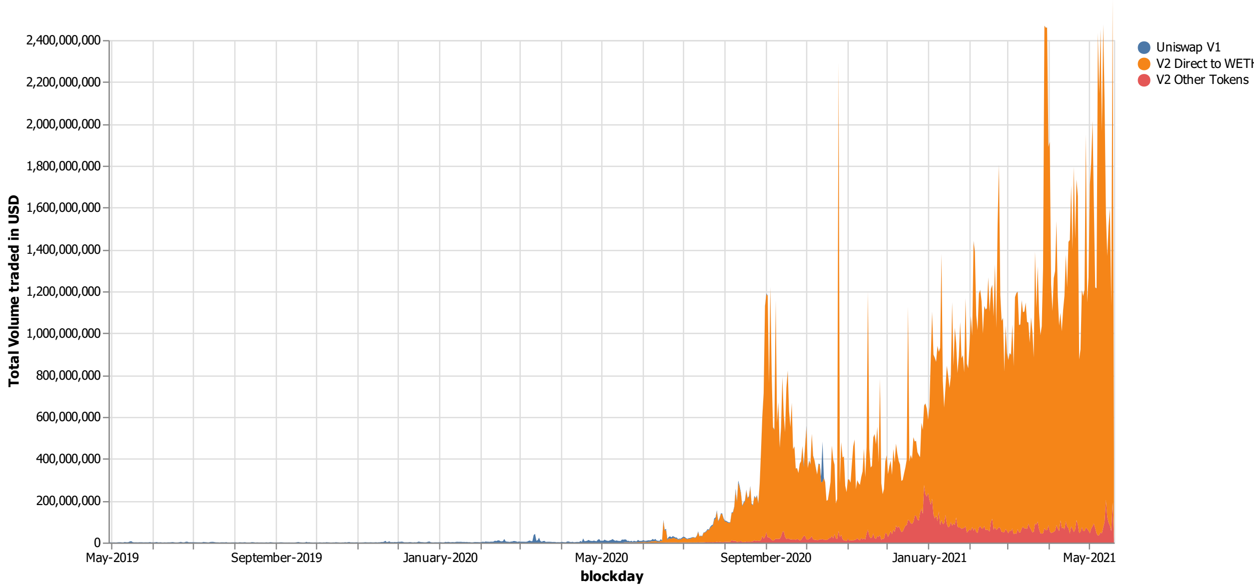

Figure 6. Trading volume on Uniswap The graph shows the daily volume of transactions including flash loans in USD capped at 2.4 billion USD.

in 1,376 pools.

4 Liquidity provision

4.1 Pool size

We have demonstrated empirically that there is heterogeneity in pools. The equilibrium pool size in Equation (16), permits us to derive comparative statics.

Corollary 1 Suppose that $\tau < \frac{\sigma}{p_0}$, then the equilibrium size of a liquidity pool is

i. Linear in the size of the liquidity trade.

ii. Is decreasing in the size of the innovation, σ.

iii. Is decreasing in informed trades, α.

In equilibrium, pool size adjusts so that a pool's per token fee revenue balances the cost from being picked off by arbitrageurs if the pool posts stale prices. When innovations, σ, are larger, the informed arbitrageur will ceteris paribus place larger trades, and losses for liquidity providers will therefore be larger. By contrast, fee revenue from the liquidity trader is independent of movements in the fundamental price. This implies that to increase the expected payoff per unit

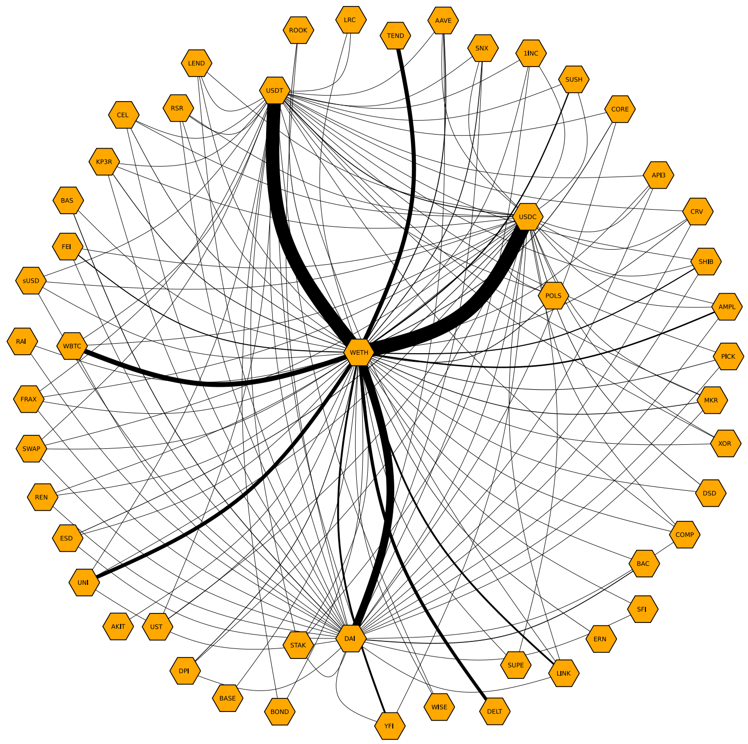

Figure 7. Network graph of pools between all tokens that are part of the 50 largest pools by volume.

of liquidity, the pool in equilibrium must be smaller. As the pool shrinks, price impact increases and so the informed arbitrageur places a smaller order. These comparative statics highlight how the AMM reaches equilibrium – as the price impact is a deterministic function of the pool size, the size of the pool adjusts to ensure that the liquidity suppliers receive the opportunity cost of their capital.

The equilibrium pool size also decreases in α, the intensity of informed trading, i.e. an arbitrageur trading because of an innovation. As α increases the fraction of liquidity traders decreases and the pool is more likely to be picked off by arbitrageurs. The pool size again shrinks in equilibrium, increasing price impact, and reducing optimal the trade size of the informed trader. Figure 8 shows the poolsize as a function of σ and α for a numerical example.

To test these predictions, we collect daily data on all 1,376 pools in our final sample. For the average exchange, we observe 208 days, while the median is 205. In Table 2 we regress pool size on price volatility and measures of uninformed trading. We control for the airdop of UNI tokens during Uniswap's liquidity minting program as they might have affected the return for liquidity providers. Between September 18, 2021 and November 17, 2021 13.5 free UNI tokens per pool per block were distributed to users in four pools.[8] Consistent with our theoretical predictions, we find that pool size decreases in token price volatility, which is our empirical proxy of the

[8] These pools are the WETH/USDT, WETH/USCD, WETH/DAI, and WETH/WBTC pools.

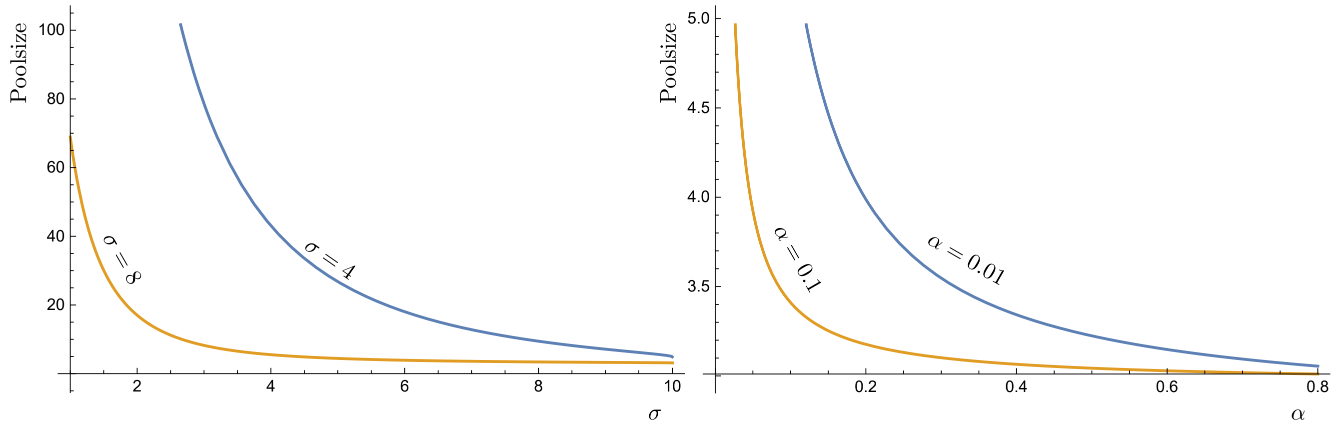

Figure 8. Poolsize as function of innovation size σ and α. Unless otherwise stated the parameters are $\alpha = 0.01, p_0 = 10, \sigma = 4, q = 3, \tau = 0.003$

size of the innovation σ. Higher token price volatility means that the pool loses more when it gets picked off, thus liquidity providers in equilibrium are compensated with higher fee revenue which is achieved by reducing the pool size.

If price innovations arrive with an exogenous intensity, then a higher number of trades in a given interval must correspond to more liquidity trading. Consistent with this idea we find in columns (2) and (3) that pool size increases in trading volume. In column (4) we confirm that pools with a higher number of trades per day are on average larger. Finally we examine reversals, which we define as a trade that is immediately followed by a trade in the opposite direction of at least 50% of the size of the original trade size. This pattern is observed whenever a liquidity trader who pushes the price away from its fundamental value is followed by an arbitrageur who brings the pool quoted price back to the appropriate relative value. Consistent with the logic of our model we find that pools with more reversals are larger.

4.2 Stability in Liquidity provision and Gas Fees

One of the characteristics of modern limit order markets is the rapid posting and cancelling of liquidity, and quote flickering. This is usually attributed to strategic liquidity provision. (Baruch and Glosten (2013) characterize equilibrium quote flickering in a limit order book.)

