Decentralized Exchanges

Presenter: QIHONG RUAN

Written by Alfred Lehar and Christine A. Parlour

March 7, 2022

Overview

- How Uniswap works? Automated Market Making (AMM) and liquidity pool
- Develop a model to compare and contrast limit order market and AMM, identifying the conditions under which an AMM is better than a limit order market
- Use high-frequency block-by-block data to verify the implications of its theory and visualize the transactions on Uniswap

1 Introduction

- Uniswap is a decentralized exchange that launched in November 2018
- A system of smart contracts on the Ethereum blockchain
- A novel model of liquidity provision, Automated Market Making (AMM), different from a limit order market
- Committed liquidity supply tops 3 billion USD among various cryptocurrencies
- Transactions worth over 700 million USD per day

Automated Market Maker (AMM)

- 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, a bonding curve
- The convexity implies that larger orders have a larger price impact
- All liquidity demanders pay a proportional fee to the liquidity suppliers

Two Key Differences between an AMM and a Limit Order Market

- 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
- 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 the price impact that maximizes their profits

The Equilibrium Effect of these Two Key Differences in a Volatile Market

- Risk neutral liquidity suppliers, a liquidity demander and an arbitrageur all interact
- In both markets, liquidity suppliers may be adversely selected as liquidity is posted before any potential asset innovation
- In a limit order book market, competing liquidity suppliers post prices to trade off adverse selection risk against profitable liquidity supply
- In the AMM, liquidity suppliers trade off liquidity fees and adverse selection of an informed arbitrageur

- In the limit order market liquidity suppliers retain all the price impact revenue from supplying liquidity
- In the AMM, arbitrageurs obtain this benefit and liquidity suppliers only earn the fees
- 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 assets that have lower volatility (and hence adverse selection) the AMM can be more effective (i.e., is cheaper) at providing liquidity

Literature about the Optimality of Market Rules

- This paper is related to the literature about the optimality of market rules
- Lawrence R. Glosten (1998) studies the optimality of time priority (FIFO) rule shows that pro rata rationing changes the marginal payoff to liquidity provision and hence the posted volume
- Richard Haynes & Esen Onur (2020) discuss the price efficiency of pro rata rationing using a natural experiment from the Treasury Futures market

Literature on Liquidity Provision

- This paper is also linked to the large literature on liquidity provision in limit order markets
- Since Lawrence R. 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
- 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, & Mayhew (2009) documents such rationing in options markets

Literature on Competition that does not Lead to Cheaper Liquidity

- The rise of high frequency traders has generated research into competition that does not lead to cheaper liquidity
- Biais, Foucault, & Moinas (2015): market competition on speed generates both positive and negative externalities
- Hendershott & Riordan (2019): high limit order submission and cancellation, consistent with strategic liquidity provision

The Model Setting is Different

- This paper also connects to the papers analyzing the theoretical properties of constant function market makers
- Ito & Aoyagi (2021) present a model in which traders choose between a centralized 'market maker' market and an automated market maker
- By contrast, we take the population of traders in each market as given

The Focus is Different

- Most closely related to our work is Agostino Capponi & Ruizhe Jia (2021)
- They present a model and test of an AMM with a focus on competition among arbitrageurs, allowing 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 focus on the tradeoff between liquidity fees and adverse selection

Our Data are more Extensive

- Andreas Barbon Angelo Ranaldo (2021) compares transaction costs and price efficiency on various decentralized exchanges and Binance for five different token pairs, concluding that due to the fixed cost of gas fees, trading costs on the DEX are higher than the CEX
- Their data only permits them to use the median gas fee over the entire sample, which as we point out, overestimates the cost of transacting

Constant Product and Automated Market Making

- We describe the market making mechanics on Uniswap V2
- Eth, the native cryptocurrency on the Ethereum Blockchain
- We use Eth as the numeraire, and other coin as "token"
- A liquidity provider deposits both Eth ad the token into the pool
- The deposit ratio is determined by the existing ratio, which implicitly defines the Eth price of the token
- Liquidity provider receives a proportional amount of a liquidity token
- 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
- Although adverse selection, providing liquidity is potentially profitable because each trade faces a tax of 30bps which is redeposited into the pool

- To buy the token, deposit Eth into the pool, and withdraw the token
- The amount that he has to deposit or withdraw depends on the bonding curve

