Decentralized Exchanges

Alfred Lehar*
Haskayne School of Business
University of Calgary

Christine A. Parlour†
Haas School of Business
UC Berkeley

May 10, 2021

Preliminary and VERY Incomplete
Please do not Circulate or Post

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.

†email:parlour@berkeley.edu.

Decentralized Exchanges

Preliminary and incomplete

Abstract

Uniswap is one of the largest decentralized exchanges with a liquidity balance of over 3 billion USD and daily trading volume of over 700 million USD. It is designed as a system of smart contracts on the Ethereum blockchain, and is a new model of liquidity provision, so called automated market making. We collect and analyze data on all 19 million Uniswap interactions from 2018 to the current time. For this new market, we analyze returns to liquidity provision and returns. We document return chasing in liquidity provision and cross-sectional heterogeneity in returns to liquidity.

1 Introduction

Uniswap is a decentralized exchange that launched in November 2018. Currently, it has a liquidity balance of over 3 billion USD in various cryptocurrencies and facilitates transactions worth over 700 million USD per day. One of the striking features of this successful exchange is that instead of a centralized limit order book, it uses a different 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.

Briefly, an automated market maker (AMM) is a mechanical protocol for supplying and demanding liquidity. In a limit order book, potentially strategic traders submit quantities at price at which they are willing to trade. In an AMM, agents supply liquidity to a pool, which comprises two distinct sets of assets. An agent supplying liquidity adds both assets in proportion to the two existing pools. An agent demanding liquidity supplies one of the assets, and removes the other one according to a predetermined downward sloping, convex relationship. The convexity implies that larger orders have a larger price impact. In addition, all liquidity demanders pay a proportional fee. At any point of time, the terms of trade are defined by the ratio of the two assets in the pool. If these terms of trade differ from "market price," then it is assumed that arbitraguers will enter the market and adjust the size of the pool until it is in equilibrium.

The design of automated market making arose from computer science. The most striking feature of the protocol is that it unbundles liquidity supply from price formation. There are various other differences between liquidity supply in a limit order market and that by an AMM. In both, passive liquidity is subject to adverse selection, that is a trade is more likely if an asset is mispriced. Thus, for any pair of tokens that she supplies to the market, she is more likely to retain the one that has become worth less. Different from a limit order book is that the returns to liquidity provision are shared pro-rata. There is therefore no time or price priority. In a traditional limit order book, liquidity is replenished by competing liquidity suppliers. Under automated market making, liquidity is replenished by arbitrageurs and the benefits accrue to the original liquidity providers. Finally, as liquidity suppliers all share in the rewards to liquidity, these markets allow us to observe latent liquidity.

We have assembled a very detailed data set of 43,349,198 interactions with the Uniswap smart contract. These allow us to identify all flows into and out of 36,958 liquidity pools as well as all the trades of tokens. The preponderance of liquidity provision is for wrapped Ether and US dollar stable coin pools. We can trace how liquidity is both supplied and demanded at this market. Similar to most financial networks Uniswap displays a core-periphery structure.

Our detailed data set also allows us to calculate the returns to liquidity provision for a range of crypto-assets. We find that returns are on average positive over our sample and are positively related to the risk that liquidity holders face. Returns vary widely by pool and over time. We conclude that there are time varying returns to liquidity provision. These returns are higher when there is more liquidity provided consistency with a liquidity externality. Our data is consistent with an equilibrium pool size – for large pools an increase in liquidity flows leads to future liquidity withdrawals, while for smaller pools growth in pool size lead to more liquidity additions. Similar to mutual fund flows we document that there appears to be "yield chasing" in liquidity pools. High past returns lead to future inflows while low past returns lead to future outflows. Liquidity use is also persistent.

We compare prices and volume for tokens listed on both Uniswap and Binance and find that prices are remarkably close. Pricing error is smaller when trading volume is somewhat 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. We find that in general price impact on Uniswap is low and stady, while price impact on Binance is generally higher and varies a lot over time.

A few papers have analyzed the theoretical properties of constant function market makers. In a general framework, Guillermo Angeris & Tarun 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, Guillermo Angeris, Hsien-Tang Kao, Rei Chiang & Charlie Noyes (2019) presents a more specific analysis of Uniswap. Similarly, Jun Aoyagi (2020) characterizes the effect of information asymmetry on these types of markets and shows that the equilibrium liquidity supply size is stable.

2 Automated Market Making

A general analysis of constant function market makers appears in Angeris & Chitra (2020); while Angeris et al. (2019) examine the Uniswap protocol specifically. To provide a context for our empirical results, we first describe how agents provide liquidity to the system, and are remunerated, and then describe how trades occur.

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. An agent who makes such a deposit receives a proportional amount of a liquidity token. This third token is specific to the pool and represent 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 in value. Liquidity providers can redeem their liquidity tokens at any time and get their share of the liquidity pool paid out in equal value of ETH and tokens.. Providing liquidity is potentially profitable because each trade faces a tax of 30bps which is redeposited into the pool. Of course, in keeping with any form of passive liquidity there is the possibility of being adversely selected.

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. At any time, the implied price quoted by the pool is given by the ratio of the amounts of the two coins. In this case, someone who is interested in buying 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 $E_1 - E_0$, and the amount of Eth in the pool drops.

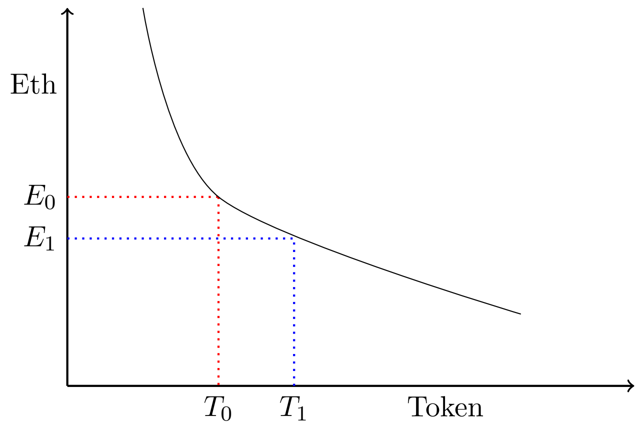

Figure 1. A bonding curve

In this way, the terms of trade are mechanically determined by the bonding curve algorithm. So, if T is the amount of tokens and E the amount of ETH in the contract's liquidity pool, then the terms of trade are set such that for any post trade quantities before any fee revenue T', E'

$$k := T' \cdot E' = T \cdot E. \tag{1}$$

In other words, the product of the Token and ETH quantities is always on the bonding curve. For each pool, k, depends on the amount of liquidity that has been deposited in the pool.

The previous clarifies the terms of trade absent the liquidity fee. Of course, this 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 κ, which benefits liquidity holders.[1] Thus the effective amount of ETH that gets traded is $(1-\kappa)e$. This leads to a post trade, but before fee revenue liquidity pool balance of $E' = E + (1-\kappa)e$. Following the logic of the bonding curve (1), the post trade token balance must be

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

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

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

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

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

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

To derive a measure of price impact later we define the zero quantity terms of trade of a pool, p^0, as the terms a pool would charge for an infinitesimal trade without fees ($e \to 0, \kappa \to 0$) which is just the ratio of the ETH reserve over the token reserve,

$$p^0 = E/T \tag{5}$$

Notice that the liquidity fee generates what is essentially a tick size that is distinct from the volume-induced that the trader pays when he moves long the bonding curve, then

$$\lim_{e \to 0} \frac{p^{tot}}{p^0} = \frac{ET}{ET(1-\kappa)} = \frac{1}{1-\kappa} \tag{6}$$

That is when buying tokens, traders have to pay a fixed spread of $\frac{1}{1-\kappa}p^0$. Similarly for token sales traders have to pay a fixed spread of $(1-\kappa)p^0$.

The price that a trader gets is determined by the bonding curve. Further, the price impact of a marginal increase in the order is $\partial p/\partial e = 1/T$. As the liquidity pool grows both token and Ether have to be added in equal value, thus growing both T and E. The price impact is therefore decreasing in the size of the liquidity pool. 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 large pool. Because Uniswap has a unique mapping of trading quantity to price the graph shows the exact amount that is traded at a certain 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.

2.1 Numerical Examples

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 $\kappa = 0.003$.

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* $p = \frac{e}{T} + \frac{E}{(1-\kappa)T} = \frac{10}{10} + \frac{100}{(1-0.0003)10} = 11.0301$

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*

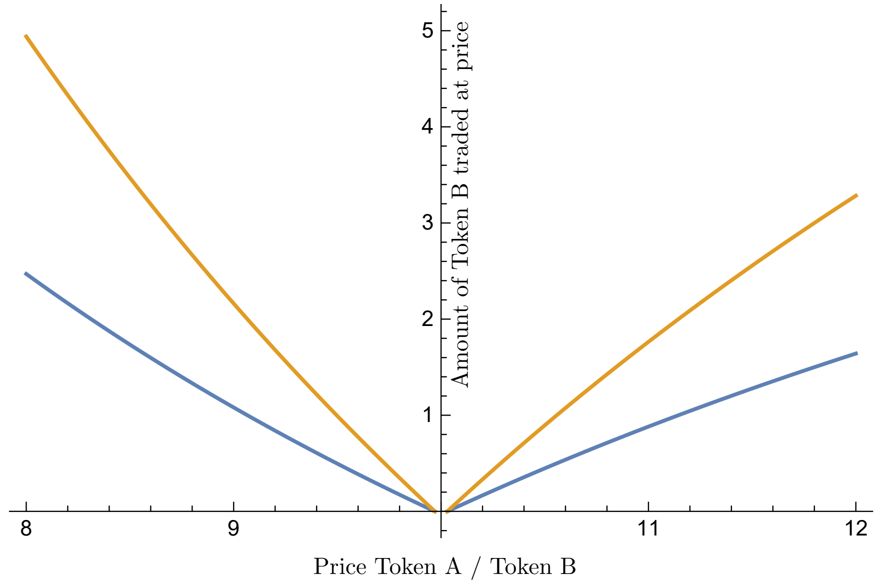

Figure 2. Uniswap orderbook depth The graph shows how many Token B could be bought or sold at a given price for a large (orange) and small (blue) liquidity pool, respectively. The parameters are: $\kappa = 0.003$ and $T = 20, E = 200$ for the large pool, and $T = 10, E = 100$ for the small pool.