In principle, liquidity provision in an AMM could be as strategic as in a limit order market. To investigate if liquidity providers engage in high frequency strategic liquidity provision we track liquidity providers by the wallet address where they store their liquidity tokens. Out of our 48 million observations we find only 1,801 events where the same person deposited (withdrew) liquidity and withdrew (deposited) liquidity in the same pool within 50 blocks. Most of these withdrawals or additions are small, the median size is USD 146.75 or 0.0159% of the pool, which has a negligible impact on trading costs. We find only 18 observations in total that are in

	(1)	(2)	(3)	(4)	(5)
Volatility	-0.512***		-0.440***	-0.508***	-0.455***
	(0.101)		(0.0761)	(0.0894)	(0.0852)
Airdrop	250.5***	219.5***	218.9***	224.4***	221.0***
	(24.49)	(25.92)	(25.90)	(31.21)	(32.52)
Volume (USD)		0.620***	0.617***		
		(0.170)	(0.170)		
Reversals				35.85***	
				(10.34)	
Number Trades					5.305***
					(1.651)
R^2	0.217	0.316	0.320	0.257	0.259
Observations	263,750	279,040	263,750	263,750	263,750

Table 2. Regression explaining pool size as a function of price volatility and measures of uninformed trading. Pool size is the daily average size of the liquidity pool measured in million USD. Airdrop is a dummy set to one if transactions in that pool were rewarded with UNI tokens at that time, Volatility is the daily standard deviation of block by block price changes of the pool in percent. Volume (USD) is the daily trading volume in million USD, Number Trades is the number of trades for a pool per day in Thousands, and Reversals are defined as a trades that are immediately followed by an opposite trade of at least 50% of the size of the original trade measured in thousands. Standard errors are clustered by pool. One, two, and three stars indicate significance at the 10%, 5%, and 1% level, respectively.

pools with over USD 10,000, where the deposit or withdrawal is for more than USD 1,000, and where the gap between deposit and withdrawal is no more than 5 blocks. We conclude that liquidity providers on Uniswap do not engage in high frequency strategic liquidity provision and withdrawals.

We note that the cost structures differ between the AMM and a limit order market. In a centralized exchange, assets are held by the exchange, and so supplying liquidity across different assets can be done almost costlessly as this is effectively a local book-keeping entry. By contrast, in the non-custodial UniSwap system, liquidity has to be withdrawn from one pool and added to another. Both legs of these transactions incur transaction costs known as "gas fees."

Gas fees are the transaction costs on Ethereum blockchain. Gas is quoted in units of gwei, where 1 Gwei is 10^{-9} Eth. The amount of gas used for a specific transactions consists of a predictable component and a state contingent component. The predictable component comes from the organization of the Ethereum Virtual Machine. Any transaction comprises a known number of individual low-level computational steps. Fixed gas amounts are assigned to each type of operation and are independent of market conditions – the total gas used for a transaction is simply the sum of each computational step. In addition to this fixed amount, a user offers an amount of Eth per unit of gas consumed by the program as an incentive to the miners to be included. This incentive amount varies with market conditions and demand for transactions. Given the price of Ethereum, the gas fee can be translated into USD.

There are three characteristics of gas fees that are relevant for our analysis. First, the average gas fee is low, second the state contingent fee can be high and finally, the incremental fee relative to trading on Binance is low. The high state contingent fee acts as a pre-commitment not to withdraw liquidity in stressed market events.

Liquidity pools are not subject to the short term evaporation of liquidity that we observe in limit order markets. Indeed, we only observe 445,136 liquidity withdrawals which are 1.17% of all the interactions with Uniswap liquidity pools in our sample. Many withdrawals are either small or are from small and illiquid pools. It is also noteworthy that liquidity in pools also does not get suddenly withdrawn in extreme market events. On May 19, 2021 Ether dropped from over USD 3,400 to 2,014, a 41% decline. To put this in perspective, on October 19, 1987 the S&P 500 index only dropped by 20.5%. The orange area in Figure 9 shows minute by minute pricing data from Binance. We examine liquidity withdrawals from the large USD stablecoin pools that trade against ETH, USDT, TSDC, and DAI. The blue line shows aggregate withdrawals from these pools in percent of pre-event poolsize. When Ether reached its lowest price, only about 2% of liquidity was withdrawn and users could trade with minimal price impact. Most withdrawals of liquidity happen around 53 minutes or 245 blocks after the big decline when prices have already recovered. Overall, only 17% of the liquidity gets withdrawn in this extreme price movement, which could also reflect reduced expectations for future trading activity and thus fee revenue.

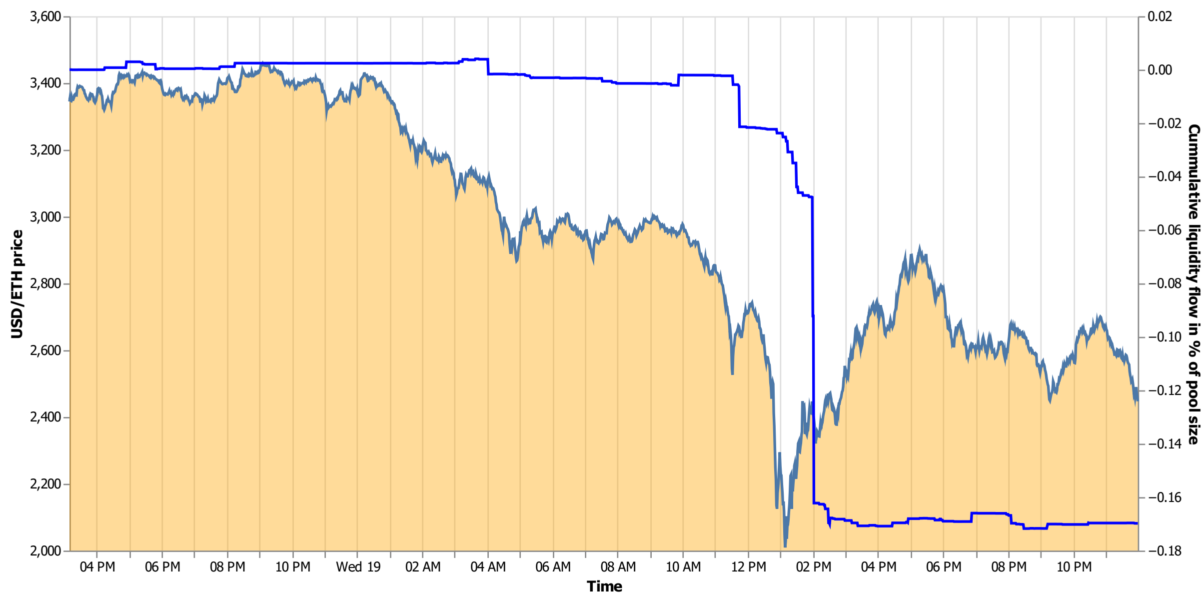

Figure 9. Liquidity withdrawals and extreme market events. The graph shows the minute by minute price of Ether from Binance around March 19 (in orange) and the aggregate percentage liquidity withdrawals from the four largest ETH - USD stablecoin pools on Uniswap.

Rapid withdrawls of liquidity supply as we highlighted in our model are a feature of modern limit order markets. Further, such withdrawls have contributed to "flash crashes." By contrast, liquidity suppliers in Uniswap pools are very stable. In as much as deep and constant liquidity is socially beneficial, the design of the AMM is effective. We reiterate that Ethereum gas fees act as a commitment device for liquidity suppliers to remain in pools.

5 Ranking Exchanges

5.1 Theory

To compare the limit order market with an automated market maker we focus on the total trading cost for a liquidity trader, which consist of fees and price impact. Recall, each liquidity supplier has undertaken private investment and is a monopolist with probability $\gamma^*(1-\gamma^*)$. The liquidity trader thus faces a monopolist with probability $2\gamma^*(1-\gamma^*)$, which we denote by η and enjoys competitive liquidity provision with probability $1-\eta$.

Proposition 7 The expected cost to the liquidity trader is

$$E(c^{limit}) = (1-\eta)\left(\alpha\sigma + \frac{\alpha}{(1-\alpha)^2}\Gamma(\alpha,\sigma)\right) + \eta\sigma,$$

where $\Gamma(\alpha,\sigma) = \sigma\left[1-\alpha^2+2\alpha\ln(\alpha)\right]$

In the AMM, the expected cost to the liquidity trader is

$$E(c^{AMM}) = p_0(1+\tau)\left(\frac{\lambda^b+\lambda^s}{2}\right) - p_0,$$

where $\lambda^b > 1$ and $\lambda^s > 1$ are constants.

The limit order market does not always dominate the automated market maker. When informed trading is low, the liquidity pool is large, price impact is low and thus cheaper for the liquidity trader. The price impact in the limit order market does not decrease in informed trading to the same extent because liquidity provision is not as competitive. Liquidity providers seek to find opportunities that allow them to extract rents from imperfectly competitive liquidity provision. In our model the private cost of obtaining market power, $I(\cdot)$, captures this incentive in a stylized way. When a is high, so obtaining market power is more costly, then liquidity provision is more competitive. We link $I(\cdot)$ to several stylized facts in financial markets. Market makers invest heavily in high frequency trading, with the main objective to carve out a niche with less competition by, for example, receiving information faster than competitors or by detecting institutional traders earlier than rivals. Such investment is wasteful from the perspective of a liquidity trader. The cost can also be seen as costs that deter entry to liquidity provision, leading to the high concentration of liquidity provision we see in financial markets today.

Our model also implies that price impact in limit order markets is more volatile because the trader does not know ex-ante if he will face a monopolist liquidity provider or a competitive market. For automated market makers price impact is known ex-ante and can be very low when the pool is sufficiently large. We summarize this intuition below.

Observation 8 The limit order market does not dominate the automated market maker.

i. There are robust parameter ranges for which the automated market maker dominates the limit order market.

ii. Conditional on trading quantity q, for a pool in equilibrium, price impact is more volatile in the limit order market than the AMM.