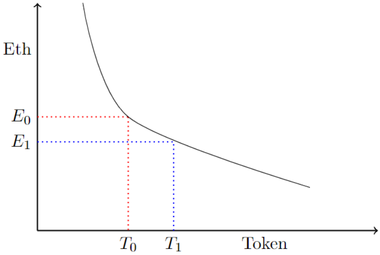

- The trader deposits $T_1 - T_0$ of the token into the pool, and he would receive $E_0 - E_1$ Eth
- Arbitrageurs trade in the opposite direction to return the ratio to equilibrium

- Constant product $k := T_1 \cdot E_1 = T_0 \cdot E_0$
- For each pool, the constant k, depends on the amount of liquidity that has been deposited inthe pool up to this point
- If more liquidity is posted, the constant changes

Assessing Liquidity Fees

- Suppose that an agent wants to trade e Eth in exchange for tokens. The exchange collects a fee τ, which benefits liquidit holders
- The effective amount of Eth that gets traded is $(1-\tau)e$
- Before fee revenue liquidity pool balance $E' = E + (1-\tau)e$
- Following the bonding curve, the post trade token balance must be

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

- The amount of token t that the trader receives is

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

- The terms of trade expressed in Eth/token is

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

- The terms of trade have a natural interpretation as a spread
- Suppose the fundamental value of the token denominated in Eth is p_0, then in equilibrium $p_0 = \frac{E}{T}$
- 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

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

- When buying tokens, traders have to pay a fixed spread of $\frac{1}{1-\tau} p_0$. 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 $\partial p/\partial e = \frac{1}{T}$. As the liquidity pool grows, the price impact of a fixed order size decreases
- The blue line: a small pool; the orange line: a large pool
- Larger orders face worse marginal prices as they do in a limit order book

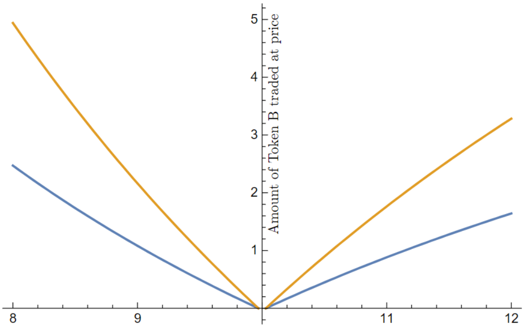

Framework

- 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 potentially informed trader monitors the market and trades whenever profitable
- A liquidity trader trades a fixed quantity q and is equally likely to buy or sell, at any price $p \in [p_0 - \sigma, p_0 + \sigma]$
- 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

- Two liquidity suppliers as it is the minimum required for competition
- 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
- We consider the case where informed trader buys. The case where informed trader sells is symmetric

Limit Order Market

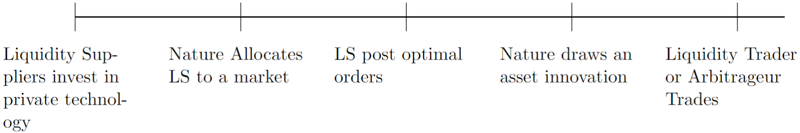

Figure 3. Sequence of Events in the Limit Order Market

- The amount that the liquidity trader trades is fixed, q, and so this is also the amount that the liquidity suppliers post
- 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
- A monopolist liquidity supplier in the market, will post a sell price of $p_0 + \sigma$, and a buy price of $p_0 - \sigma$ 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 trade
- In this way, rivalrous liquidity provision will make them aggressively undercut. The symmetric equilibrium is in mixed strategies
- 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 $[p_0 + \alpha\sigma, p_0 + \sigma]$, where

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

- A symmetric expression holds for competing liquidity buyers. Each competing liquidity suppliermakes zero profits

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

Automated Market Maker

- In the AMM or bonding curve market, liquidity suppliers choose a market 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 was in the limit order market
- The second thing to observe is that the AMM requires committed capital
- We will start our analysis under the assumption that the price in the bonding curve market is equal to the equilibrium price, or $b_0 = p_0$. $\frac{E_0}{T_0} = b_0$, $E_0 T_0 = k$, $(E_0 + e)(T_0 + t) = k$
- k is the constant of the bonding curve. 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

