THE BLOCK · Research

Pooled Liquidity Provision in DeFi: Concentrated Liquidity

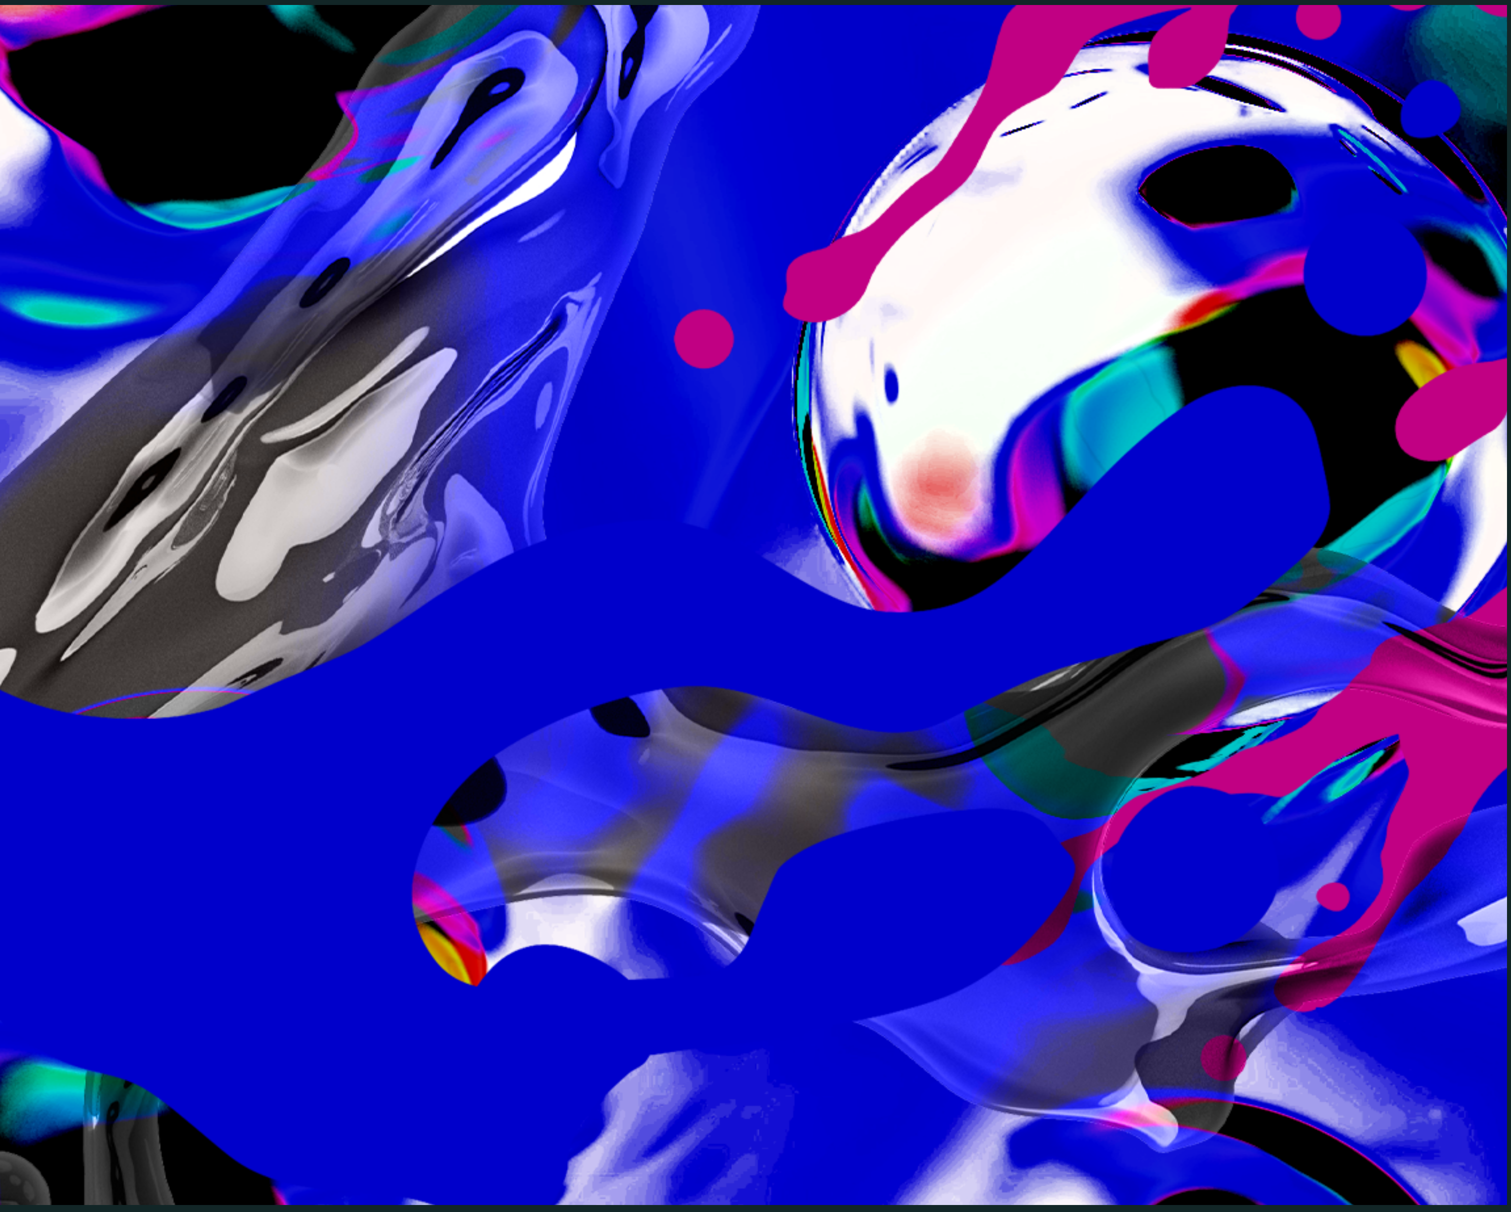

COMMISSIONED BY ORCA

Executive Summary

This report provides insights into concentrated pooled liquidity provision on Automatic Market Makers (AMMs). AMMs have revolutionized decentralized exchanges (DEXes) and have turned out to be a major use case in DeFi, to the extent that they rival liquidity on centralized exchanges (CEXes). After Uniswap v3's introduction of concentrated liquidity in 2021, it quickly gained market share among DEXes. Compared to infinite price range liquidity provision, concentrated liquidity can be much more capital efficient. However, the flipside is an increasing risk of impermanent loss for liquidity providers.

Concentrated liquidity providers are given more degrees of freedom to set their pool position. This leads to several different strategy setups, some of which may be better confined to sophisticated investors, or active liquidity management protocols. That seems to apply, for example, to liquidity provision of volatile token pairs, which may require frequent rebalancing to be profitable. Automatic liquidity management protocols may have advantages because of lower transaction cost once their fund size increases, but also because they claim to offer more sophisticated management for retail liquidity providers. However, it remains to be seen whether their total return –including protocol fees, impermanent loss, and 'compounded' contract risk– offers value for money and adequately compensates liquidity providers.

Empirical results for concentrated liquidity provision are so far rather disappointing. In most studies it does not consistently outperform simple buy and hold strategies. This is mainly driven by impermanent loss. One possible exception are token pairs, which feature low risk of impermanent loss, such as stablecoins. The main advantage from higher capital efficiency of concentrated liquidity so far seems to accrue to swappers, who get much less slippage for their trades, even compared to centralized exchanges. Going forward that situation will likely redress, be it through more efficient strategies for liquidity provision, or a reduction in liquidity provision until liquidity providers are adequately rewarded for the risks they take.

About

Commissioned by Orca, the DEX for people, not programs.