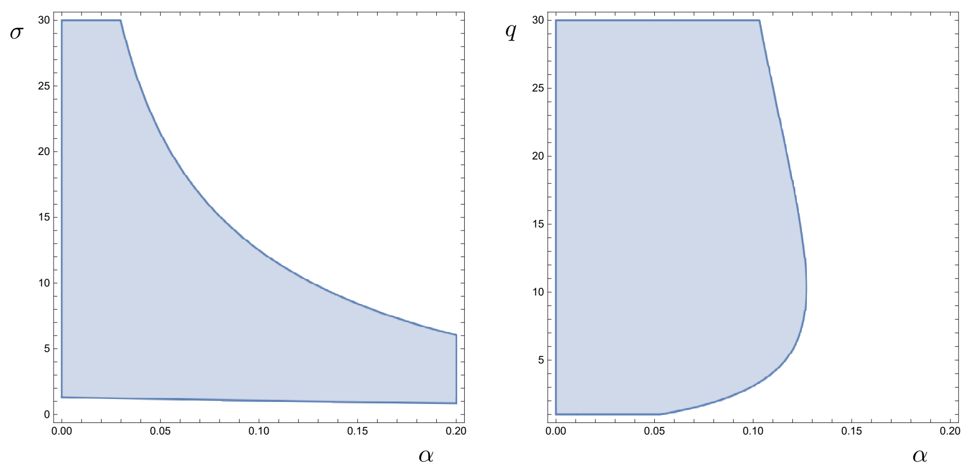

Figure 10. Region of the parameter space where AMMs dominate limit order markets. The graph shows the region of the parameter space where trading costs for the liquidity trader in the liquidity pool are lower than in the limit order market. Unless otherwise stated the parameters are $P_0 = 100, \sigma = 10, q = 10, a = 45, \tau = 0.003$

Figure 10 depicts the region of the parameter space where total trading costs for the liquidity trader in the liquidity pool are lower than in the limit order market. The left panel shows that automated market makers are the better trading venue for the liquidity trader when either the innovation in prices or the intensity of informed trading are sufficiently small. In those cases the pool is large and the liquidity trader can trade without much price impact. The right panel shows that for a given intensity of informed trading liquidity pools dominate when the trade size of the liquidity trader is neither to small, because then the fee revenue would be too small, nor too large, because then the pool would grow to a size where it can be picked off to much.

5.2 Empirical Analysis

Our framework does not incorporate market access fees. However, various papers such as Capponi and Jia (2021) and Barbon and Ranaldo (2021) emphasize gas fees as an important determinant of UniSwap market quality. Fully quantifying the difference between on-chain markets (such as Uniswap) and off-chain markets (such as Binance) is complex. First, a token on-chain has different attributes than a token off-chain. Second, off-chain venues have their own market access fees and in some case make-take fees.

Binance, like all 'centralized' crypto exchanges retains custody of the traded assets. Uniswap, in contrast, is non-custodial, meaning that the user can initiate a trade directly out of their personal wallet and keep custody of traded assets until they are swapped in an atomic transaction. Binance based tokens have to be transferred out of Binance custody into a user's wallet before they can be used in other protocols. Similarly, any tokens that a user wants to trade on the exchange must be transferred out the user's personal wallet into the exchange wallet and the exchange has to give the user credit for these assets in their own internal ledger before they can be traded. By contrast, Uniswap traded tokens can be efficiently used in other DeFi protocols, for yield farming or as collateral. In short, owning a token on Binance has a different use value to owning a token on chain. The difference in value should account for the opportunity cost of being off-chain and having access to the different protocols, and the direct cost of being exposed to Binance credit risk.

To compare equivalent token positions we evaluate a swap transaction to two tokens transfers, which would be required to move a token in and out of a centralized exchange. To measure gas costs we pick a sample of 200,000 blocks staring from block 11,668,958. Gas prices are recorded per transaction and each transaction can contain multiple swaps or multiple transfers or even other interactions with DeFi platforms. To generate a conservative estimate, we pick the 10 lowest gas amounts for transactions that contain swaps or transfers, respectively. These transactions are most likely to contain a 'minimalist' transaction that is, for example, just a simple token transfer. We manually verify on Etherscan that this is indeed the case. In these data, a Uniswap V2 swap uses on average 49.973 units of gas and a transfer uses 19.830 units of gas.[9] A swap therefore uses on average 10.313 gas units more than two token transfers.

Market conditions might make it very lucrative to use one DeFi protocol over another one at a given time or users might have different preferences to get their transaction included in a timely manner. To best capture preferences of Uniswap users, on each day we compute the median gas price for all Uniswap transactions, multiply that with the incremental gas amount of 10.312 gas units, and convert that to USD. Figure 11 illustrates our findings. The average incremental fee is 61 cents, the median incremental fee is 4.22 cents. On some days the incremental fee can be very high. Indeed, the maximum in our sample is USD 11.63 on May 11, 2021 when Eth first broke through USD 4,000. For the second part of our sample, in 2021, the median incremental fee was USD 1.88. Finally, observe that our calculation does not include any withdrawal fees that an exchange may charge. Binance, for example, charges token specific fees that are hard to include in a comparison. Binance charges a fee of 0.005 Eth for Eth withdrawals to the Ethereum blockchain, which at the time of writing the paper would be equivalent to USD 21.24.

In addition to market access fees, price level and price volatility are important for liquidity demanders. To compare Uniswap to traditional exchanges we collect minute by minute trading data from Binance, one of the largest crypto-exchanges by volume. Many of the 1,251 token

[9]We observe some very small variations in gas amounts due to minor differences in the implementation of ERC-20 token contracts.

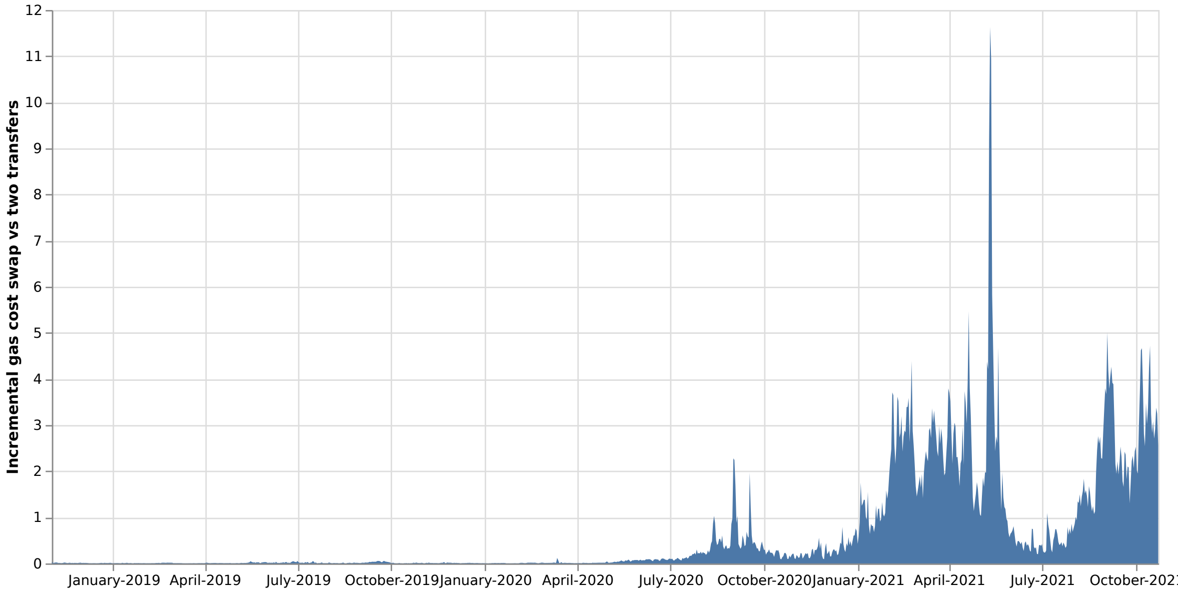

Figure 11. Incremental Fees of a swap transaction over two token transfers measured in USD.

pairs listed on Binance trade against fiat currencies. We find 384 token pairs that trade on both, Uniswap and Binance, however many of these pairs are very infrequently traded. We eliminate all pairs with an average daily volume of less than 10 ETH on either market and an average daily Uniswap pool size of less than 10 ETH. We treat WETH and ETH as identically given the easy and cheap conversion. We end up with 27 token pairs that are cross listed on Uniswap and Binance.

Pricing differences between Uniswap and Binance are small except in the startup-phase of the Uniswap pool when liquidity is scarce. Figure 12 shows the pricing in blue and pool size (in orange) for the USDC/ETH pool. When the pool starts, as long as the poolsize is below 100 ETH, pricing errors are huge reaching over 40%. This is not surprising as a small invariant k will cause a very steep bonding curve (see equation 1). Once the poolsize is above 700 ETH, the pricing difference stays below 1% with an average of -0.026% for this pool.