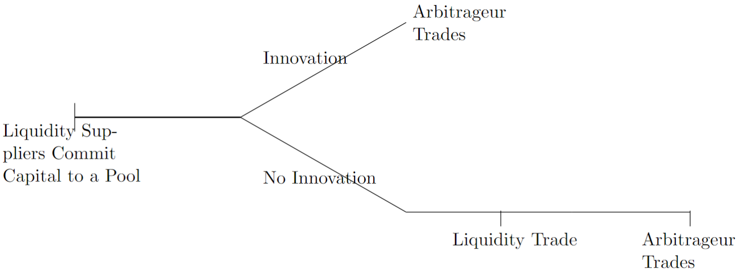

Figure 4. Sequence of Events in the Automated Market Maker

- If the liquidity trader buys tokens, they will remove q tokens, and will buy these with Eth and add in e_l^b to the Eth pool

$$(E_0 + e_l^b)(T_0 - q) = E_0 T_0$$

$$e_l^b = \frac{E_0 T_0}{T_0 - q} - E_0$$

- The fees paid by this noise trader are τe_l^b

- An arbitrageur will reverse this trade, adding tokens and remove Eth
- The arbitrageur also pays the fee
- the payoff to liquidity provision is twice the fee paid by the liquidity trader, or $2\tau e_l^b$
- With probability $(1-\alpha)$, there is a liquidity event and the fee revenue for liquidity provision is:

$$2\tau p_0 q(\frac{T_0}{T_0^2 - q^2})$$

- 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
- She will buy t^b tokens and pay e_I^b for them such that $(T_0 - t^b)(E_0 + e_I^b) = E_0 T_0$
- Her profit function is

$$\pi_I^b = (p_0 + \sigma)t^b - (1+\tau)[\frac{E_0 T_0}{T_0 - t^b} - E_0]$$

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

$$(p_0 + \sigma) - (1+\tau)\frac{E_0 T_0}{(T_0 - t^b)^2} = 0$$

- The optimal transaction amount

$$t^b = T_0 - \sqrt{\frac{(1+\tau)E_0 T_0}{p_0 + \sigma}}$$

$$e_I^b = \sqrt{\frac{E_0 T_0 (p_0 + \sigma)}{(1+\tau)}} - E_0$$

- In order to have $t^b \geq 0$, we should have $\tau \leq \frac{\sigma}{p_0}$, so that the transaction cost is low relative to the information
- After the innovation, absent informed trading, the Eth value of the total supplied capital would be $E_0 + (p_0 + \sigma)T_0$
- Given the informed trade, the Eth value of the supplied capital is

$$E_0 + e_I^b + (p_0 + \sigma)(T_0 - t^b) = \sqrt{E_0 T_0 (p_0 + \sigma)}(\frac{2+\tau}{\sqrt{1+\tau}})$$

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

$$\frac{2+\tau}{\sqrt{1+\tau}}\sqrt{T_0 E_0 (p_0+\sigma)} - (E_0 + (p_0+\sigma)T_0)$$

- This change in value corresponds to "picking off" risk, in the sense that the informed trader rebalances the amount of Eth and Tokens to reflect the value in the wider market
- However, the arbitrageur pays a liquidity fee for an amount equal to $\tau(\sqrt{\frac{E_0 T_0 (p_0+\sigma)}{(1+\tau)}} - E_0)$
- Balancing the picking-off risks and liquidity provision benefits, the liquidity providers choose the equilibrium size of the pool

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

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

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

- The parameter restrictions 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 information, and will also not trade to ensure that the price implied by the pool corresponds to the true value of the token
- Assume that these conditions hold in empirical analysis

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 handover custody of their digital assets to a third party
- the source code for many DEXs is public and users can verify that the code is not fraudulent

- Uniswap V1 was launched in November 2018, and the first pool is ETH/MKR
- Uniswap V2 was launched on May 18, 2020
- The update allows direct trade of any ERC-20 token pairs and includes Tether (USDT)
- 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
- Uniswap has no designated operator, V1 pools cannot be deleted from the blockchain and exist in parallel to V2 pools
- UniSwap liquidity pools are not certified, users should verify the address to trade with the correct tokens