Orca is the simplest swap and best concentrated liquidity experience in DeFi. Swim on over to orca.so and experience it first hand.

Researched by The Block and The Block Research

The Block is an information services company founded in 2018. Its research arm, The Block Research, analyzes an array of industries including digital assets, fintech, and financial services.

Contact

Email: research@theblockcrypto.com

Twitter: @theblockres

Author

Marcel Bluhm - Research Analyst

Twitter: @bluhmus

Table of Contents

Section 1: Introduction **........ 5**
Section 2: AMM Approaches and Concentrated Liquidity **........ 8**
First Generation Approaches: Key Invariants 8
Second Generation Approaches: Concentrated Liquidity 10
Section 3: Concentrated Liquidity Provision and Empirical Evidence **........ 17**
Strategies for Concentrated Liquidity Provision 17
Empirical Evidence 19
Conclusion and Outlook **........ 22**

Section 1: Introduction

Exchanging tokens on DEXes without the need for intermediaries such as banks, brokers and custodians was one of the first key use-cases that emerged in decentralized finance (DeFi). The largest decentralized exchange is Uniswap, which was set up in its first version (v1) in 2018, and operates on the Ethereum (compatible) ecosystem. In 2021 Uniswap v3 launched its concentrated liquidity feature, which is the focus of this report.

Concentrated liquidity became quickly adopted as a leading solution in the crypto DEX space. Since then, several variants to this model have also emerged on other crypto ecosystems such as Solana. While DEX activity ebbs and flows with overall activity in the crypto space (with a peak during 'DeFi-Summer' 2020), there has been a clear upward trend. Figure 1.1 displays the ratio of decentralized to centralized crypto exchange volumes.

Figure 1.1: DEX to CEX Spot Trade Volumes (%)

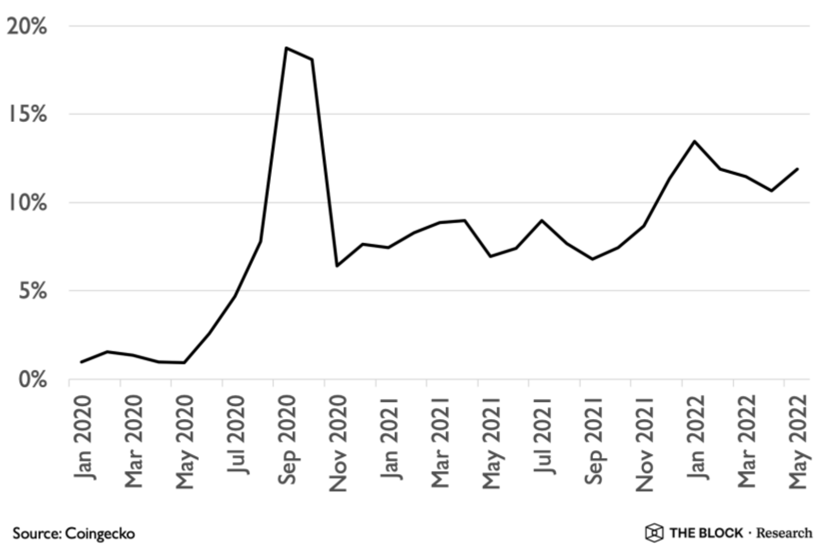

The DEX-CEX volume ratio was negligeable at the beginning of 2019, but it has increased to 12% and is on an upward trajectory. DEXes are an important building block for most investment activities, which require exchanging and operating in different tokens and stablecoins. Therefore, they are at the heart of numerous chains' sprawling ecosystems.

Figure 1.2 displays the major DEXes by volume.[1] Total DEX volume was $131bn in May 2022. The largest DEX by volume is Uniswap v3 ($61bn volume in May). In addition, two more exchanges in the top five by volume provide concentrated liquidity (Orca, $5bn in May) or use a very similar concept (Curve, $14bn).[2] These figures provide an indication for

[1] On the figure, the 16 Others are: Uniswap v1, Uniswap v2, Balancer, 0x, Synthetix, dYdX, Kyber, Bancor Network, Loopring, Serum DEX, Dodo, Astroport, Trader Joe, SpiritSwap, Raydium, and Solidly.