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* ($\kappa e = 0.003 \cdot 10 = 0.03$) *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 $e = \frac{(1-\kappa)tE'}{T'+t(1-\kappa)} = \frac{0.997 \cdot t \cdot 110}{9.09339+0.997t}$ 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. $\kappa = 0$, 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.*

3 Framework

Consider a market with 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 potentially informed trader monitors the market and trades when profitable, otherwise a passive trader, who trades a fixed quantity q, arrives. The passive trader is equally likely to buy or sell. The passive trader is willing to trade his q shares at any price $p \in [p_0 - \sigma, p_0 + \sigma]$, else they will not trade.

There are two rational liquidity suppliers who potentially enter the market before the passive trader and post a price that optimally trades off the surplus they can extract from him against the possibility of being "picked off" by an informed trader. We focus on the case of two liquidity suppliers as it is the minimum required for competition. To simplify the exposition, we only detail the case where the informed trader buys. The case where the informed trader sells is symmetric.

We envisage rational liquidity suppliers searching over profitable asset markets to supply liquidity. Specifically, with probability γ a liquidity supplier arrives at the asset market and chooses an order to post. To make search more efficient, rational liquidity suppliers can invest $I > 0$ in a market monitoring technology. After this investment, they are more likely to identify liquidity provision opportunities, specifically, they identify the market with probability $\hat{\gamma} > \gamma$, where we denote the incremental improvement by Δ_γ.

3.1 Limit order market

The sequence of events in the limit order market is as follows: First, nature determines the number of limit order submitters, then the limit order submitters post their orders. Second, nature determines the new asset value. If there is no information event, the liquidity trader arrives and trades against the best quote or randomizes if indifferent. If there was an information event, then, the informed trader trades against the best quote. 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 is profitable, i.e., $2q$.

Given that liquidity traders are searching for profitable trading opportunities, they may be alone or competing in a market. 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 sole liquidity supplier in the market, will post a price of sell price of* $p_0 + \sigma$ *and obtain a profit of* $(1 - \alpha)\sigma$.

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. However, the possibility of adverse selection means that they will not charge the lowest possible price, but rather randomize.

Lemma 2 *Consider two competing liquidity suppliers in the market. In symmetric, mixed strategy equilibrium each will choose a distribution over prices $F(\cdot)$ over $[p_{\min}, p_{\max}]$, where*

$$F(p) = \frac{(p-p_0)(1-\frac{\alpha}{2}) - \sigma\frac{\alpha}{2}}{(p-p_0)(1-\alpha)},$$

and

$$p_{\min} = p_0 + \sigma\left(\frac{\frac{\alpha}{2}}{1-\frac{\alpha}{2}}\right) \qquad p_{\max} = p_0 + \sigma.$$

Consistent with standard intuition, there is a positive spread in this market, which is larger, the larger the possibility of adverse selection, which increases in σ. Going forward, we define

$$\begin{aligned} p_{min} - p_0 &= \sigma\left(\frac{\frac{\alpha}{2}}{1-\frac{\alpha}{2}}\right) \\ &= s_\sigma, \end{aligned}$$

which is the half spread that is increasing in the size and probability of an innovation.

Proposition 3 *There are threshold investment costs $\underline{I} < \bar{I}$, so that if $I \leq \underline{I}$ both liquidity traders will pay to monitor the market, if $\underline{I} < I < \bar{I}$ only one will pay to monitor the market, and if $I > \bar{I}$, none will pay to monitor the market.*

To determine the cost to the liquidity trader, observe that he will execute against the best quote when the enters the market, or $c(q) = \min[p_i, p_j]$. Let $F_c(x)$ denote the cdf of the best price, then

Lemma 4 *The expected transaction price for the liquidity trader is*

$$E(c) = \left(\frac{\frac{\alpha}{2}(1-\delta)}{(1-\alpha)}\right)^2 \left(-s_\sigma + \frac{\sigma^2}{s_\sigma} + 2\sigma \ln\frac{s_\sigma}{\sigma}\right)$$

Similarly, we can calculate the expected profits to the informed trader from this competitive supplied liquidity. We obtain:

$$\begin{aligned} \pi^I &= \frac{\alpha}{2}\left\{(1-\delta)\delta\, q + (1-\delta_j)\delta_i q + (1-\delta_i)(1-\delta_j)2q\right\} Ep^* \\ &= \alpha\, q(1-\delta)Ep^* \end{aligned}$$

Where, the expected price is given by

$$\begin{aligned} Ep^* &= \int_{p_{min}}^{p_0+\sigma} [1 - \frac{(p-p_0)(1-\frac{\alpha}{2}(1+\delta_i)) - \sigma\frac{\alpha}{2}(1-\delta_i)}{(p-p_0)(1-\alpha)}]dp \\ &= \frac{\alpha(1-\delta)(-s_\sigma - \sigma - \sigma\ln(-s_\sigma) + \sigma\ln(-\sigma)}{2(1-\alpha)} \end{aligned}$$

Notice that the arbitrageur would be willing to trade an arbitrary amount up to the new price, $p_0 + \sigma$, but is unable to do so due. Liquidity providers restrict the amount that they post to minimize the chance of trading with the informed trader.

3.2 Bonding curve market

In the bonding curve market, liquidity suppliers commit quantities of both Eth and Tokens. Then, with probability $\frac{\alpha}{2}$ there is an asset innovation and the informed trader removes tokens from the pool and deposits Eth. (Recall, we are characterizing the case in which there is a positive asset innovation.) With probability $1-\alpha$, a liquidity trader arrives and buys q Tokens. As a result of this trade, the quoted Eth price of tokens is too high, and so the arbitrageur enters the market and sells to the liquidity providers until the Eth price of tokens reverts to the equilibrium amount.

Assume that investors have each committed E_0 Eth and the equivalent total amount of tokens T_0. Two identities will be useful going forward, first, the ratio of Eth to tokens, in equilibrium, is the Eth price of tokens, or

$$\frac{E_0}{T_0} = p_0. \tag{7}$$

Second, any transactions must occur along the curve, defined by

$$E_0 T_0 = k, \tag{8}$$

where k is the constant of the bonding curve.

To facilitate comparison with the limit order market, observe that Eth is the numeraire, and so the amount committed to trade is the number of tokens (T_0). We seek to determine the equilibrium provision of this quantity, which will also determine the costs to the liquidity trader and profits of the informed trader. To proceed, we establish how the relative balance of Eth and the token change in response to the various types of asset innovations and trader arrivals.

First suppose that a liquidity trader arrives. Recall, that this trader demands q tokens, and will add in Δ_E^ℓ to the Eth pool. The Eth cost for doing this trade is determined by the bonding curve, so

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

Δ_E^ℓ can be further decomposed into a price per token and a liquidity fee that accrues to the liquidity suppliers. Specifically, $\Delta_E^\ell = (1+\tau)p^{trans}$, so the transaction price for q units is simply $p^{trans} = \frac{\Delta_E^\ell}{(1+\tau)}$. The payoff to the liquidity providers is therefore $\frac{\tau}{(1+\tau)}\Delta_E^\ell$.

Given that $\Delta_E^\ell > 0$, the ratio $\frac{E_0+\Delta_E^\ell}{T_0-q} > \frac{E_0}{T_0}$. The pool is now quoting an Eth price of the token that is too high. This is an arbitrage opportunity, and therefore the arbitrageur will add tokens and remove Eth so that the pool returns to equilibrium. The value of the pool before and after the liquidity trader is the same. However, both the liquidity trader and arbitrageur have paid a transaction fee of $2\frac{\tau}{1+\tau}\Delta_E^\ell$. This is the payoff to liquidity provision.[2]

Now suppose that there was a positive innovation event so that an informed trader arrives. Since the pricing is deterministic she will trade an amount that maximizes her profit. Specifically, the informed trader will buy tokens until the new pool price is equal to $p_0+\sigma$, net of the transaction costs. She implements the trade by sending Δ_E^I Eth to the pool and removing Δ_T^I tokens in return so that,

$$\frac{(E_0 + \Delta_E^I(1+\tau))}{(T_0 - \Delta_T^I)} = p_0 + \sigma. \tag{9}$$

Pool balances post trade are again on the bonding curve, so,

$$(E_0 + \Delta_E^I(1+\tau))(T_0 - \Delta_T^I) = E_0 T_0.$$

We can solve for the amount of the traded tokens, and the amount paid so

$$\Delta_T^I = T_0 - \sqrt{\frac{k}{p_0+\sigma}} \tag{10}$$

$$\Delta_E^I = \frac{\sqrt{k(p_0+\sigma)} - E_0}{1+\tau}. \tag{11}$$

Before the asset innovation, the liquidity supply was worth $E_0 + p_0 T_0$. After the innovation, the liquidity supply is worth

$$\begin{aligned} & E_0 + \Delta_E^I(1+\tau) + (p_0+\sigma)(T_0 - \Delta_T^I) \\ = \; & \sqrt{k(p_0+\sigma)} + (p_0+\sigma)\sqrt{\frac{k}{p_0+\sigma}} \\ = \; & 2\sqrt{k(p_0+\sigma)} \end{aligned}$$

Therefore, the change in value for liquidity suppliers after a change in the value of the asset is:

$$E_0 + p_0 T_0 - \left(2\sqrt{k(p_0+\sigma)}\right) \tag{12}$$

$$= 2E_0 - \left(2\sqrt{k(p_0+\sigma)}\right) \tag{13}$$

[2] Technically, the arbitrageur faces a different size pool than the liquidity trader as the liquidity fee has been paid into the pool. We do not consider this incremental effect.