We examine determinants of price differences between Binance and Uniswap for the broader sample in Table 3. We examine the absolute percentage pricing error defined as the absolute value of the price differential between Binance and Uniswap divided by the price on Binance. Pricing error is lower for large pools, which are the ones that have more liquidity and are also the more commonly traded tokens. When fx-volatility is high, arbitrageurs find it harder to keep keep up with price changes and we see prices diverging across markets. Higher volume, measured here as volume on Binance is associated with smaller price differences. On Uniswap liquidity is defined by the poolsize and independent of trading activity, on Binance, however, higher liquidity is trading volume. Relative volume is defined as volume on Uniswap over combined volume. A negative coefficient on relative volume and a positive coefficient on squared relative volume indicates that the pricing error is u-shaped in relative volume, i.e. it is high when

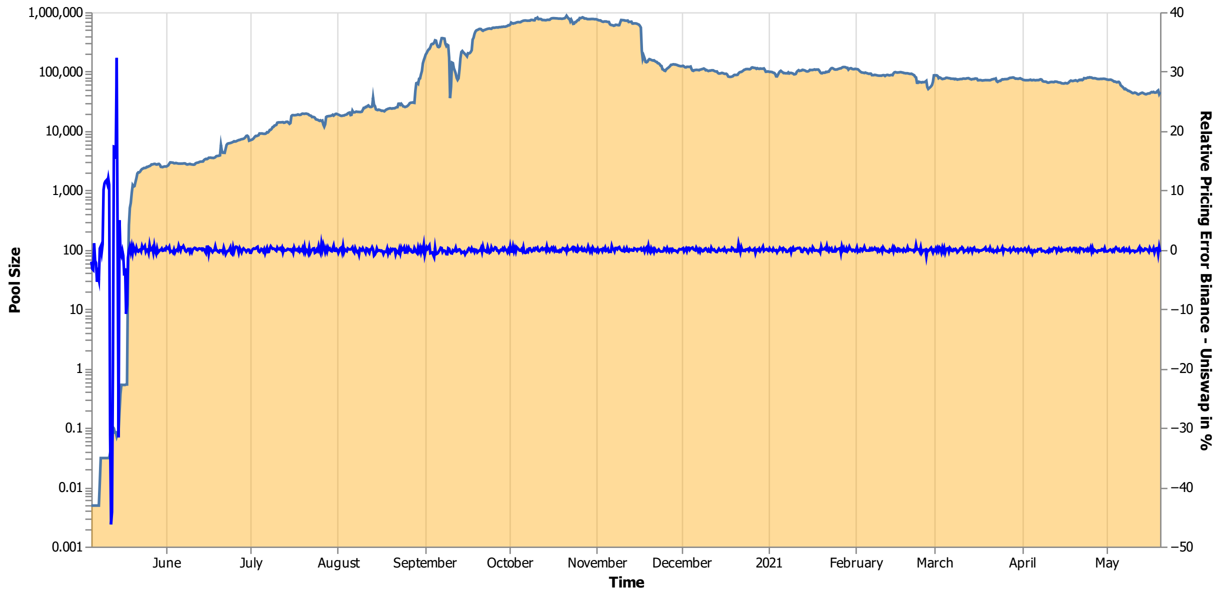

Figure 12. Pricing error and pool size Pricing difference for the USDC/ETH pair when comparing Binance to Uniswap in percent of the Binance price (blue line, right axis) and pool size of the Uniswap USDC/ETH pool (orange, log-scale, left axis).

most trading activity is concentrated on one exchange and lower when both exchanges have a somewhat even share of trading volume. Pricing errors are larger for tokens with very low prices relative to ETH. This might be similar to a penny stock effect, as some of the tokens trade at prices with four or five leading zeros. Traders might not realize that a price difference of 0.00001 ETH can be a huge percentage difference. Price differences also increase in gas prices. When mining costs, i.e. gas prices, are high small price differences are not profitable to arbitrage away.

In columns (8) and (9) of Table 3 we examine price impact as an explanatory variable. We define price impact as the absolute price change over trading volume computed over one minute intervals. For the regression we average the price impact measures on a daily basis. We have to control for pool size as price impact, i.e. the curvature of the bonding curve, is in Uniswap mechanically related to pool size. We find that higher price impact is associated with higher price differences between exchanges, which is intuitive as price impact reduces profits for arbitrageurs. Figure 13 shows our measures for price impact for the USDC/ETH pair. We can see that price impact on Binance almost always exceeds that on Uniswap. Binance's price impact also varies a lot over time while the price impact on Uniswap stays pretty much constant. We also compute the theoretical price impact for Uniswap which we derive analytically from the bonding curve in Equation (1) assuming zero fees. We find that our analytical measure of price impact on Uniswap (green line) corresponds closely to our empirical measure (orange line).

Figure 14 shows the trading volume of the USDC/ETH pair on Binance and Uniswap, respectively. We can see that trading volume is remarkably correlated across the two markets, which is surprising given that tokens have to be moved back and forth trough on-chain transactions

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Pool size	-1.185***						-0.145	-1.263***	0.0871
	(0.220)						(0.138)	(0.161)	(0.155)
Volatility		0.266***					0.182**		0.182**
		(0.0660)					(0.0676)		(0.0698)
Volume Binance			-0.691***				-0.624***		-0.617***
			(0.0951)				(0.107)		(0.120)
Relative Vol Uniswap				-3.895***			-2.912***		-3.310***
				(0.731)			(0.661)		(0.761)
(Relative Vol Uniswap)^2				3.020***			2.427***		2.720***
				(0.692)			(0.679)		(0.786)
Binance Price					-0.0191***		-0.00345		-0.0332***
					(0.00305)		(0.00448)		(0.00979)
Gas Price						1.687**	1.223**		1.595***
						(0.670)	(0.466)		(0.471)
Binance Price Impact								311.2	1021.9***
								(243.7)	(256.4)
Uniswap Price Impact								2.147***	1.885***
								(0.0520)	(0.0869)
R^2	0.0143	0.109	0.00687	0.0842	0.00289	0.0101	0.152	0.0311	0.183
Observations	4,255	4,189	4,255	4,255	4,255	4,255	4,189	4,240	4,174

Table 3: Regression explaining absolute pricing errors between Uniswap and Binance markets. Absolute pricing error is the absolute daily median of the minute by minute price differences between Binance and Uniswap observed prices in percent, Pool size is the size of a liquidity pool in million ETH, Volatility is the daily standard deviation of block by block price changes of the price in a pool in percent. Volume Binance is the daily volume on Binance in million ETH, Relative Vol. Uniswap is the fraction of combined volume that is observed on Uniswap, Binance Price is the median daily observed price on Binance, Gas price is the median daily gas price for uniswap transactions in TWEI ($10^1 2$ Wei). Standard errors are clustered by pool. One, two, and three stars indicate significance at the 10%, 5%, and 1% level, respectively.

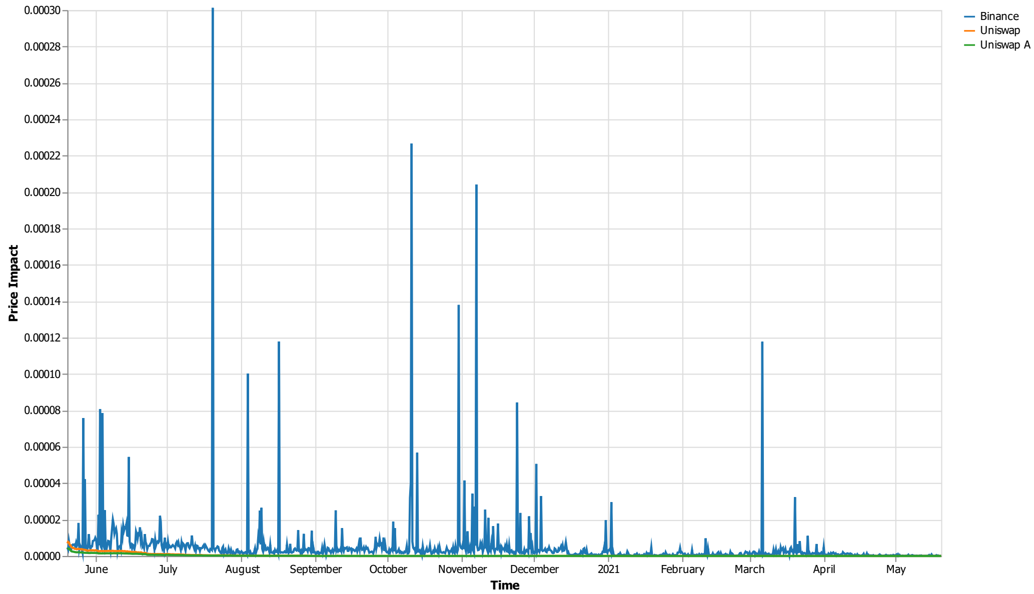

Figure 13. Price Impact of USDC/ETH on Uniswap (orange, green) and Binance (blue). Price impact is computed as change in price over volume (green and blue lines) as well as analytically as the price change for a marginal unit bought using the bonding curve formula (green line).

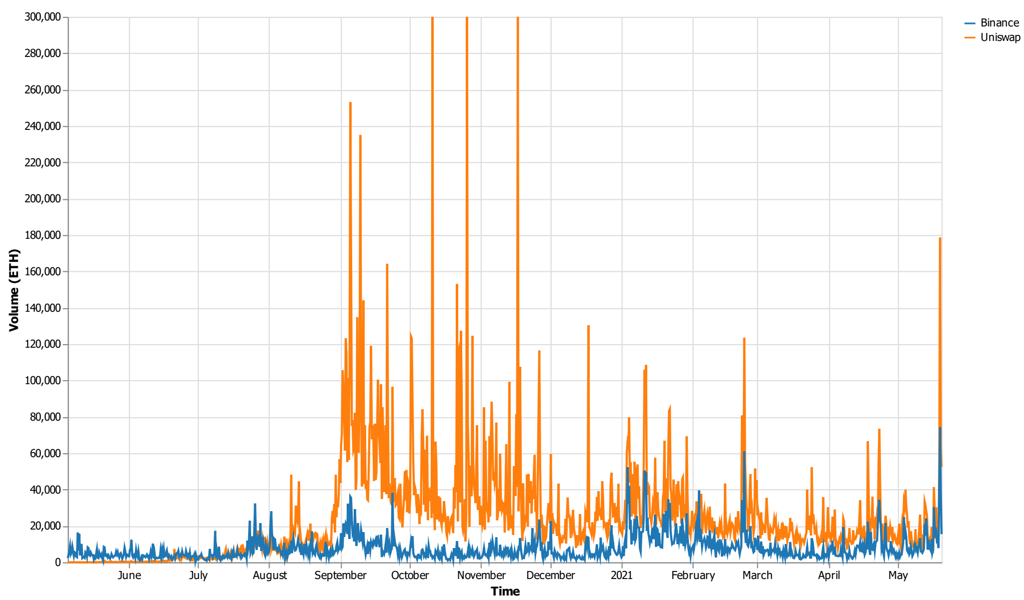

Figure 14. Trading volume of USDC/ETH on Uniswap (orange) and Binance (blue). The graph shows the trading volume excluding flash loans in ETH. Trading volume is aggregated over rolling eight hour intervals.

between the two markets. We can also see that Uniswap is gaining market share over time and eventually more trading is happening on Uniswap relative to Binance.