[2] The volume figures for Curve and Orca contain transactions based on both concentrated and non-concentrated liquidity.

the significance of AMMs to the crypto space, as well as the growing importance of concentrated liquidity, which is described in more detail in Section 2.

Figure 1.2: DEX Volume ($bn)

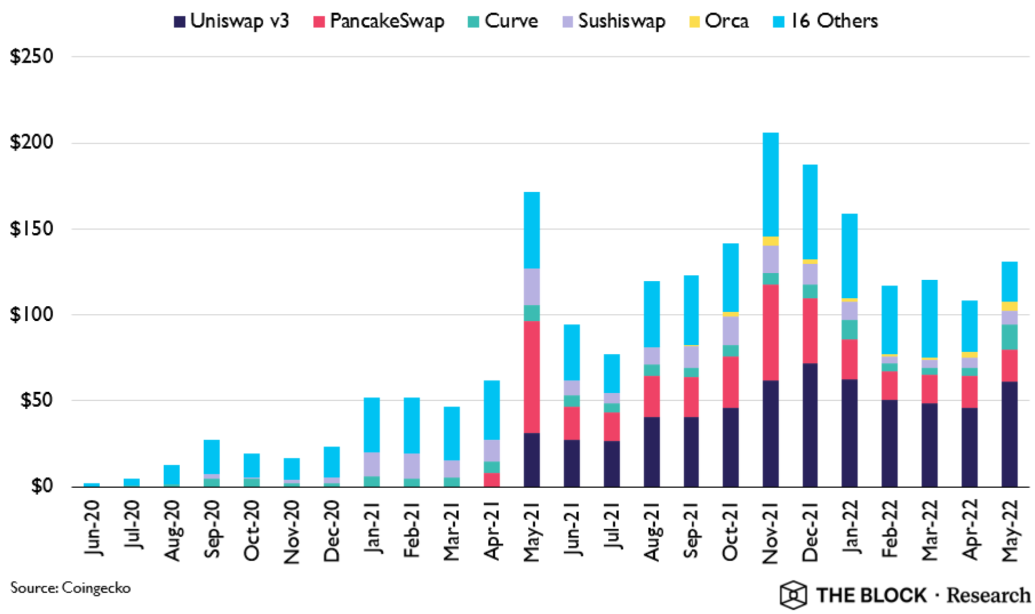

The remainder of this section outlines the concepts of AMMs and liquidity pools, which are key to analyzing pooled (concentrated) liquidity provision in DeFi.

Automatic Market Makers

Pooled liquidity provision is strongly tied to the concept of automatic market making, which operates differently from traditional financial markets. Traditional markets rely on an order book of buyers and sellers, which provide bid and ask prices with a spread. For example, when a buyer on an exchange with a central limit order book such as Binance finds an acceptable ask price for a token she wants to purchase, the underlying token is exchanged between her and the market participant (usually a market maker) offering the token at that price, which then becomes the new market price.

AMMs allow permissionless trading of self-custodied digital assets 24/7 on-chain. Since liquidity on-chain can be insufficient,[3] and transactions can be expensive (e.g. to adjust the order book when the price moves), AMMs use a liquidity pool, which helps to automate trades without relying on specific bids or asks from market participants. On an AMM there is no direct interaction between a buyer and a seller, but each transaction takes place between a market participant and a liquidity pool. This allows DEXes to operate without a central intermediary but to use smart contracts on a distributed ledger instead. For this

[3] Liquidity indicates how fast one asset can be converted into another without impacting the market price: Markets with deep liquidity allow for quick and large trades without affecting the market price.

reason, the blockchain (or layer) the DEX operates on is an important consideration and impacts the user experience regarding transaction cost and throughput significantly.