Thus, the overall payoff to liquidity provision (for the entire pool) is

$$
\begin{aligned}
& (1-\alpha)2\Delta_E^{\ell}\ \frac{\tau}{1+\tau}+\frac{\alpha}{2}\left(2\sqrt{k(p_0+\sigma)}-2E_0\right) && (14)\\
= \quad & (1-\alpha)2\left(\frac{k}{T_0-q}-E_0\right)\ \frac{\tau}{1+\tau}+\frac{\alpha}{2}\left(2\sqrt{k(p_0+\sigma)}-2E_0\right) && (15)
\end{aligned}
$$

The equilibrium size of the liquidity pool is implicitly defined by

$$
(1-\alpha)2\left(\frac{k}{T_0-q}\right)\ \frac{\tau}{1+\tau}+\alpha\left(\sqrt{k(p_0+\sigma)}\right) \quad = \quad T_0p_0\left((1-\alpha)\ \frac{\tau}{1+\tau}+\alpha\right)
$$

which yields liquidity supply of

$$
\begin{aligned}
T_0 \quad &= \quad q+\frac{\sqrt{2\alpha^2k(p_0+\sigma)+4p_0\left((1-\alpha)\ \frac{\tau}{1+\tau}+\alpha\right)(1-\alpha)2\frac{\tau}{1+\tau}-2qp_0\alpha\left(\sqrt{k(p_0+\sigma)}\right)\left((1-\alpha)\ \frac{\tau}{1+\tau}+\alpha\right)}}{2p_0\left((1-\alpha)\ \frac{\tau}{1+\tau}+\alpha\right)}\\
&= \quad q+\frac{\sqrt{2\alpha^2k(p_0+\sigma)+\left((1-\alpha)\ \frac{\tau}{1+\tau}+\alpha\right)\left[4p_0(1-\alpha)2\frac{\tau}{1+\tau}-2qp_0\alpha\sqrt{k(p_0+\sigma)}\right]}}{2p_0\left((1-\alpha)\ \frac{\tau}{1+\tau}+\alpha\right)}
\end{aligned}
$$

There are a few things to notice about this. First, the liquidity supply is not constrained by the strategic considerations of the the liquidity suppliers. Specifically, because of adverse selection liquidity suppliers collectively post a maximum of 2q. By contrast as the predetermined price impact is shared across all liquidity suppliers, more liquidity may be posted on the DEX.

The profits to the informed trader are

$$
\begin{aligned}
\pi^I \quad &= \quad \Delta_T^I(p_0+\sigma)-\Delta_E^I\\
&= \quad \left(T_0-\sqrt{\frac{k}{p_0+\sigma}}\right)(p_0+\sigma)-\left(\frac{\sqrt{k(p_0+\sigma)}-E_0}{1+\tau}\right)\\
&= \quad T_0(p_0+\sigma)+\frac{E_0}{1+\tau}-\sqrt{k(p_0+\sigma)}\ \frac{2+\tau}{1+\tau}
\end{aligned}
$$

Notice, this differs in the payoff to the informed trader in the limit order market.

In the limit order market, liquidity provision is rivalrous, and the liquidity suppliers trade off the benefits of trading with the passive trader against the potential adverse selection costs. This implies that there will be a positive spread even with continuous prices and there will be excess volatility as orders. By contrast the automated market maker shares all benefits with liquidity provision among all liquidity suppliers. Further the amount of adverse selection is restricted by the bonding curve.

3.3 UniSwap

Decentralized exchanges (DEX) are smart contracts mostly deployed on the Ethereum blockchain. By interacting with the smart contract users can exchange digital tokes and currencies in an automated, trustless way. Users initiate an exchange by posting an Ethereum transaction that sends token or cryptocurrency are to be sold to the contract and calls a function of the smart contract to perform the exchange. The smart contract then sends other tokens or cryptocurrency back. Since transactions on Ethereum are atomic, meaning that 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 at any time. The source code for many DEX is public and users can verify that the code is not fraudulent and perfectly predict the smart contract's behavior.

Uniswap which was launched in November 2018 at Devcon 4 and the first pool allowed swaps between Eth and Mkr. Uniswap is open source, functions as a public good, and has no owner or operator. Similar to other decentralized exchanges ("DEX") it employs smart contracts.

In its first release, Uniswap V1, it allows the exchange of ERC20 tokens against Ether (ETH) by interacting with a smart contract, which we will refer to as a *liquidity pool*.[3] Each pool allows the trading of exactly one token against ETH. In case that there no pool exists for a specific token it can be created for free by anyone calling 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.

Uniswap V2 was launched on May 18, 2020 and allows the direct trading of ERC 20 token pairs. It provides several benefits over V1 such as a broader set of permissible tokens (such as Thether USDT for example), enhanced functionality, and better oracle functionality.[4] Because Uniswap has no owner V1 pools cannot be deleted from the blockchain and exist in parallel to V2 pools. As we document below most V2 pools trade tokens against wrapped Ether (WETH), which is a ERC 20 representation of Ether (ETH), the native currency on the Ethereum blockchain. To simplify the exposition of the paper we illustrate the mechanics of liquidity provision and trading in a pool which is trading a Token against ETH.

4 Data

Using a factory contract on Ethereum, anyone can create a UniSwap liquidity pool. First, we obtained a list of all UniSwap V1 and V2 liquidity pools from factory contract transactions. In our sample we have a total of 36,958 individual liquidity pools, consisting of 3,937 V1 pools

[3] Most digital tokens that are traded on Ethereum follow the ERC20 specification.

[4] Oracles provide information to other smart contracts that they need as input for their program. Many smart contracts use Uniswap as price feed to obtain current token prices, pretty much in the same way that many traders in mainstream financial markets use Bloomberg. There have been instances of price manipulation where traders placed huge orders on Uniswap V1 exchanges to push the price in a certain way and then take advantage of other smart contracts relying on this misleading information, for example to inflate the value of collateral against which they borrow. Once prices revert to normal they default on the undercollateralized loan. Uniswap V2 improves its oracle functionality by providing a moving average of past prices.

and 33,021 V2 pools. We then matched transactions into and out of these liquidity pools with block-by-block transactions on the Ethereum block chain. Our data thus comprises 43,349,198 transactions on Uniswap from its inception on November 2, 2018 until April 29, 2021.

We note that, in contrast to traditional exchanges on which there are listing requirements, UniSwap liquidity pools are not certified. Indeed, some of the token pools are misleading. For example five different tokens in our sample have the ticker symbol USDC. A naive user, who does not verify the smart contract address could be tricked into buying a worthless coin with the same ticker.[5] While fake tokens exits, trading activity in these pools is limited and will not affect our results. We provide detailed information on fake tokens in Appendix B.

Pools may differ in size but also in the volume of transactions. Figure 3 shows the size of all uniswap liquidity pools in ETH. The volume is broken down in subgroups by volume. The 10 largest exchanges by total volume over the sample period are in the *high volume* group (blue), the next 190 exchanges are in the *mid volume* group (green), and the remaining exchanges are in the *low volume* group (orange). Most of the liquidity is concentrated in the high volume group. These exchanges are mostly V2 exchanges and are responsible for most of Uniswap's growth.

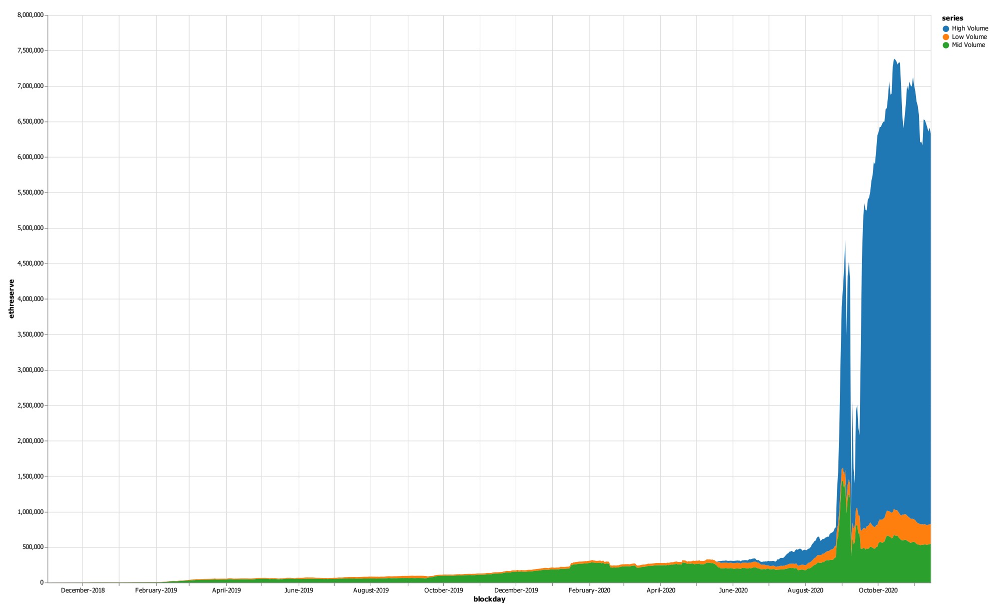

Figure 3. Relative size of liquidity pools in ETH broken down by volume.

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

[5]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 copy-cat token at the address 0x1bffb8a3fede9f83a3adc292ebf1716d40b220c1, which has a total of 10 trades and the size of the pool never exceeded 50 ETH.