Figure 15. Intraday prices for the USDC/ETH pair on October 21, 2020 The graph shows minute-by-minute prices of the USDC/ETH pair on Binance and Uniswap.

Figure 15 shows intraday prices of the USDC/ETH pair on one day, October 21, 2020. The patterns is typical for most days in our sample. It seems that often Binance prices are leading Uniswap prices. and that Binance prices are more volatile that the prices on Uniswap.

6 Conclusion

In 1971, Fischer Black wrote two articles for the Financial Analyst's Journal speculating on whether computers or "automation" could ever replace human interaction in financial markets (Black (1971a), Black (1971b)). In these papers, he argued that market liquidity was constrained by the size of a market maker's inventory and suggested that a solution would be to have more direct participation from other market participants.

The Uniswap experiment does not merely increase the supply of liquidity by relaxing market makers' inventory constraints. It also changes how the benefits and costs of liquidity provision are shared among market participants. In a limit order market, the exchange sets price and time priority, which effectively determine the interaction of liquidity demand and supply. Price impact in a limit order market is the endogenous outcome of the interaction. By contrast, on an AMM the price impact is programmatically determined by a a bonding curve. However, we note that this impact is conditional on the size of the pool. Thus, the equilibrium price impact arises as the pool size adjusts, so that liquidity suppliers trade off potential adverse selection against fee revenue.

The automated Uniswap protocol has clearly been successful. Further, in July 2021, Swarm a Berlin based DeFi company announced that it was launching an AMM that was fully licensed by the German regulator, BaFin. One of the reasons for the uptake of these liquidity sharing protocols is that we have demonstrated both theoretically and empirically that pools adjust to tradeoff the benefits and costs of liquidity provision. Further, compared with a centralized exchange there are some token pairs for which the AMM provides liquidity more efficiently than a centralized exchange.

We note that Uniswap has recently introduced a V3. Their re-design will give liquidity suppliers partial ability to associate their liquidity to price ranges. This change re-introduces competition between liquidity suppliers. In as much as the AMM is effective if it reduces competition between liquidity suppliers, these changes may drive out non-strategic liquidity suppliers.

References

Angeris, Guillermo, and Tarun Chitra, 2020, Improved Price Oracles: Constant Function Market Makers, Working Paper.

Angeris, Guillermo, Hsien-Tang Kao, Rei Chiang, and Charlie Noyes, 2019, An Analysis of Uniswap markets, Working Paper.

Aoyagi, Jun, 2020, Liquidity Provision by Automated Market Makers, Working Paper.

Barbon, Andreas, and Angelo Ranaldo, 2021, On The quality of Cryptocurrency Markets, working paper.

Baruch, Shmuel, and Lawrence R. Glosten, 2013, Fleeting Orders, working paper.

Biais, Bruno, Thierry Foucault, and Sophie Moinas, 2015, Equilibrium Fast Trading, Journal of Financial Economics 116,No2., 292–313.

Black, Fisher, 1971a, Toward a Fully Automated Stock Exchange-Part 1, Financial Analysts Journal 27, 28–35.

Black, Fisher, 1971b, Toward a Fully Automated Stock Exchange, Part 2, Financial Analysts Journal pp. 86–87.

Brogaard, Jonathan, Terrence Hendershott, and Ryan Riordan, 2019, Price Discovery without Trading: Evidence from Limit Orders, Journal of Finance 74, No4., 1621–1161.

Capponi, Agostino, and Ruizhe Jia, 2021, The Adoption of Blockchain-based Decentralized Exchanges, working paper.

Glosten, Lawrence R., 1994, Is the Electronic Open Limit Order Book Inevitable?, Journal of Finance 49, No4., 1127–1658.

Glosten, Lawrence R., 1998, Precedence Rules, working paper.

Hao, Jia, Avner Kalay, and Stewart Mayhew, 2009, Ex-Dividend Arbitrage in Option Markets, The Review of Financial Studies 23(1), 271–303.

Haynes, Richard, and Esen Onur, 2020, Precedence Rules in Matching Algorithms, Journal of Commodity Markets 19.

Ito, Yuki, and Jun Aoyagi, 2021, Co-Existing Exchange Platforms, working paper.

Park, Andreas, 2021, The Conceptual Flaws of Constant Product Automated Market Making, Available at SSRN 3805750.

A Measuring volume

Dollar volume is a natural volume measure on US exchanges, and there is typically a recognized common price. On Uniswap, the US dollar is not the numeraire; most of the trading in the Uniswap system is against WETH or USD stablecoins. Further, the existence of flash loans and flash swaps during so called oracle attacks mean that there can be large price discrepancies across trading venues.

To compute volume, we take the WETH part of the trade and convert it to USD using Binance minute by minute data. We use the Binance price because the internal UniSwap price is sometimes purposely distorted. As we have indicated, large trades such as oracle attacks can push relative token prices out of equilibrium. Most of the remaining pools trade against a USD stablecoin. For those pools we convert stablecoin to USD. For all remaining pools we search for all pools where one of the tokens trades against WETH and convert using the prices from the pool with the highest volume. We then use external data from Binance to convert the ETH to USD.

As a concrete example of price distortions caused by oracle attacks, a December 17, 2020 trade against Warp Finance is in our sample.[10] The initiator borrowed about 200 million USD in flash loans from Uniswap and dYdX to manipulate the DAI/ETH price by trading 48.58 million DAI against 342,252 WETH. DAI is a USD stablecoin so one side of the trade roughly corresponds to 48 million USD. One ETH was worth about USD 643 at the time, making the WETH side of the trade worth about 220 million USD. The relatively complex trade involved depositing DAI and WETH into a liquidity pool, then using the liquidity tokens as collateral in Warp Finance. Finally, using a large trade to distort the DAI WETH price which led the Warp code to overvalue the deposited collateral. Smart contract nuances aside, this trade was designed to distort the DAI/WETH price which means that using contemporaneous prices may distort the volume. Thus in the example above we would measure the volume as 220 million USD. Our approach does not systematically inflate volume because oracle attacks can happen in both directions and all trades that push prices out of equilibrium also have an opposing trade that brings prices back to equilibrium.

Many other contracts such as lending platforms rely on decentralized exchanges as price feeds or oracles to determine, for example, the value of the collateral in relation to the face value of an outstanding loan. Attackers can exploit poorly written code of such lending platforms for financial gain. Typically large trades are used to move prices in the bonding curve market that the lending contract uses as oracle, making the smart contract believe that the collateral is very valuable. Then the attacker borrows against the collateral, brings the price on the bonding curve market back to equilibrium, and walks away from the, now under-collateralized, loan. One such attack happened on Harvest finance, a yield farming cooperative similar to a floating NAV money market fund. User could deposit tokens with Harvest finance in return for fAsset tokens (e.g. depositors of USDC receive a fUSDC token). The underlying tokens were invested in high yielding liquidity pools and the revenues are shared with the holders of fAsset tokens.

On October 26 an attacker borrowed 18 million USDT and 50 million USDC on Uniswap and

[10] See transaction 0x8bb8dc5c7c830bac85fa48acad2505e9300a91c3ff239c9517d0cae33b595090.

converted over 17 million USDT to USDC on curve.fi (an automated market maker similar to Uniswap that specialized in trading stablecoins).[11] This temporarily increased the price of USDC in the curve.fi pool, which is used a price oracle for Harvest finance. The attacker then deposited USDC with Harvest, in exchange for an inflated number of fUSDC tokens. Specifically, the price manipulation caused the price of fUSDC to temporarily drop to 0.9712 from 0.98 before the attack. Then the attacker changed 17 million USDC back to USDT on curve.fi and sold his fUSDC tokens back to Harvest finance at 0.9833 as Harvest finance's smart contract updated the price based on the new information from curve.fi. The net profit of this attack was 619,408 USDC. The hacker then repeated the process 17 times and also attacked other Harvest finance pools for a total profit of USD 24 million.[12]

Data Issues:

We also observe some transactions that trade tokens back and forth without apparent reason. For example in one transaction on March 30, 2021 somebody created a new Uniswap exchange for a token named SCAMMY, borrowed WETH worth 220 million USD in a flash loan on dYdX, injected half as liquidity to the pool, and then traded 50 times the other half of the funds back and forth for a total volume of USD 5 billion.[13] The trader then withdrew the liquidity and repaid the flashloan. There is no obvious profit motive for this transaction. One possibility is that the trader tried to generate high fee revenue to place the token in a leading position at one of the yield farming websites in order to attract investment to this scam token (although naming the token scammy is not helpful for this purpose). We include such events in graphs and summary statistics. For our econometric analysis we winsorize the data and thus eliminate such outliers.

Ticker symbols on Uniswap are not protected. Anyone can create a token and assign the ticker symbol of a popular token like WETH or USDC. Tokens are uniquely identified by their address, e.g. 0xc02aaa39b223fe8d0a0e5c4f27ead9083c756cc2 for WETH, which is not easy to work with. Most people therefore use tickers and are exposed to copycat tokens. Table 4 lists fraudulent versions of popular tokens. We can see that, for example, Yearn Finance (YFI) has 17 copycat tokens that use the same ticker symbol. A total of 328 transactions were done on Uniswap with copycat tokens which is small compared to 257,728 transactions in the legitimate token. Overall we find that there amount of trading in fake tokens is small and will not affect our findings.