- 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 blockchain
- 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 of liquidity from a pool, and 45,481,500 trades of tokens
- A Uniswap transaction is a set of instructions that are processed in the same block
- We focus on the analysis on liquidity supply on both a limit order market and an AMM. So we do not consider market access fees like gas fees

- Transactions are finalized if they are incorporated onto the Ethereum blockchain, and anything that happens in the same block effectively happens at the same time
- In our data, 56,606 transactions combine liquidity additions or removals with swaps or flash swaps
- 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

Ten Largest Exchanges for Uniswap V1 and V2

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

Number of transactions on Uniswap

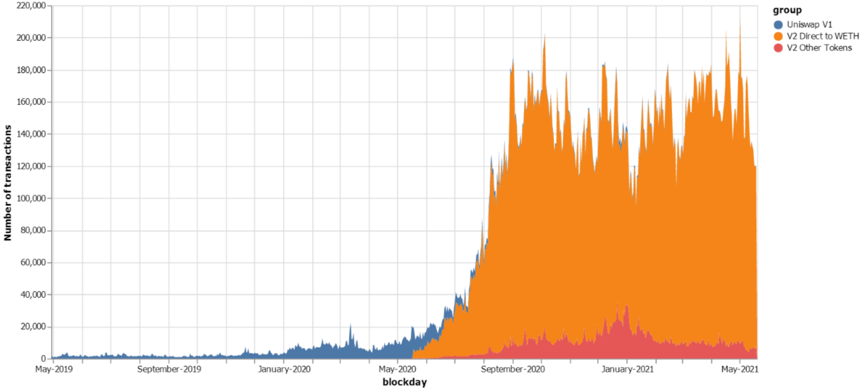

Trading Volume on Uniswap

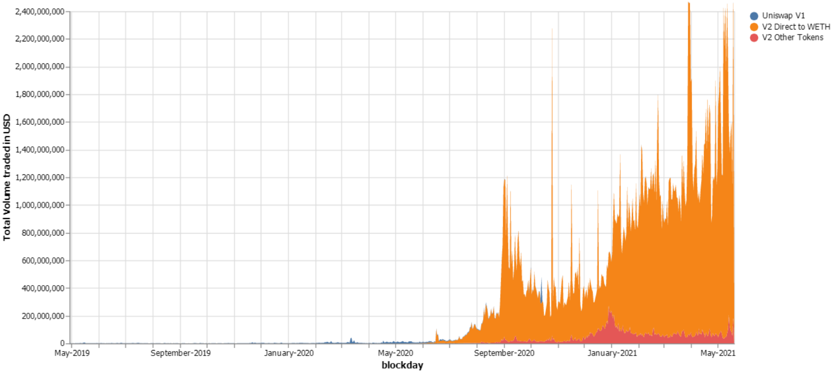

Network graph between all tokens of the 50 largest pools

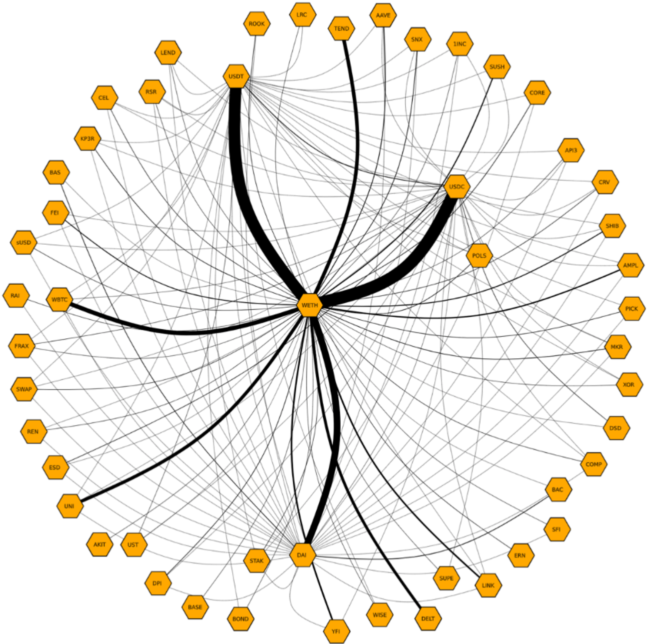

Liquidity Provision

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

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

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

- Empirically there is heterogeneity in pools
- We derive comparative statics for pool size. Suppose that $\tau < \frac{\sigma}{p_0}$, then the equilibrium size of a liquidity pool is
  1. Linear in the size of the liquidity trade
  2. Is decreasing in the size of the innovation, σ
  3. Is decreasing in informed trades, α