Token 1		Token 2		Number Transactions	Volume (ETH)	Volume (USD)	Pool size (ETH)
		Panel A: Uniswap V2					
Wrapped Ether	WETH	Tether USD	USDT	7,262.7	84,122	61,970,290	227,782
USD Coin	USDC	Wrapped Ether	WETH	5,632.0	81,914	63,189,134	209,443
Dai Stablecoin	DAI	Wrapped Ether	WETH	3,040.0	48,091	34,716,533	169,411
Uniswap	UNI	Wrapped Ether	WETH	2,520.8	33,181	26,172,330	54,116
Wrapped BTC	WBTC	Wrapped Ether	WETH	958.2	30,096	22,236,270	298,367
yearn.finance	YFI	Wrapped Ether	WETH	924.5	21,389	9,733,820	27,646
Fei USD	FEI	Wrapped Ether	WETH	285.6	19,248	41,958,876	512,350
Tendies Token	TEND	Wrapped Ether	WETH	153.9	16,551	23,246,953	728
SushiToken	SUSHI	Wrapped Ether	WETH	951.5	16,087	7,186,316	78,652
Wrapped Ether	WETH	Ampleforth	AMPL	1,073.2	14,652	6,101,790	38,886
		Panel B: Uniswap V1					
Ether	ETH	Dai Stablecoin	DAI	550.3	2,777	521,430	9,423
Ether	ETH	HEX	HEX	224.4	1,874	399,122	22,701
Ether	ETH	Maker	MKR	144.8	1,347	267,577	10,519
Ether	ETH	USD Coin	USDC	255.6	1,297	277,424	7,086
Ether	ETH	LoopringCoin V2	LRC	21.1	1,012	376,052	800
Ether	ETH	Sai Stablecoin v1.0	SAI	202.3	985	223,721	5,047
Ether	ETH	Synthetix Network Token	SNX	129.0	723	140,262	3,536
Ether	ETH	Synth sETH	sETH	45.5	596	114,721	26,762
Ether	ETH	UniBright	UBT	111.8	289	60,788	636
Ether	ETH	Pinakion	PNK	41.9	202	61,761	1,537

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. We exclude pools with less than 5,000 total transactions. .

caused most traders and liquidity providers to converge on the new protocol. The largest pool in terms of total volume traded is Tether (USCT) - Wrapped ETH with an aggregate volume (over all days) of over 21.4 billion USD. This pool also has the highest number of total trades in our sample over 2.5 million trades. The most liquid pool is FEI-WETH with an average size of over 512 thousand ETH.

From the Ethereum blockchain we observe 1,027,802 liquidity injections into a pool, 551,796 withdrawals of liquidity from a pool, and 41,317,816 trades of tokens. The remaining 451,784 transactions are either complex transactions that combine liquidity additions or removals with swaps or flash swaps. Briefly, flash swaps were introduced in Uniswap V2 and allow a user to borrow any amount up to the total liquidity available in a pool, so long as 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. Several other protocols also offer such '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 0.3 percent fee on amounts borrowed.

While some of the pools are very active, many are not: 29,953 pools in our sample have fewer than 100 transactions. Figure 4 shows the number of trades. 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).

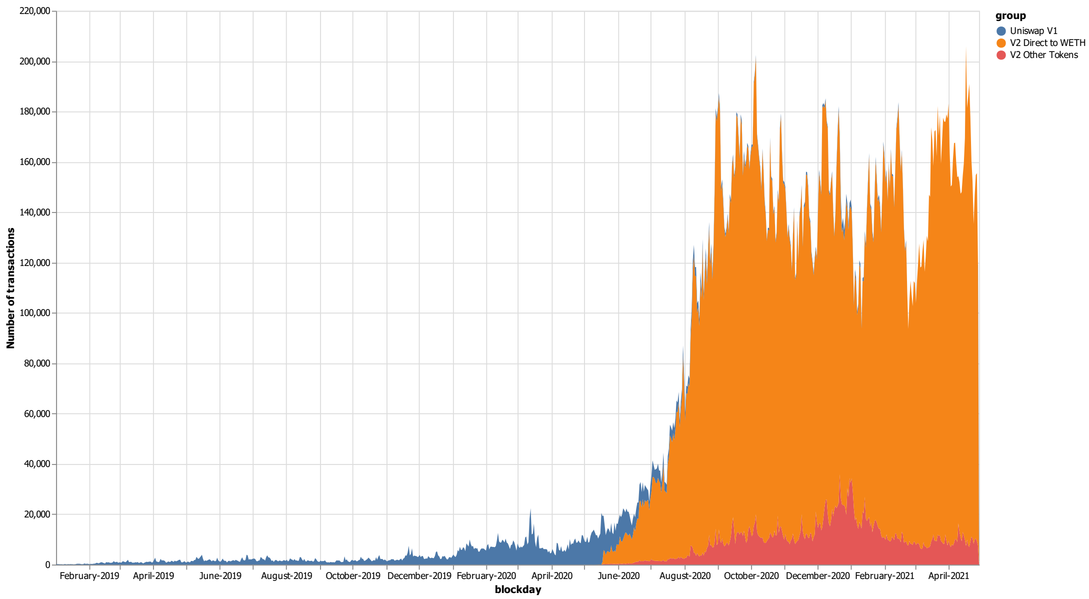

Figure 4. Number of transactions on Uniswap.

Figure 5 illustrates the trading volume per day. Trading volume varies more than the number of trades. Again V2 pools trading to WETH dominate in terms of trading volume. 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, 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. Usage of Flash Swaps varies a lot in our sample. Out of the 359 days where V2 was deployed, flash swaps occurred on 330 days. The median flash swap volume per day was 130,897 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 an attack on the platform Warp finance. Many large trades are part of an explout 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.[6] The median trade size in our sample is 838.2 USD. 30.5% of trades are below 0.5 ETH and 13.9% are below 100 USD. Computing volume is not straightforward

[6] 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.

in Bonding curve markets as attackers often deliberately push markets out of equilibrium. We provide details on our methodology in Appendix C.

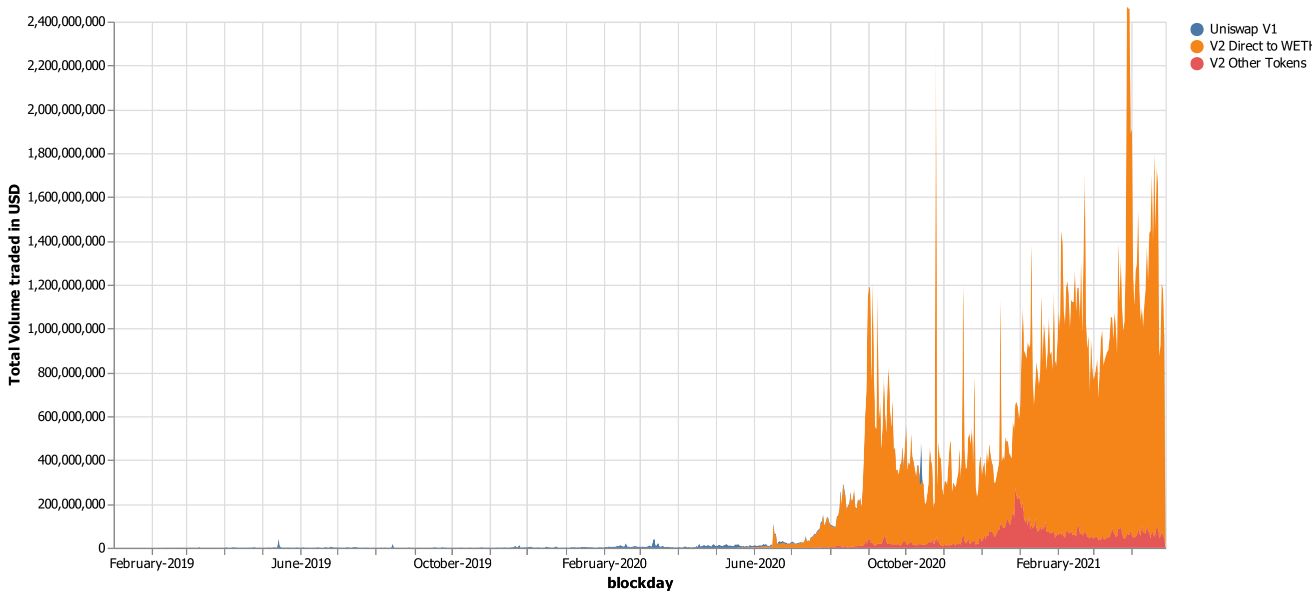

Figure 5. Trading volume on Uniswap The graph shows the daily volume of transactions including flash loans in USD capped at 2.4 billion USD.

Users interact with the family of Uniswap smart contracts by posting transactions on the Ethereum blockchain. Apart for limits on transaction size given by the Ethereum network there is no theoretical limit on how many interactions with Uniswap liquidity pools can be done in one transaction. In our sample 81.06% of Ethereum transactions only have one interaction with one liquidity pool, another 15.92% have two interactions. 1615 transactions or 0.0061% 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.[7]

Figure 6 presents the network of pools between all tokens that are part of the 50 largest pools by volume. The thickness of the line corresponds to the trading volume between the tokens and the color of the token-markers is proportional to the log of the depth of the liquidity pools for that token with red marking the most liquid tokens. We can see that Wrapped Ether (WETH) takes a central position in the Uniswap network. For our whole sample of 28,654 tokens we find that 27,796 tokens, or 97%, trade directly against WETH. The second highest number of tokens, 1,479, trade against USDT. The highest volume and the most connections are between WETH and USD stable coins such as USDT, USDC, and DAI. 2,799 tokens or 9.76% of tokens

[7]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.

are trading directly against these three stablecoins. The Uniswap network has a core-periphery structure similar to many other financial networks. 27,220 tokens or 95% of tokens trade only against one other token.