B Numerical Example of an AMM

Assume that the fair exchange rate for a token is 10 ETH/token and a sole liquidity provider contributed $E = 100$ ETH and $T = 10$ tokens to the liquidity pool for which he gets 100 liquidity tokens in return. Suppose that the fee is $\tau = 0.003$.

[11] see transaction 0x35f8d2f572fceaac9288e5d462117850ef2694786992a8c3f6d02612277b0877.

[12] See 'Harvest Flashloan Economic Attack Post-Mortem', medium.com, Oct 26, 2020. https://medium.com/harvest-finance/harvest-flashloan-economic-attack-post-mortem-3cf900d65217

[13] See transaction 0xa8c00a56cf2455241bbc4b5ef9e3f9e761cdbd7909847ab8274dcd9bd1dded6a.

Ticker	Number of fraudulent tokens	Fraudulent transactions	Nonfraudulent transactions
WETH	4	5	30,624,081
USDT	17	91	3,062,694
USDC	6	78	2,977,465
DAI	4	7	1,967,084
UNI	14	463	550,393
AMPL	3	15	380,315
WBTC	7	5	350,009
LINK	9	59	349,431
HEX	4	6	271,709
YFI	17	328	257,728
SNX	7	2,388	232,924
MKR	1	332	230,044
SUSHI	6	90	217,765
SAI	3	22	184,963
KP3R	18	227	184,644

Table 4. Fraudulent tokens. The table shows the number of fraudulent tokens, i.e. tokens with the same ticker symbol as popular tokens but with a different address. Number of fraudulent tokens is the number of fraudulent tokens found as part of a Uniswap liquidity pool. Fraudulent transactions are the number of transactions in liquidity pools with these fraudulent tokens. Non-fraudulent transactions is the number of transactions in the original token in Uniswap pools.

Example 1 A trader wants to buy tokens for $e = 10$ ETH. He gets $\frac{0.997eT}{0.997e+E} = \frac{0.997 \cdot 10 \cdot 10}{0.997 \cdot 10+100} = 0.90661$ tokens in return. The pool collects a fee of $0.003e = 0.003 \cdot 10 = 0.03$ ETH. The new token balance post trade is $10 - 0.90661 = 9.09339$ tokens.

The post trade ETH balance equals the old balance plus what the trader gave for tokens plus the fee revenue $100 + 0.997e + 0.003e = 100 + 9.97 + 0.03 = 110$. The average price the trader got is $p = \frac{e}{T} + \frac{E}{(1-\tau)T} = \frac{10}{10} + \frac{100}{(1-0.0003)10} = 11.0301$

Note that the invariant k as defined in Equation 1 is the same pre and post trade only without fees, i.e. $10 \cdot 100 = 9.09339(100+0.997 \cdot 10) = 1000$. Because the fee gets credited to the liquidity pool after the trade, the invariant increases to $9.09339 \cdot 110 = 1000.27$. The next trade will be priced based on this new invariant. The new mid-price is $p^0 = 12.0967$. In response to a buy order, the mid-price moved up.

When redeeming her liquidity tokens, the liquidity provider would receive whatever is in the pool, which is now 110 ETH and 9.09339 token.

Consider two cases:

(a) True price is 10 ETH/token Had she kept her initial investment of 100 ETH and 10 tokens in a private wallet it would now be worth $100 + 10 \cdot 10 = 200$ ETH.

When she redeems the liquidity token, she would obtain a total of $110 + 9.09339 \cdot 10 = 200.9339$ ETH and makes a profit of 0.9339 ETH. This is the sum of the trading fee ($\tau e = 0.003 \cdot 10 = 0.03$) and the gain from selling to the trader at an average price above the true price.

(b) True price is 12.0967 ETH/token Had she kept her initial investment of 100 ETH and 10 tokens in a private wallet it would now be worth $100 + 10 \cdot 12.0967 = 220.0967$

When she redeems the liquidity token she gets $110 + 9.09339 \cdot 12.0967 = 220$. She loses 0.0967 ETH, in which the gain from the trading fee is more than offset by the loss from the exchange selling tokens at stale prices.

As with any passive liquidity provider, the Uniswap pools present a free option to the market. That is if the quantities in the pool are such that the terms of trade differ from the true value, arbitrageurs are more likely to pick off stale liquidity. This logic is reflected in the previous example. However, liquidity demanding trades are always valuable. The liquidity suppliers receive a fee for the liquidity demanding order and also receive a fee when equilibrium is replenished by arbitrage traders. They only face potential losses if there has been a permanent value change in the token.

If the price of a token moves away from the fundamental value because of a large order, an arbitrageur will initiate an offsetting trade and bring the mid-price of the exchange back to the fundamental value. Such short term deviations from the fundamental price of a token are beneficial to liquidity holders. In a pool without fees liquidity-providers will gain zero on such a trading pattern. Trades are always priced in such a way that the amount of ETH and tokens are on the bonding curve before and after any trade (see Equation 1). Thus a move from (E, T) to (E', T') and then back to (E, T) will leave the liquidity providers at exactly the same point they started from. Many crypto-traders refer to gains or losses while the pool is off equilibrium at value (E', T') as impermanent loss. With positive fees liquidity traders benefit from such short term deviations as they collect a proportional fee for both trades. In our empirical analysis we will estimate such short term deviations from a fundamental value as reversals.

Finally, even though arbitrageurs replenish the liquidity pool after a large, they pay a fee for doing it. This differs from a traditional limit order book in which liquidity is replenished by rivalrous liquidity suppliers.

Example 2 Continue Example 1, case (a) and suppose that an arbitrageur brings back the price closer to the fundamental value. Assume that the arbitrageur can buy tokens at the fundamental value of 10, sells them to the pool at price $p(t)$, and chooses the optimum amount of tokens t to sell to the pool to maximize profit, $\pi = t(p(t) - 10)$. By sending t tokens to the pool he will obtain $e = \frac{(1-\tau)tE'}{T'+t(1-\tau)} = \frac{0.997 \cdot t \cdot 110}{9.09339+0.997t}$ ETH in return, resulting in a price $p(t) = e/t$. Solving for the optimal t that maximizes the arbitrageur's profit we find $t = 0.895648$ which is smaller then the amount of token sold by the pool in Example 1 because (i) the invariant k has changed after the first trade due to the fee revenue and (ii) fees make it optimal for the arbitrageur to sell a smaller amount back to the pool. It is easy to verify that without fees, i.e. $\tau = 0$, the invariant does not change after the first trade and the arbitrageur would sell exactly the same amount of tokens back to the pool that the pool sold in the previous trade, and the new mid-price of the pool would exactly equal the fundamental value. With fees, however, the arbitrageur optimally sells $t = 0.895648$ tokens to the pool for which he receives 9.836 ETH, leaving the pool with a new balance of 100.164 ETH and 9.98904 token. The new pool mid-price is 10.0274, which deviates slightly from the fundamental value of 10. Liquidity providers value their token holdings at the fundamental value of 10 and hold a total of $100.164 + 99.8904 = 200.054$ which is higher than their initial investment of 200.

C Proofs

Proof of Lemma 1

If the liquidity supplier is alone in the market, and posts at prices $\underline{p}$ at which he buys and $\bar{p}$ at which he sells, with $\underline{p} < \bar{p}$ then

i. With probability $(1-\alpha)$ a noise trader arrives. If the noise trader is a buyer, the liquidity supplier sells to him at $\bar{p}$ and obtains a payoff of $(\bar{p} - p_0)$. Symmetrically, if the noise trader is a seller, the liquidity supplier buys from him at $\underline{p}$ and obtains a payoff of $(p_0 - \underline{p})$.

ii. With probability α there is an information event. If the informed trader buys, the liquidity supplier obtains a payoff $\bar{p} - (p_0 + \sigma)$ and if the informed trader sells, he obtains a payoff of $(p_0 - \sigma) - \underline{p}$.

His overall profit is then

$$(1-\alpha)\left[\frac{1}{2}(\bar{p} - p_0) + \frac{1}{2}(p_0 - \underline{p})\right] + \alpha\left[\frac{1}{2}(\bar{p} - (p_0 + \sigma)) + \frac{1}{2}((p_0 - \sigma) - \underline{p})\right] \tag{17}$$

Clearly, he will post a sell price at which to sell of $\bar{p} = p_0 + \sigma$, and a price at which to buy of $\underline{p} = p_0 - \sigma$, to obtain a profit of $(1-\alpha)\sigma$.

■

Proof of Lemma 2

A limit order submitter in competition chooses a sell price p_i to maximize his expected profits, which comprises:

i. With probability $(1-\alpha)$ the noise trader arrives. Limit order i gets his sell order filled with probability $\frac{1}{2}(1 - F_j(p_i))$ and obtains a payoff of $(p_i - p_0)$

ii. With probability $\frac{\alpha}{2}$ there is an information event in which the asset value jumps up. Trader i will trade with the informed trader for a payoff of $(p_i - p_0 - \sigma) \leq 0$.