Liquidity Pools

Liquidity pools are supplied by liquidity providers, which receive transaction fees as a reward. For example, a pool may consist of two tokens that can be exchanged against each other by market participants, such as WBTC against USDC. Transaction fees are paid by market participants, who use the pools to exchange tokens. To fix ideas, consider a person that wants to exchange WBTC against USDC. Using an AMM, she can provide WBTC into the pool in exchange for USDC, which she obtains from the pool at an exchange rate that is determined by a mathematical formula.[4] A certain percentage of the exchanged amount then accrues as a transaction fee to the pool's liquidity providers to incentivize capital provision.

AMMs have revolutionized DEXes and have turned out to be a major use case in DeFi to the extent that they rival liquidity on centralized exchanges.[5] The remainder of this report dives into pooled (concentrated) liquidity provision (Section 2) and gives insights into strategies for concentrated liquidity provision and its empirical evidence (Section 3). Section 4 provides closing thoughts and an outlook for pooled liquidity provision.

[4] Section 2 shows different formulae that have been developed for AMMs, usually to provide capital efficiency for different use-cases.
[5] See "Liao and Robinson: [The Dominance of Uniswap v3 Liquidity]".

Section 2: AMM Approaches and Concentrated Liquidity

This section provides background on key AMM setups, the concept of impermanent loss, and then turns the focus on concentrated liquidity. It also gives an overview of concentrated liquidity AMMs operating on different blockchain ecosystems.

First Generation Approaches: Key Invariants

Several AMM approaches, which rely on different mathematical formulas to price assets, have emerged over the last few years. Formulas are typically represented as invariants to express a relation between the reserves of two (or more) tokens.[6] Table 2.1 provides an overview of key 'first generation' AMM approaches.[7] The invariants displayed cover the simplest case of two tokens, x and y, which can be exchanged. k is a constant and a is a scaling parameter.

Table 2.1: First Generation AMM approaches[8]

Invariant	Description	Popular Dex for Invariant
$xy=k$	Constant product formula: The product of token quantities x and y remains constant before and after trades. Most AMMs, including Uniswap V2, Pancakeswap, and Sushiswap are based on this design, which is effective for token pairs featuring a considerable price range.	Uniswap
$x+y=k$	Constant sum formula: The sum of x and y remains constant, which results in fixing the token pair price with no slippage. This design is used for assets that trade in a very tight price range, such as stable coins pegged to the same unit of account. It is less robust if one of the tokens depegs, which results in the pool quickly becoming imbalanced.	Mstable
$a(x+y)+xy=k(a)$	Hybrid formula: The hybrid formula combines the constant product and constant sum invariants and aims to provide the advantages of both approaches. It can operate like a constant sum formula, when two tokens are expected to trade nearby their peg, but converges to the constant product formula if the peg does not hold for one of the tokens (see Figure 2.1).	Curve v1

Notes: x are the reserves of one asset, which can be exchanged against the reserves of another asset in the pool, y. k is a constant, and a is a scaling parameter.
Source: "YieldSpace: An Automated Liquidity Provider for Fixed Yield Tokens". THE BLOCK · Research

[6] See 'YieldSpace: An Automated Liquidity Provider for Fixed Yield Tokens'.
[7] Further details on the AMM design space are provided in "The Block Research: Exploring Constant Function Market Maker Designs".
[8] The table outlines the three most established key invariants. Further approaches exist, see "The Block Research: Exploring Constant Function Market Maker Designs".

The **constant product formula** was implemented in the Uniswap protocol on Ethereum. Note that the formula can be extended to feature pools with more than two tokens.[9] This formula is most effective for tokens, which can trade with each other over extensive price ranges, such as, for example, a pair consisting of a stablecoin and a more volatile token (e.g. USDC/WBTC). Compared to other invariants, slippage can be relatively high.[10]

The **constant sum formula** by contrast is optimal for assets that are expected to always trade in a very narrow price range with each other, such as two stablecoins that are pegged to the same currency (e.g. USDC/USDT). MStable protocol follows this approach on Ethereum and Polygon.