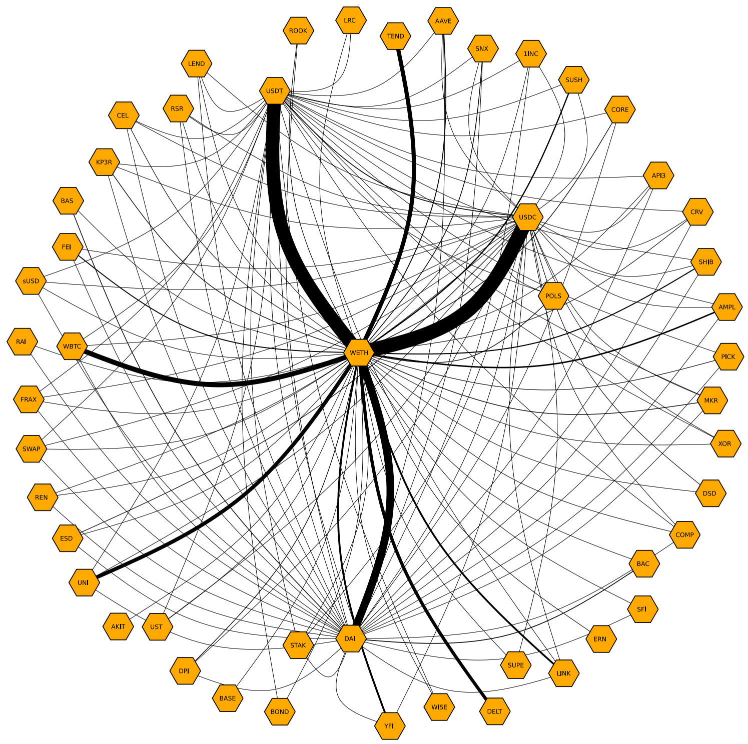

Figure 6. Network graph of pools between all tokens that are part of the 50 largest pools by volume.

5 Liquidity provision

Because the benefits to providing liquidity are shared pro rata, and the size of the liquidity pool may increase or decease as we compute daily returns to liquidity providers similarly to returns for mutual funds where in- and outflows of liquidity can occur. We assume that pools are arbitrage free, i.e., the true market price of a token equals the mid-price of the pool.[8] Under this condition the value of the T tokens in the liquidity pool is $Tp_0 = TE/T = E$, which is the same as the value of the ETH balance. The total ETH value of the pool is then $2E$.

Let λ_t be the the amount of liquidity added or removed in period t. Then the return for liquidity providers in period t is

$$l_t = \frac{2E_t - \lambda_t - 2E_{t-1}}{2E_{t-1}}. \tag{16}$$

[8] If the true market price would not equal to the mid-pool price and arbitrageur could make an instant profit and would push the mid-pool price of the token to the true market price.

Figure 7 exhibits profits accruing to liquidity providers. The blue line shows aggregate fee revenue on Uniswap over time. As volume grows, especially after the introduction of V2, so does fee revenue. The orange line shows the cumulative return for all liquidity providers across all Uniswap pools. The graph at point T is computed as $(\prod_{t=1}^{T}(1 + l_t)) - 1$.

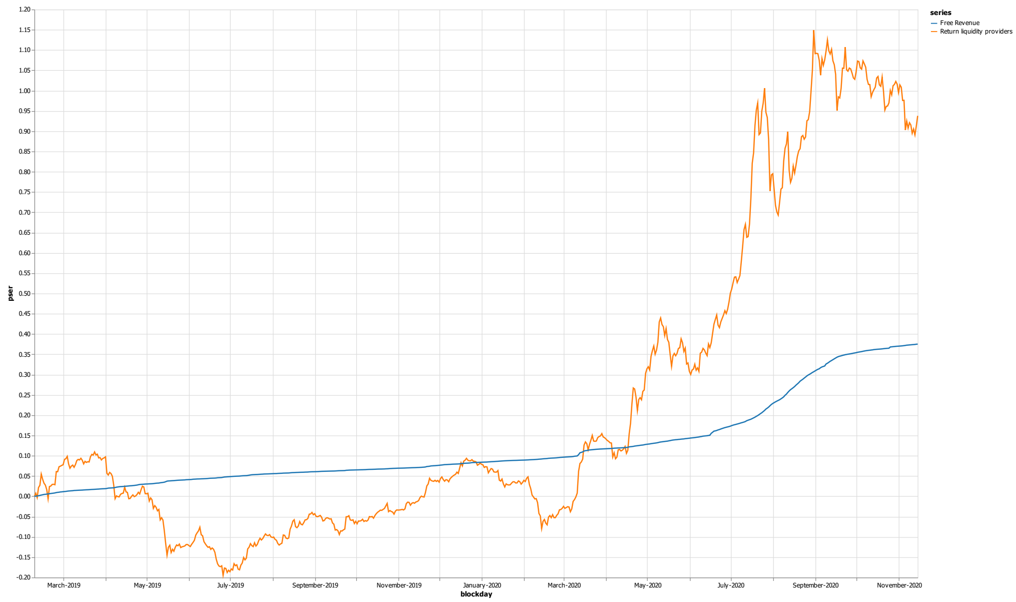

Figure 7. Aggregate Profit for liquidity providers in Uniswap across all Pools.

Figure 8 shows cumulative returns for liquidity providers broken down by average pool size. Average pool size is very unevenly distributed. The largest 4 pools have on average the same balance as all the other pools together. In total, 17, 018 pools have an average balance of less than 100 ETH. To create somewhat even buckets we define tier 1 as the largest 5 pools, tier 2 as the next 5, tier 3 as the next 10, tier 4 as the next 30, tier 5 as the next 100, and tier 6 comprises all remaining pools. Interestingly, profitability is highest for liquidity providers of the second tier, pools that are large but not at the very top. This stylized fact is consistent with a setting where these pools are provide enough liquidity to attract traders but not too big so that there is price impact and fee revenue per liquidity provided remains attractive.

To understand the cross section of expected returns for liquidity providers we collect daily returns of all liquidity pools with a balance of at least 20 ETH with at least 50 observations. Figure 9 plots average return over return volatility for large (orange), medium (blue), small (green) pools. Pool returns are very noisy which is part driven by the very short sample period, yet a positive relationship between risk and average return is visible. In Table 2 we regress average returns for liquidity providers on the volatility of the returns for liquidity providers as well as average pool volume and size. We compute the average return as the average of daily returns for liquidity providers over the sample period. We drop pools with less than 50 days of data and with an average pool balance of less than 50 ETH. The average sample size of the remaining pools is 127 days, the median 89. Volatility liq providers is the standard deviation of the daily returns, average volume is the average daily volume per pool over the whole sample

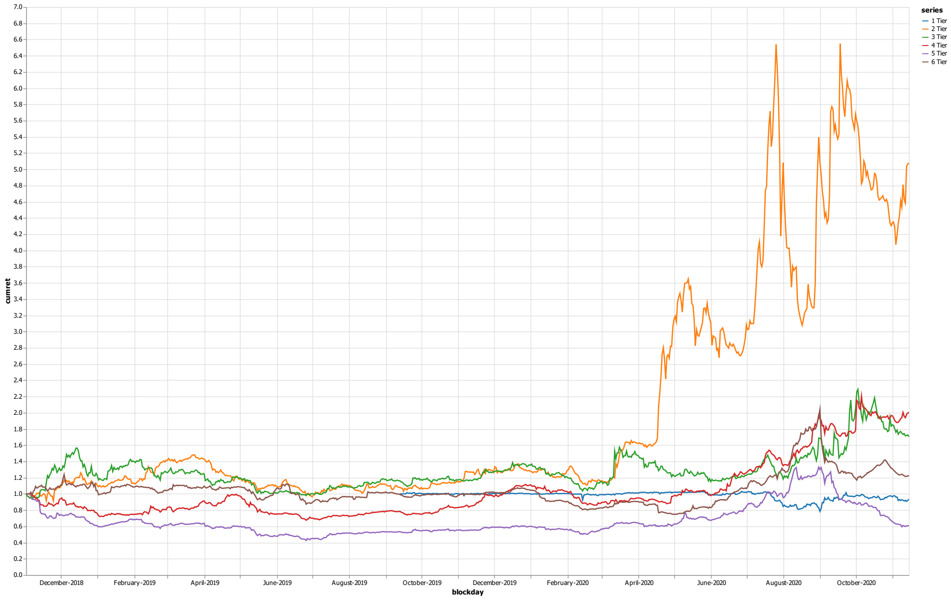

Figure 8. Cumulative returns by average pool size

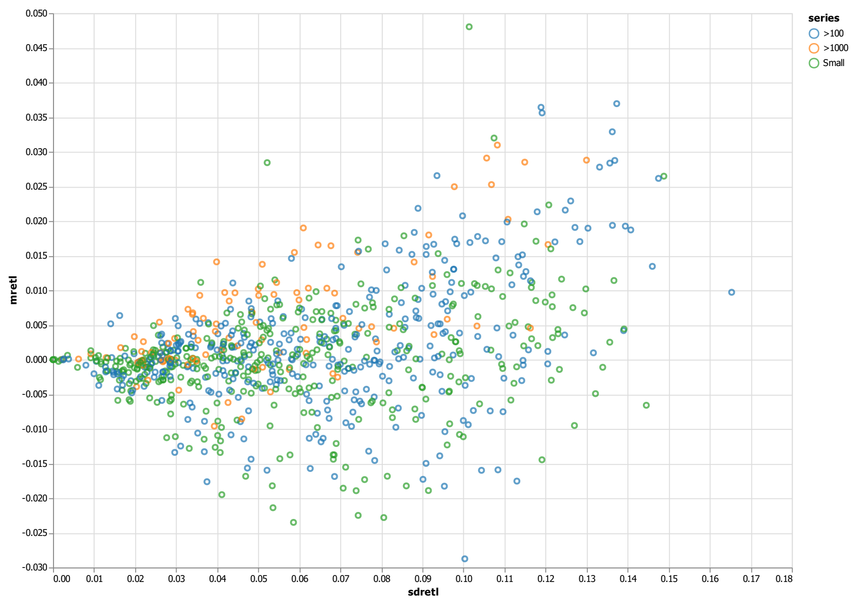

Figure 9. Risk and average returnm of liquidity provision

period, and average pool size is the average daily pool size measured in ETH over the sample period. Columns (1) and (2) include the whole sample. We can see that pools that are riskier for liquidity providers pay a higher average return. Volume is positively related to return which is intuitive as liquidity providers collect a fee that is proportionate to volume. Coulumn (1) shows the results for large pools with an average balance of more than 1,000 ETH, column (2) for median pools, and column (3) for small pools with an average balance of less than 100 ETH. Risk is a driving factor of returns for all subsamples.