- To test these predictions, we collect daily data on all 1,376 pools in our 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

	(1)	(2)	(3)	(4)	(5)
Volatility	-14646277.8***		-14193779.4***	-13976232.8***	-15986481.5***
	(1907118.6)		(1742450.0)	(1615636.4)	(2208943.9)
Volume (USD)		0.255***	0.255***		
		(0.0739)	(0.0739)		
Number trades				3051.9**	
				(1521.2)	
Reversals					18963.9**
					(9073.6)
R^2	0.000925	0.0498	0.0507	0.0338	0.0264
Observations	263,750	279,040	263,750	263,750	263,750

- 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
- Ethereum gas fees act as a commitment device for liquidity suppliers to remain in pools

Ranking Exchanges

- 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
- 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 $1-\eta$
- The expected cost to the liquidity trader is

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

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

- The expected cost to the liquidity trader is

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

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

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

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

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

- The limit order market does not dominate the AMM
  1. There are robust parameter ranges for which the automated market maker dominates the limit order market
  2. Conditional on trading quantity q, for a pool in equilibrium, price impact is more volatile in the limit order market than the AMM

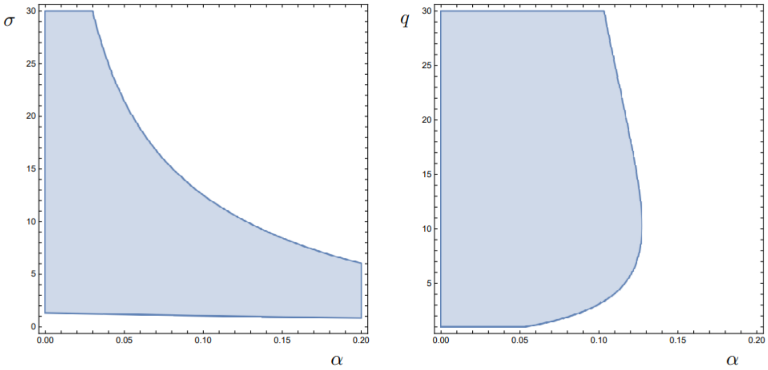

- The left panel shows that automated market makers are the better trading venue for the liquidity trader when either the innovation in prices or the intensity of informed trading are sufficiently small
- The right panel shows that for a given intensity of informed trading liquidity pools dominate when the trade size of the liquidity trader is neither too 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

Empirical Analysis

- We end up with 27 token pairs that are cross listed and frequently traded on Uniswap and Binance
- When the pool starts, as long as the pool size 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
- Once the poolsize is above 700 ETH, the pricing difference stays below 1% with an average of -0.026% for this pool

USDC/ETH Pool Size and Pricing Error

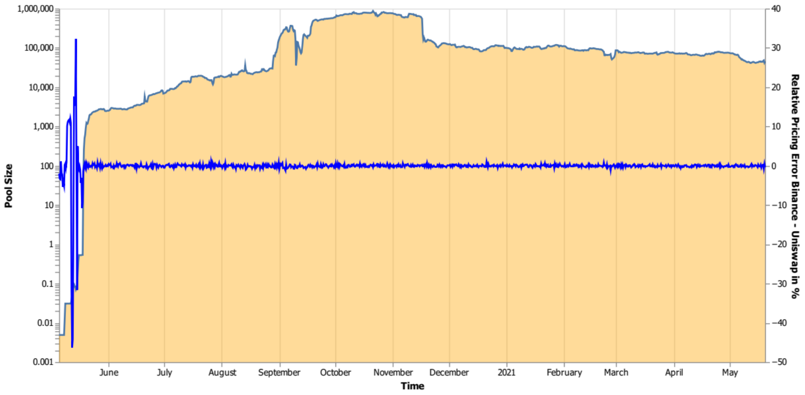

Determinants of price differences between Binance and Uniswap

	(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

- 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
- Decreases in pool size and volume
- Increases in fx-volatility, gas price, and price impacts
- Non-monotonic in relative volume