The expected profit for the liquidity provider upon posting a sell order with price p_i is then

$$\pi_i(p_i) \quad = \quad \frac{(1-\alpha)}{2}(1 - F_j(p_i))(p_i - p_0) + \frac{\alpha}{2}(p_i - p_0 - \sigma).$$

In equilibrium, it has to be that each price is offered with some probability and it is not optimal to deviate from that price. So, the first order condition for any optimal price satisfies

$$\frac{\alpha}{2} + \frac{(1-\alpha)}{2}(1 - F_j(p)) - \frac{(1-\alpha)}{2}(p - p_0)F_j'(p) = 0. \tag{18}$$

We can solve the differential equation in (18) with the boundary condition $F^s(p_0 + \sigma) = 1$, to get the symmetric equilibrium schedule:

$$F^s(p) = \frac{(p - p_0) - \alpha\sigma}{(p - p_0)(1 - \alpha)}.$$

The minimum price p^s_{min} that the limit order submitters are willing to offer can be solved from $F(p^s_{min}) = 0$, to obtain

$$p_{min} = p_0 + \alpha\sigma \tag{19}$$

By symmetry the schedule for the buy orders is given by

$$F^b(p) = \frac{(p_0 - p) + \alpha\sigma}{(p_0 - p)(1 - \alpha)},$$

where $p^b_{min} = p_0 - \sigma$ and $p^b_{max} = p_0 - \alpha\sigma$

■

Proof of Proposition 3

The ex ante profit for the limit order trader i is

$$\gamma_i(1 - \gamma_j)(1 - \alpha)\sigma q - I(\gamma_i)$$

The reaction function is defined by

$$(1 - \gamma_j)(1 - \alpha)\sigma q - 2a\gamma_i = 0$$

In symmetric equilibrium, each liquidity supplier will choose a private level of investment

$$\gamma^* = \frac{(1 - \alpha)\sigma q}{2a + \sigma q(1 - \alpha)}$$

$\blacksquare$

Proof of Lemma 4

The fees accruing to the liquidity providers from accommodating the liquidity trader follows from the text, and are τe_ℓ^s and τe_ℓ^b,

where

$$\begin{aligned} e_\ell^b &= \frac{E_0 T_0}{T_0 - q} - E_0 \\ e_\ell^s &= E_0 - \frac{E_0 T_0}{T_0 + q}. \end{aligned}$$

The arbitrageur reverses the trade and pays the same fee in Eth. The overall expected fee revenue is then

$$\begin{aligned} &= 2\tau \left(\frac{1}{2} \frac{E_0 T_0}{T_0 - q} - \frac{1}{2} \frac{E_0 T_0}{T_0 + q} \right) \\ &= \tau E_0 T_0 \left(\frac{T_0 + q - (T_0 - q)}{(T_0 - q)(T_0 + q)} \right) \\ &= 2\tau p_0 T_0^2 \left(\frac{q}{T_0^2 - q^2} \right), \end{aligned}$$

where the last line uses the fact that $p_0 = \frac{E_0}{T_0}$.

Reversing the trade is profitable for the arbitrageur if the AMM price including transaction costs differs from the fundamental value or

$$\begin{array}{lll} e_\ell^b(1-\tau) \geq p_0 & \text{or } \dfrac{T_0 q}{(T_0 - q)} > \dfrac{1}{1-\tau} & \text{to reverse a liquidity buy order} \\ e_\ell^s(1+\tau) \leq p_0 & \text{or } \dfrac{T_0 q}{(T_0 + q)} < \dfrac{1}{1+\tau} & \text{to reverse a liquidity sell order,} \end{array}$$

which yields the conditions,

$$\begin{array}{ll} \tau < 1 - \dfrac{T_0 - q}{T_0 q} & \text{To reverse the liquidity buy} \\ \tau < \dfrac{T_0 + q}{T_0 q} - 1 & \text{To reverse the liquidity sell} \end{array}$$

Define $\tau^* = \min \left[1 - \frac{T_0 - q}{T_0 q} \, , \, \frac{T_0 + q}{T_0 q} - 1 \right]$. Then, for τ, τ^*, this trade is profitable for the arbitrageur. $\blacksquare$

Proof of Lemma 5

The change in value to the liquidity providers if there is a positive asset innovation follows directly from the arguments in the text.

If there is a negative asset innovation, the post innovation value of the committed capital is $E_0 + T_0(p_0 - \sigma) = 2E_0 - \sigma T_0$. If there is a negative asset innovation, then the informed trader sells up to the point of no arbitrage, and the trade occurs on the bonding curve, so

$$
\begin{aligned}
(E_0 - e_I^s)(T_0 + t^s) &= E_0 T_0. && (20)\\
e_I^s &= E_0 - \frac{E_0 T_0}{T_0 + t^s} && (21)
\end{aligned}
$$

The profits of the informed trader are given by

$$
\begin{aligned}
\pi_I^s &= (1-\tau)e_I^s - (p_0 - \sigma)t^s && (22)\\
&= (1-\tau)\left(E_0 - \frac{E_0 T_0}{T_0 + t^s}\right) - (p_0 - \sigma)t^s. && (23)
\end{aligned}
$$

The first order condition is

$$
(1-\tau)\frac{E_0 T_0}{(T_0 + t^s)^2} - (p_0 - \sigma) \quad = \quad 0. \tag{24}
$$

Note, this represents optimal trading if $t^s = \sqrt{\frac{(1-\tau)E_0 T_0}{p_0 - \sigma}} - T_0 \geq 0$, or $\frac{\sigma}{p_0} \geq \tau$.

Thus, we obtain

$$
\begin{aligned}
t^s &= \sqrt{\frac{(1-\tau)E_0 T_0}{p_0 - \sigma}} - T_0\\
e_I^s &= E_0 - \sqrt{\frac{E_0 T_0 (p_0 - \sigma)}{1-\tau}}.
\end{aligned}
$$

Hence, the value of the committed capital after the asset value has dropped and the informed trader has traded is

$$
E_0 - e_I^s + (T_0 + t^s)(p_0 - \sigma) \quad = \quad \left(\frac{2-\tau}{\sqrt{1-\tau}}\right)\sqrt{E_0 T_0 (p_0 - \sigma)}.
$$

Thus, the change in total value of the committed capital is

$$
\left(\frac{2-\tau}{\sqrt{1-\tau}}\right)\sqrt{E_0 T_0 (p_0 - \sigma)} - 2E_0 + \sigma T_0. \tag{25}
$$

The liquidity cost paid by the informed trader is $\tau\left(E_0 - \sqrt{\frac{E_0 T_0 (p_0 - \sigma)}{1-\tau}}\right)$, so the total payoff to liquidity provision if the asset value drops is

$$\left(\frac{2-\tau}{\sqrt{1-\tau}}\right)\sqrt{E_0T_0(p_0-\sigma)}-2E_0+\sigma T_0+\tau\left(E_0-\sqrt{\frac{E_0T_0(p_0-\sigma)}{1-\tau}}\right) \tag{26}$$

$$= \left(\frac{2-\tau}{\sqrt{1-\tau}}\right)\sqrt{T_0^2p_0(p_0-\sigma)}-2p_0T_0+\sigma T_0+\tau\left(T_0p_0-\sqrt{\frac{T_0^2p_0(p_0-\sigma)}{1-\tau}}\right) \tag{27}$$

$$= 2\sqrt{1-\tau}\sqrt{T_0^2p_0(p_0-\sigma)}-(2-\tau)p_0T_0+\sigma T_0 \tag{28}$$

■

Proof of Proposition 6

From Lemmas 4 and 5, the overall payoff to liquidity provision (for the entire pool) and using the fact that $p_0 = \frac{E_0}{T_0}$ is

$$(1-\alpha)2\tau p_0T_0^2\left(\frac{q}{T_0^2-q^2}\right)$$
$$+\ \frac{\alpha}{2}\left(2T_0\sqrt{(1+\tau)}\sqrt{p_0(p_0+\sigma)}-(2+\tau)p_0T_0-\sigma T_0\right)+\frac{\alpha}{2}\left(2T_0\sqrt{1-\tau}\sqrt{p_0(p_0-\sigma)}-(2-\tau)p_0T_0+\sigma T_0\right)$$

Or,

$$\frac{(1-\alpha)2\tau p_0T_0^2q}{T_0^2-q^2}+T_0\alpha\left(\sqrt{p_0(p_0+\sigma)(1+\tau)}+\sqrt{(1-\tau)p_0(p_0-\sigma)}-2p_0\right) \tag{29}$$

The liquidity suppliers have deep pockets with a zero opportunity cost of capital, and are indifferent between committing extra value when the value to liquidity provision is zero. Thus, the equilibrium size of the liquidity pool is implicitly defined by setting Equation 29 equal to zero, which yields,

$$0 = (1-\alpha)2\tau p_0T_0\left(\frac{q}{T_0^2-q^2}\right)+\alpha\omega, \tag{30}$$

where

$$\omega = \left(\sqrt{p_0(p_0+\sigma)(1+\tau)}+\sqrt{p_0(p_0-\sigma)(1-\tau)}-2p_0\right)$$

Thus,

$$T_0 = q\left[\sqrt{1+\frac{(1-\alpha)^2\tau^2p_0^2}{\alpha^2\omega^2}} - \frac{(1-\alpha)\tau p_0}{\alpha\omega}\right]. \quad (31)$$

A finite pool trades off the fee revenue against the picking off risk, else pool size is infinite. It also has to be large enough so that $T_0 > q$. Clearly, $T_0 > q$, if and only if $\frac{(1-\alpha)\tau p_0}{\alpha\omega} < 0$, which holds if $\omega < 0$.

Omega reaches a maximum value at zero.

$$\begin{aligned}
\frac{d\omega}{d\tau} &= \frac{\sqrt{p_0(p_0+\sigma)}}{2\sqrt{(1+\tau)}} - \frac{\sqrt{p_0(p_0-\sigma)}}{2\sqrt{1-\tau}}, \\
\frac{d^2\omega}{d\tau^2} &= -\frac{\sqrt{p_0(p_0+\sigma)}}{4(1+\tau)^{3/2}} - \frac{\sqrt{p_0(p_0-\sigma)}}{2(1-\tau)^{3/2}} < 0
\end{aligned}$$

$$\begin{aligned}
\frac{d\omega}{d\tau} &= 0 \\
\implies \tau &= \frac{\sigma}{p_0}.
\end{aligned}$$

At $\tau = \frac{\sigma}{p_0}$, $\omega = 0$. Thus, ω is strictly negative for $\tau \neq \frac{\sigma}{p_0}$

Further, to ensure that arbitrageur finds it profitable to trade, from the previous lemmas, $\tau \leq \frac{\sigma}{p_0}$.

■

Proof of Corollary 1

[i.] Immediate.