	(1)	(2)	(3)	(4)	(5)
Volatility liq. providers	0.0944***	0.0925***	0.195***	0.0881***	0.0553***
	(0.00833)	(0.00832)	(0.0232)	(0.0146)	(0.0139)
Average volume		0.000000387***	9.34e-08	0.00000374***	-0.000000286
		(9.58e-08)	(7.97e-08)	(0.000000897)	(0.00000644)
Average pool size		-6.24e-08**	-1.23e-08	0.00000259	0.0000469**
		(2.64e-08)	(2.16e-08)	(0.00000207)	(0.0000188)
R^2	0.131	0.150	0.457	0.247	0.0820
Observations	853	853	102	369	382
height					

Table 2. Regression explaining the crossection of average returns for liquidity providers. One, two, and three stars indicate significance at the 10%, 5%, and 1% level, respectively.

Table 3 presents the results of a regression of the annualized percentage return for liquidity providers on several contemporaneous variables. We understand that some relationships are mechanical, driven by the design of the Uniswap contract, and therefore we use this simply to understand the Uniswap contract. Since our sample is dominated by a huge number of very small pools an analysis of the whole sample will merely be a reflection of that group. We therefore divide the sample in four groups. Daily returns of large pools with ETH balance greater than 10,000 are in column (1), with a balance between 1,000 and 10,000 in column (2), returns in column (3) are from pools with a balance between 100 and 1,000, and returns from all remaining pools are analyzed in column (4).

	(1)	(2)	(3)	(4)
Pool size	-0.0000500	0.00322*	0.00674	1.356***
	(0.0000438)	(0.00194)	(0.0124)	(0.143)
FX-return	0.504***	0.500***	0.498***	0.493***
	(0.00289)	(0.00111)	(0.000617)	(0.000520)
Std.Dev Fx Rate	-3784.1	-2579.5***	-1384.6***	-696.3***
	(4294.5)	(370.5)	(142.0)	(39.73)
Liquidity Flow	156.4***	84.29***	144.5***	111.7***
	(17.24)	(12.16)	(8.010)	(6.558)
rel. Volume	37.11***	18.94***	6.775***	-5.068***
	(11.18)	(2.482)	(1.226)	(0.956)
Reversals	-476.2**	320.9***	172.6***	268.3***
	(205.6)	(49.13)	(20.20)	(16.33)
R^2	0.929	0.943	0.936	0.892
Observations	2,408	12,511	45,369	110,292
height				

Table 3. Regression explaining daily return for liquidity providers. One, two, and three stars indicate significance at the 10%, 5%, and 1% level, respectively.

For all but the smallest pools, the poolsize is insignificant. This is consistent with the idea that those pools are close to their optimal size. In the subsample of smaller pools, the larger ones are more profitable for liquidity providers, which is consistent with the idea that pools need a certain minimum size to be operational. A key driver of returns is the FX rate. Liquidity providers hold an equal amount of each token. Following the logic of Example 1, an increase of the token relative to the numeraire causes the value of the pool to increase. Since half of the pool consists of the numeraire token we expect a coefficient of 0.5. Volatility of the exchange rate has a negative impact on liquidity holders' returns because there is a greater chance that arbitrageurs take advantage of stale prices that are quoted by the pool. This effect is not statistically significant for large pools where such losses are spread among a larger group of liquidity providers. Contemporaneous liquidity flow is associated with higher returns for liquidity providers in for all pool sizes. This is a bit counter-intuitive as mechanically liquidity inflows mean that the pool's fee revenue has to shared among more liquidity providers. Think this finding is due to some persistence in flows and profitability. When profitable pools attract more liquidity and pools remain profitable for some period of time. We define relative volume as trading volume over poolsize and find a positive relationship to profitability, which is to be expected as fees are collected in proportion to volume.

Finally we examine reversals, which we define as a trade that is immediately followed by an opposite trade of at least 50% of the size of the original trade. Reversals are often executed by arbitrageurs who want to bring back the price to the fundamental value. In the regression we use reversals as a fraction of total trades. We find that they are associated with higher returns for liquidity providers, which is intuitive because they provide fee revenue without any fundamental price change. For large pools we see a negative relationship, which is counter-intuitive. We believe that is due to measurement problems. In large pools we observe 1705 trades on average per trade compared to 7.31 trades for the other pools. With such a high frequency it is less likely to observe reversals because during the time an arbitrageur composes a reversal trade another trader might interject an order. consistent with this idea we observe only 10.7% of orders to be reversals in large pools compared to 20.2% in all other pools. In lieu of the negative association of reversals and profitability in large pools we see a larger coefficient on relative volume. One explanation is that such interjected reversals do not count as reversal but still contribute to volume.

Table 5 presents results explaining liquidity flow. Absent from the largest pools, who probably have reached their optimal size, larger pools tend to attract more liquidity providers. Liquidity flows into pools with higher token volatility and, except for the smallest pools, where more reversal trades, which are more profitable for liquidity providers, occur. Lagged return for liquidity providers is positively related to liquidity provision consistent with the commonly observed practice of yield farming. Several websites post recent returns on liquidity pools in various protocols including Uniswap and users seem to make their decisions based on these rankings.

	(1)	(2)	(3)	(4)
Pool size	-4.33e-08	0.00000476***	0.0000354***	0.000189***
	(4.93e-08)	(0.00000126)	(0.00000583)	(0.0000236)
Std.Dev Fx Rate	39.59***	1.270***	1.522***	0.130***
	(4.708)	(0.240)	(0.0680)	(0.00702)
Lagged rel. Liquidity Flow	-0.0513***	-0.00134	0.00482*	0.00654***
	(0.0151)	(0.00471)	(0.00250)	(0.000928)
Lagged rel. Volume	-0.0112***	-0.000964***	-0.000711***	-0.0000150***
	(0.00220)	(0.000162)	(0.0000650)	(0.00000103)
Lagged Reversals	3.775***	0.555***	0.100***	-0.00734***
	(0.299)	(0.0420)	(0.0132)	(0.00236)
Lagged ret. liq. providers	0.126*	0.0353**	0.0347***	0.00631***
	(0.0690)	(0.0164)	(0.00734)	(0.00213)
R^2	0.106	0.0206	0.0152	0.00337
Observations	2,335	12,136	44,521	240,065
height				

Table 4. Regression explaining liquidity flows. One, two, and three stars indicate significance at the 10%, 5%, and 1% level, respectively.

6 Latent liquidity

Decentralized exchanges allow us to observe latent liquidity. Because of the mechanical way prices are set as a function of trading and because liquidity providers cannot instantly withdraw their funds a trader can anticipate the price they will get. To measure the relationship between offered and used liquidity we compute for each liquidity pool the liquidity use by dividing the daily trading volume by the average daily pool balance. Figure 10 plots average liquidity use ratios for large, medium, and small pools as measured by overall average pool balance. Tier 1 comprises the largest 20 pools, tier 2 the next 80, and tier 3 all remaining pools. Liquidity use is fairly correlated among the large pools. For the largest 20 pools two principal components explain 69.5% of the variation in liquidity use.

In Table 5 we analyze liquidity use as a function of several lagged indicators. Liquidity use is increasing in pool size for smaller pools, which makes sense as price impact decreases with the size of a liquidity pool. For large pools price impact is already low and an increase in poolsize will not affect the price traders can get by much. liquidity use increases in FX volatility. Many crypto traders engage in margin trading where they, for example, post tokens as collateral on platforms such as Maker in exchange for DAI (a USD stablecoin). They then convert the DAI on Uniswap to buy more tokens and thus build a levered position in the token. Higher volatility in crypto markets increases the risk of these levered positions and many traders go through uniswap to reduce their leverage. Liquidity use increases in lagged liquidity flow for all but the smallest pools. This is intuitive as additional liquidity reduces price impact and thus makes it more attractive to trade in that pool. Finally liquidity use seems to be persistent. Lagged liquidity use is positively related to current liquidity use.

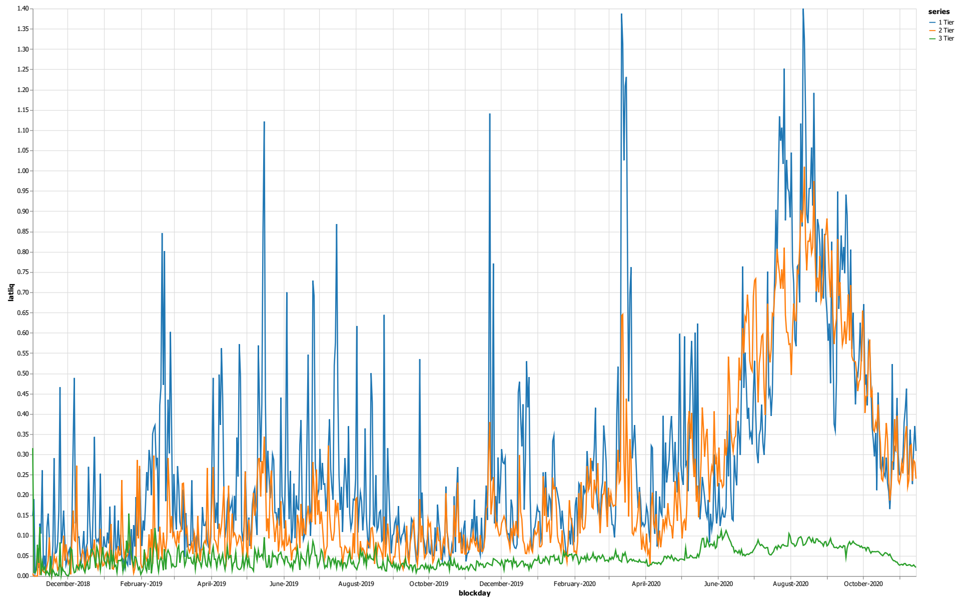

Figure 10. Liquidity use Uniswap. The graph shows the volume traded per day as a fraction of the average poolsize on that day for three tiers of liquidity pools.