Price Impact of USDC/ETH on Uniswap (orange, green) and Binance (blue)

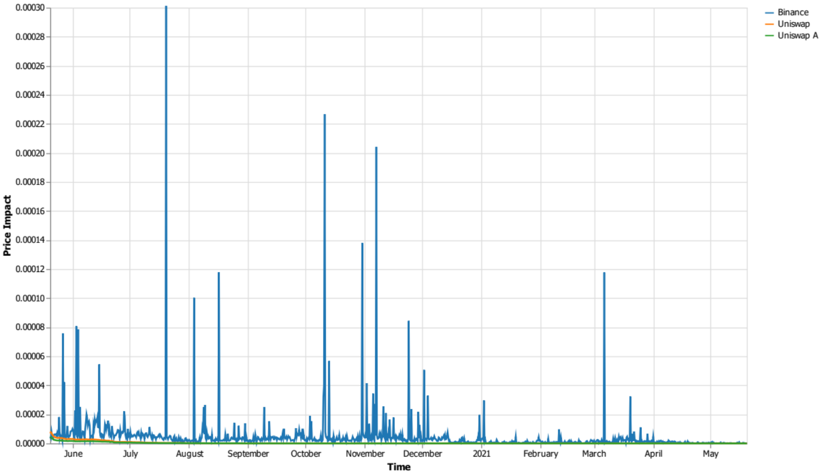

Trading Volume of USDC/ETH on Uniswap (orange) and Binance (blue)

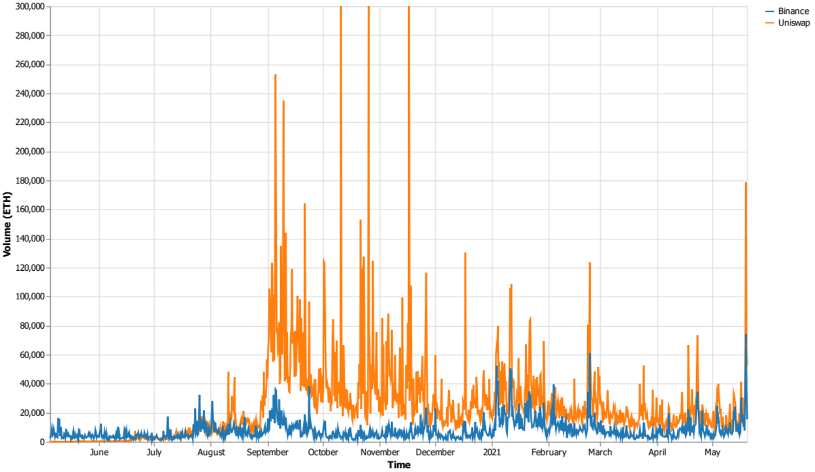

- Remarkably correlated across the two markets
- Uniswap is gaining market share over time

Intraday Prices for the USDC/ETH Pair on 10/21/2020

- Binance prices are often leading Uniswap prices
- Binance prices are more volatile that the prices on Uniswap

Conclusion

- 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
- Price impact in a limit order market is the endogenous outcome of the interaction between liquidity demand and supply, while on an AMM the price impact is programmatically determined by a bonding curve
- Liquidity suppliers trade off potential adverse selection against fee revenue to adjust the equilibrium pool size
- Uniswap V3 will give liquidity suppliers partial ability to associate their liquidity to price ranges
- This change re-introduces competition between liquidity suppliers, which may drive out non-strategic liquidity suppliers

Comments

- I have learned how AMM works, interesting data of Uniswap and Binance, and the framework to compare different trading rules
- This paper ranks limit order market maker and AMM based on the total trading cost of liquidity trader, while the trading cost of limit order market depends on the parameters of private investment, which are not relevant to the AMM. The change of these parameters may alter the half-and-half result of limit order market v.s. AMM
- Technically, when the market maker is monopolistic, they can set a price even higher than $p + \sigma$ to gain more profits
- What if we change the ranking criterion such as including the cost to build centralized exchanges or decentralized blockchains? Then which system is socially optimal under what conditions?
- This paper is about the differences between limit order market and AMM, but does not directly touch the notion of decentralization since an AMM can be organized by a centralized exchange as well. How to model the value of decentralization, such as self-custody and being permissionless?