[ii.]

$$\frac{dT_0}{d\omega} = q\left[\frac{(1-\alpha)p_0\tau}{\alpha\omega^2} - \frac{(1-\alpha)^2p_0^2\tau^2}{\alpha^2\omega^3\sqrt{\left(\frac{(1-\alpha)^2p_0^2\tau^2}{\alpha^2\omega^2}+1\right)}}\right] \quad (32)$$

Notice that

$$\frac{(1-\alpha)p_0\tau}{\alpha\omega^2} > \frac{(1-\alpha)^2p_0^2\tau^2}{\alpha^2\omega^3\sqrt{\left(\frac{(1-\alpha)^2p_0^2\tau^2}{\alpha^2\omega^2}+1\right)}},$$

so $\frac{dT_0}{d\omega} > 0$. Further,

$$\frac{d\omega}{d\sigma} = \frac{1}{2}\left[\frac{\sqrt{p_0(1+\tau)}}{\sqrt{(p_0+\sigma)}} - \frac{\sqrt{p_0(1-\tau)}}{\sqrt{(p_0-\sigma)}}\right]. \tag{33}$$

$\frac{d\omega}{d\sigma} < 0$ if $\tau < \frac{\sigma}{p_0}$.

[iii.]

$$\frac{dT_0}{d\alpha} = \frac{p_0\tau\left(1+\frac{p_0\tau(1-\alpha)}{\sqrt{p_0^2(1-\alpha)^2\tau^2+\alpha^2\omega^2}}\right)}{\alpha^2\omega} < 0,$$

as $\omega < 0$.

■

Proof of Proposition 7

The Limit Order Market:

If two traders are in the market, and the liquidity trader buys, the distribution over the lower of the two prices is given by

$$\begin{aligned} F_{min}(p) &= 1 - \Pr(p > x) \\ &= 1 - \Pr(p_i > x, p_j > x) \\ &= 1 - \Pr(p_i > x)\Pr(p_j > x) \\ &= 1 - [1 - F^s(p)]^2 \end{aligned}$$

Where the last two lines follow from the fact that the distributions are independent and identical in symmetric equilibrium. Thus, the cumulative distribution of the minimum price is given by

$$F_{min}(p) = 1 - \left(1 - \frac{(p - p_0 - \alpha\sigma)}{(p - p_0)(1-\alpha)}\right)^2$$

Given the distribution of the minimum prices, the expected transaction price (cost to buy) is determined as:

$$
\begin{aligned}
Ec^b(q) &= p_{min} + \int_{p_{min}}^{p_0+\sigma} 1 - F_c(p)dp \\
&= p_{min} + \int_{p_{min}}^{p_0+\sigma} \left(\frac{\alpha(p-p_0-\sigma)}{(p-p_0)(1-\alpha)}\right)^2 dp \\
&= p_{min} \left(\frac{\alpha}{(1-\alpha)}\right)^2 \int_{p_{min}}^{p_0+\sigma} \left(1 - \frac{\sigma}{(p-p_0)}\right)^2 dp \\
&= p_0 + \alpha\sigma + \frac{\alpha}{(1-\alpha)^2}\Gamma(\alpha,\sigma) \\
\Gamma(\alpha,\sigma) &= \sigma\left[1-\alpha^2+2\alpha\ln(\alpha)\right]
\end{aligned}
$$

$E(\Delta_p^{limit}) = \frac{\alpha}{(1-\alpha)^2}\Gamma(\alpha,\sigma) - p_0$. A symmetric expression holds for the other side of the market.

If there is a sole liquidity supplier, he places orders at $p_0 - \sigma$ and $p_0 + \sigma$. The transaction cost is therefore σ.

The probability that a limit order liquidity supplier is a monopolist is $\gamma^*(1-\gamma^*)$, therefore the probability that the passive trader faces a monopolist when there are two liquidity suppliers is $2\gamma^*(1-\gamma^*) = \eta$.

In the AMM:

$$
\begin{aligned}
e_\ell^b &= \frac{E_0 T_0}{T_0 - q} - E_0 \\
&= \frac{p_0 T_0^2}{T_0 - q} - p_0 T_0 \\
&= \frac{p_0\left(q\left[\sqrt{1+\frac{(1-\alpha)^2\tau^2p_0^2}{\alpha^2\omega^2}} - \frac{(1-\alpha)\tau p_0}{\alpha\omega}\right]\right)^2}{q\left[\sqrt{1+\frac{(1-\alpha)^2\tau^2p_0^2}{\alpha^2\omega^2}} - \frac{(1-\alpha)\tau p_0}{\alpha\omega}\right] - q} - p_0\left(q\left[\sqrt{1+\frac{(1-\alpha)^2\tau^2p_0^2}{\alpha^2\omega^2}} - \frac{(1-\alpha)\tau p_0}{\alpha\omega}\right]\right) \\
&= p_0 q\left[\frac{\left[\sqrt{1+\frac{(1-\alpha)^2\tau^2p_0^2}{\alpha^2\omega^2}} - \frac{(1-\alpha)\tau p_0}{\alpha\omega}\right]^2}{\sqrt{1+\frac{(1-\alpha)^2\tau^2p_0^2}{\alpha^2\omega^2}} - \frac{(1-\alpha)\tau p_0}{\alpha\omega} - 1} - \left[\sqrt{1+\frac{(1-\alpha)^2\tau^2p_0^2}{\alpha^2\omega^2}} - \frac{(1-\alpha)\tau p_0}{\alpha\omega}\right]\right] \\
&= p_0 q\left[\frac{\sqrt{1+\frac{(1-\alpha)^2\tau^2p_0^2}{\alpha^2\omega^2}} - \frac{(1-\alpha)\tau p_0}{\alpha\omega}}{\sqrt{1+\frac{(1-\alpha)^2\tau^2p_0^2}{\alpha^2\omega^2}} - \frac{(1-\alpha)\tau p_0}{\alpha\omega} - 1}\right] \\
&= p_0 q \lambda^b
\end{aligned}
$$

Here,

$$\lambda^b = \left[\frac{\sqrt{1+\frac{(1-\alpha)^2\tau^2 p_0^2}{\alpha^2\omega^2}}-\frac{(1-\alpha)\tau p_0}{\alpha\omega}}{\sqrt{1+\frac{(1-\alpha)^2\tau^2 p_0^2}{\alpha^2\omega^2}}-\frac{(1-\alpha)\tau p_0}{\alpha\omega}-1}\right]$$

In addition, the liquidity trader also pays the liquidity fee of τ per Eth. Thus, the total payment is $p_0 q_0(1+\tau)\lambda$. Hence, the per unit cost of trading q units is $(1+\tau)\lambda - 1$.

$$\begin{aligned}
e_\ell^s &= E_0 - \frac{E_0 T_0}{T_0+q} \\
&= p_0 T_0 - \frac{p_0 T_0^2}{T_0+q} \\
&= p_0\left(q\left[\sqrt{1+\frac{(1-\alpha)^2\tau^2 p_0^2}{\alpha^2\omega^2}}-\frac{(1-\alpha)\tau p_0}{\alpha\omega}\right]\right) - \frac{p_0\left(q\left[\sqrt{1+\frac{(1-\alpha)^2\tau^2 p_0^2}{\alpha^2\omega^2}}-\frac{(1-\alpha)\tau p_0}{\alpha\omega}\right]\right)^2}{\left(q\left[\sqrt{1+\frac{(1-\alpha)^2\tau^2 p_0^2}{\alpha^2\omega^2}}-\frac{(1-\alpha)\tau p_0}{\alpha\omega}\right]\right)+q} \\
&= p_0 q\left[\sqrt{1+\frac{(1-\alpha)^2\tau^2 p_0^2}{\alpha^2\omega^2}}-\frac{(1-\alpha)\tau p_0}{\alpha\omega} - \frac{\left(\sqrt{1+\frac{(1-\alpha)^2\tau^2 p_0^2}{\alpha^2\omega^2}}-\frac{(1-\alpha)\tau p_0}{\alpha\omega}\right)^2}{\sqrt{1+\frac{(1-\alpha)^2\tau^2 p_0^2}{\alpha^2\omega^2}}-\frac{(1-\alpha)\tau p_0}{\alpha\omega}+1}\right] \\
&= p_0 q\left[\frac{\sqrt{1+\frac{(1-\alpha)^2\tau^2 p_0^2}{\alpha^2\omega^2}}-\frac{(1-\alpha)\tau p_0}{\alpha\omega}}{\sqrt{1+\frac{(1-\alpha)^2\tau^2 p_0^2}{\alpha^2\omega^2}}-\frac{(1-\alpha)\tau p_0}{\alpha\omega}+1}\right] \\
&= p_0 q \lambda^s
\end{aligned}$$

where

$$\lambda^s = \left[\frac{\sqrt{1+\frac{(1-\alpha)^2\tau^2 p_0^2}{\alpha^2\omega^2}}-\frac{(1-\alpha)\tau p_0}{\alpha\omega}}{\sqrt{1+\frac{(1-\alpha)^2\tau^2 p_0^2}{\alpha^2\omega^2}}-\frac{(1-\alpha)\tau p_0}{\alpha\omega}+1}\right]$$

Thus, the expected Eth cost of trading tokens is

$$\begin{aligned}
&= (1+\tau)\left(\frac{1}{2}p_0 q\lambda^s + \frac{1}{2}p_0 q\lambda^b\right) \\
&= p_0 q(1+\tau)\frac{\left[\sqrt{1+\frac{(1-\alpha)^2\tau^2 p_0^2}{\alpha^2\omega^2}}-\frac{(1-\alpha)\tau p_0}{\alpha\omega}\right]^2}{\left[\sqrt{1+\frac{(1-\alpha)^2\tau^2 p_0^2}{\alpha^2\omega^2}}-\frac{(1-\alpha)\tau p_0}{\alpha\omega}\right]^2-1}
\end{aligned}$$

The trading costs are in excess of the the fundamental value of the asset, p_0.

$\blacksquare$