	(1)	(2)	(3)	(4)
Pool size	-6.23e-08	0.00000617	0.000621***	0.00308***
	(5.71e-08)	(0.00000514)	(0.0000333)	(0.0000560)
Std.Dev Fx Rate	51.18***	23.92***	20.38***	9.268***
	(5.413)	(0.980)	(0.383)	(0.0171)
Lagged rel. Liquidity Flow	0.0540***	0.0772***	0.117***	-0.0315***
	(0.0168)	(0.0189)	(0.0139)	(0.00236)
Lagged liqidity use	0.0854***	0.0451***	0.0387***	0.000152***
	(0.00236)	(0.000670)	(0.000364)	(0.00000187)
R^2	0.400	0.310	0.278	0.189
Observations	2,347	12,262	45,803	1,396,910
height				

Table 5. Regression explaining liquidity use defined as daily volume over average daily poolsize. One, two, and three stars indicate significance at the 10%, 5%, and 1% level, respectively.

7 Exchange competition

To compare Uniswap to traditional exchanges we collect minute by minute treading data from Binance, one of the largest crypto exchanges by volume. Many of the 1,251 token 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.

Arbitrage between the two markets is not instant. Binance, like all 'traditional' crypto exchanges require to have custody of the traded assets. Any tokens that a user wants to trade on the exchange need to be transferred out the users personal wallet into the exchange wallet and the exchange ha to give the user credit for these assets in their own internal ledger before they can be traded. Once they are in the system of the exchange tokens can be traded with minimal delay and high frequency. Uniswap, in contrast, is non-custodial, meaning that the user initiate a trade directly out of their personal wallet and keep custody of traded assets until they are swapped in an atomic transaction. Since Uniswap is on-chain trading is tied to the transaction processing of the Ethereum blockchain. Ethereum is designed to be faster than Bitcoin with about 10-20 seconds between blocks, however, execution of trades on Uniswap can never be as fast as on traditional exchanges.

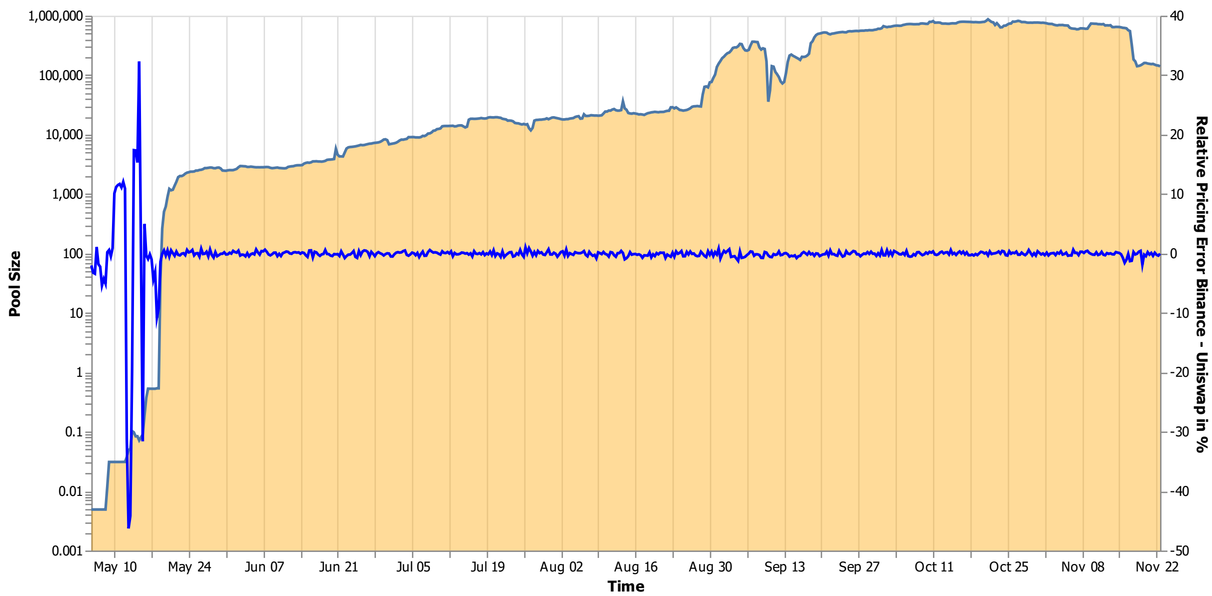

Figure 11. Pricing error and pool size Pricing difference for the USDC/ETH pair when comparing Binance to Uniswap in percent of the Binance price (blue line, right axis) and pool size of the Uniswap USDC/ETH pool (orange, log-scale, left axis).

Pricing differences between Uniswap and Binance are small except in the startup-phase of the Uniswap pool when liquidity is scarce. Figure 11 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 6. We examine the absolute percentage pricing error defined as the absolute

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Pool size	-0.00000115***						-0.000000202	-0.00000119***	-2.01e-08
	(0.000000214)						(0.000000143)	(0.000000159)	(0.000000141)
Std.Dev Fx Rate		12.86***					8.386***		8.182***
		(1.890)					(1.874)		(1.827)
Volume Binance			-0.000000700***				-0.000000744***		-0.000000733***
			(9.41e-08)				(0.000000108)		(0.000000106)
Relative Vol Uniswap				-3.908***			-3.105***		-3.340***
				(0.668)			(0.642)		(0.705)
(Relative Vol Uniswap)^2				3.073***			2.477***		2.640***
				(0.649)			(0.668)		(0.708)
Binance Price					-0.0189***		-0.00289		-0.0239***
					(0.00301)		(0.00477)		(0.00802)
Gas Price						1.57e-12**	1.54e-12***		1.90e-12**
						(7.00e-13)	(5.02e-13)		(6.88e-13)
Binance Price Impact								221.7	780.2***
								(218.4)	(203.8)
Uniswap Price Impact								2.257***	2.047***
								(0.142)	(0.0736)
R^2	0.0138	0.0761	0.00671	0.0847	0.00263	0.00916	0.139	0.0324	0.167
Observations	4,322	4,233	4,322	4,322	4,322	4,319	4,230	4,307	4,215
height									

Table 6: **Regression explaining absolute price difference between Uniswap and Binance markets. One, two, and three stars indicate significance at the 10%, 5%, and 1% level, respectively.**

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 volume may be indicative of higher liquidity. 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 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. We use the Binance Price as a reference point. 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. To trade on Uniswap users must pay the miner to record the transaction on the Ethereum blockchain. When mining costs, i.e. gas prices, are high small price differences are not profitable to arbitrage away.

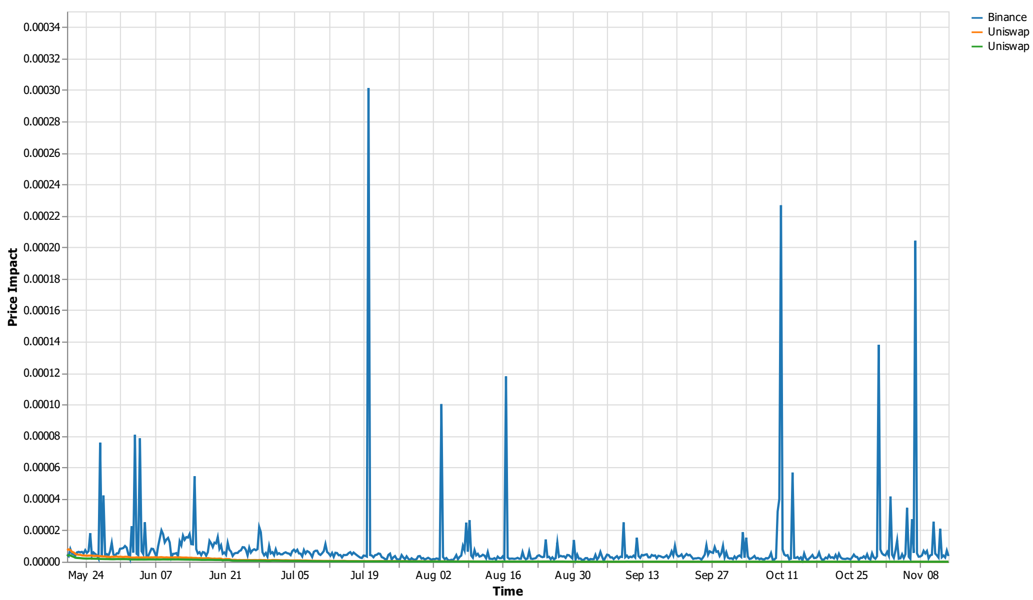

Figure 12. Price Impact of USDC/ETH on Uniswap (orange, green) and Binance (blue). Price impact is computed as change in price over volume (green and blue lines) as well as analytically as the price change for a marginal unit bought using the bonding curve formula (green line).

In columns (8) and (9) of Table 6 we examine price impact as 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 12 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).

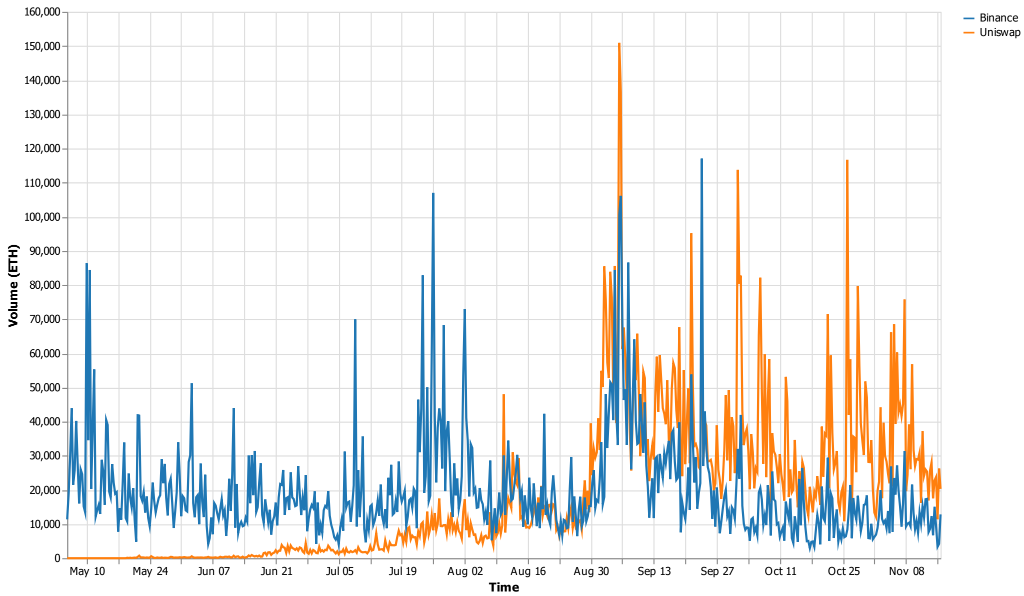

Figure 13. Trading volume of USDC/ETH on Uniswap (orange) and Binance (blue). The graph shows the trading volume excluding flash loans in ETH. Trading volume is aggregated over rolling eight hour intervals.

Figure 13 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 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 14 shows intraday prices of the USDC/ETH pair on one day, October 21, 2020. The patters 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.

8 Conclusion

In 1971, Fischer Black two articles for the Financial Analyst's Journal speculating on whether computers or "automation" could ever replace human interaction in financial markets (Fisher Black (1971*a*), Fisher Black (1971*b*)). 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

Figure 14. 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.

to have more direct participation from other market participants. The rise of "yield farming" and automated protocols such as Uniswap provide such opportunities.

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*.

Black, Fisher. 1971*a*. “Toward a Fully Automated Stock Exchange-Part 1.” *Financial Analysts Journal*, 27(4): 28–35.

Black, Fisher. 1971*b*. “Toward a Fully Automated Stock Exchange, Part 2.” *Financial Analysts Journal*, , (6): 86–87.

A Uniswap volume and the attack on Harvest Finance

On this day a hacker launched a large scale attack on Harvest finance, a yield farming cooperative. Users can deposit their tokens with Harvest finance in return for fAsset tokens (e.g. depositors of USDC receive a fUSDC token). The underlying tokens are then 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 converted over 17 million USDT to USDC on curve.fi (an automated market maker similar to uniswap that specialized in trading stablecoins).[9] This temporarily increased the price of USDC in the curve.fi pool, which is used a price oracle (i.e. source of market information) for Harvest finance. Since Harvest finance was misled by this market information it issued too many fUSDC tokens to the attacker upon their deposit of USDC 50 million. 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.[10]

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.

B Fake tokens

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 7 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.

[9] see transaction 0x35f8d2f572fceaac9288e5d462117850ef2694786992a8c3f6d02612277b0877.

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

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 7. 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 liquiditypools with these fraudulent tokens. *Nonfraudulent transactions* is the number of transactions in the original token in Uniswap pools.

C Measuring volume

Measuring volume on uniswap is not trivial fro two reasons. First, large trades, for example oracle attacks can push prices out of equilibrium making it hard to measure volume. Second, we observe transactions that trade large quantities back and forth artificially inflating volume. We show examples for both distortions and then outline our process to compute volume.

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 Warp Finance on December 17, 2020.[11] An attacker 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 40 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 whole purpose of this transaction is to cause a temporary difference in the true economic value of the two sides of the trade for the oracle attack. Yet it is not clear how to measure volume in this case. Based on the

[11]See transaction 0x8bb8dc5c7c830bac85fa48acad2505e9300a91c3ff239c9517d0cae33b595090.

fact that most of the trading in the Uniswap system is against WETH we use ETH as numeraire for computing volume. We then use external data from Binance to convert the ETH to USD. 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 alos have an opposing trade that brings prices back to equilibrium.

We also observe some transactions that trade tokes 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.[12] 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 windorize the data and thus eliminate such outliers.

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 using the prices from the pool with the highest volume.

D Proofs

Proof of Lemma 1

If the liquidity supplier is alone in the market, and posts at a price p then

i. With probability $(1-\alpha)$ a noise trader arrives. The order is filled with certainty and he obtains a payoff of $(p - p_0)$

ii. With probability $\frac{\alpha}{2}\delta$ there has been a relevant information event, and if he learns about it, in which case he adjusts his quote to $p_0 + \sigma$ and obtains a payoff of zero.

iii. With probability $\frac{\alpha}{2}(1-\delta)$ the limit order submitter is unaware that there has been an information event. 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 an order with price p_i is then

$$\pi(p) \quad = \quad (1-\alpha)(p-p_0) + \frac{\alpha}{2}(1-\delta)(p-p_0-\sigma).$$

[12]See transaction 0xa8c00a56cf2455241bbc4b5ef9e3f9e761cdbd7909847ab8274dcd9bd1dded6a.

Clearly, he will post a sell order at price $p_0+\sigma$, to obtain a profit of $(1-\alpha)\sigma$. ■

Proof of Lemma 2

A limit order submitter chooses a price p_i to maximize his expected profits, which comprises four elements:

i. With probability $(1-\alpha)$ a noise trader arrives. The other market maker j is present as well and limit order i gets his order filled with probability $(1-F_j(p_i))$ and obtains a payoff of $(p_i - p_0)$

ii. With probability $\frac{\alpha}{2}\delta$ there has been a relevant information event, and he learns about it, in which case he adjusts his quote to $p_0+\sigma$ and obtains a payoff of zero.

iii. With probability $\frac{\alpha}{2}(1-\delta)$ the limit order submitter is unaware that there has been an information event. 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 an order with price p_i is then

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

In equilibrium, it has to be that each price is offered with some probability and it is not optimal to deviate from that price. This it has to be that

$$\frac{\alpha}{2}(1-\delta) + (1-\alpha)(1-F_j(p)) - (p-p_0)(1-F_j'(p)) = 0 \qquad \forall p_i. \tag{17}$$

We can solve the differential equation in (17) with the boundary condition $F(p_0+\sigma)=1$, to get the symmetric equilibrium schedule:

$$F(p) \quad = \quad \frac{(p-p_0)(1-\frac{\alpha}{2}(1+\delta_i)) - \sigma\frac{\alpha}{2}(1-\delta)}{(p-p_0)(1-\alpha)}.$$

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

$$\begin{aligned} p_{min} &= \frac{p_0(1-\frac{\alpha}{2}(1+\delta)) + \frac{\alpha}{2}\sigma(1-\delta)}{1-\frac{\alpha}{2}(1+\delta)} \\ &= p_0 + \sigma\left(\frac{\frac{\alpha}{2}(1-\delta)}{1-\frac{\alpha}{2}(1+\delta)}\right). \end{aligned} \tag{18}$$

■

Proof of Proposition 3

Suppose that trader j invests monitoring technology, then trader i prefers to invest if

$$\begin{aligned} \gamma'(1-\gamma')(1-\alpha)\sigma - I &\geq \gamma(1-\gamma')(1-\alpha)\sigma, \\ \Delta_\gamma(1-\gamma')(1-\alpha)\sigma \geq I \end{aligned}$$

whereas, if trader j does not invest the condition is

$$\Delta_\gamma(1-\gamma)(1-\alpha)\sigma \geq I.$$

i. Suppose that $I \in [0, \Delta_\gamma(1-\gamma')(1-\alpha)\sigma]$ then both traders will invest.

ii. Suppose that $I \in (\Delta_\gamma(1-\gamma')(1-\alpha)\sigma, \Delta_\gamma(1-\gamma)(1-\alpha)\sigma]$, then either trader i or trader j invests but not both.

iii. Suppose that $I > \Delta_\gamma(1-\gamma)(1-\alpha)\sigma$ then neither trader will invest.

■

Proof of Lemma 4

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

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

$$F_c(x) = 1 - \left(\frac{\frac{\alpha}{2}(1-\delta_i)(p-p_0-\sigma)}{(p-p_0)(1-\alpha)}\right)^2$$

To determine the expected transaction price note that

$$\begin{aligned} Ec(q) &= \int_{p_{min}}^{p+\sigma} 1 - F_c(x)dx \\ &= \int_{p_{min}}^{p+\sigma} \left(\frac{\frac{\alpha}{2}(1-\delta)(p-p_0-\sigma)}{(p-p_0)(1-\alpha)}\right)^2 \\ &= \left(\frac{\frac{\alpha}{2}(1-\delta)}{(1-\alpha)}\right)^2 \int_{p_{min}}^{p+\sigma} \left(\frac{(p-p_0-\sigma)}{(p-p_0)}\right)^2 \end{aligned}$$

$$Ec(q) = \left(\frac{\frac{\alpha}{2}(1-\delta)}{(1-\alpha)}\right)^2 \left(-p_{min} + p_0 + \frac{\sigma^2}{p_{min} - p_0} + 2\sigma \ln(p_{min} - p_0) - 2\sigma \ln(\sigma)\right)$$

$\blacksquare$