Curve (v1) uses a **hybrid formula**, which combines features of both, the constant product and constant sum formulas, depending on a tuning parameter, a. The tuning parameter allows to cover a range of cases including token pools, where the price should allow considerable variation (tuning parameter set low), but also those which feature a constant relation to each other (tuning parameter set high). This design can concentrate liquidity in a narrow price range around the peg between stable pairs and thereby allows for lower slippage compared to the constant product AMM approach, while still offering liquidity off-peg.

Figure 2.1: Key Invariants

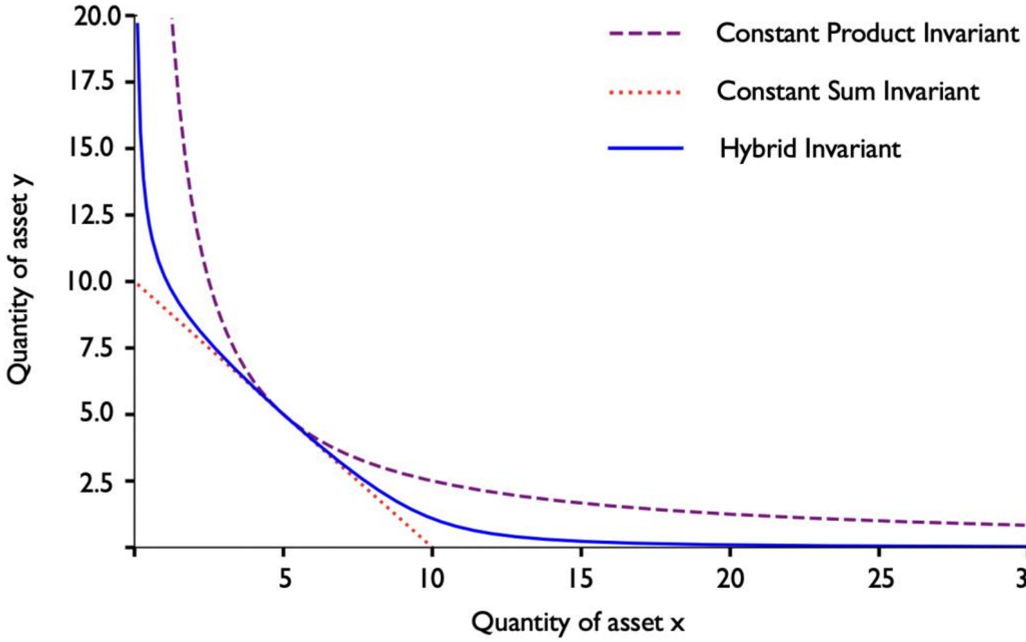

Source: "Egorov: StableSwap – Efficient Mechanism for Stablecoin liquidity"

THE BLOCK · Research

[9] See, for example, "Balancer: Whitepaper".

[10] Slippage is the difference between the expected price of a trade (such as the current market price) and the trade at which it is actually executed. There is no slippage on central order book exchanges as long as trades are executed within one tick.

Figure 2.1 visualizes the three key first-generation AMM invariants outlined above. Note that the functions' slope (or first derivative) indicates slippage when exchanging token pairs. In that example, when the pool is perfectly balanced with 5 units of each token x and y, all invariants' slopes are the same, resulting in the same slippage. For example, taking the constant sum invariant (red dotted line), a trader can exchange one unit of token x against one unit of token y, resulting in a new point on the invariant (x=4, y=6), which features the same slippage. However, while the constant sum invariant maintains an exchange ratio of one-to-one (slope remains constant), the other two invariants' exchange ratios become more unfavorable as the pool becomes more imbalanced. When imbalanced, the hybrid invariant (solid blue line) still provides less slippage than the constant product invariant (dashed purple line) over most of the trading range. Yet it also provides liquidity across a much wider price range than the constant sum invariant.

Impermanent Loss

All liquidity providers, including those on centralized exchanges, face the risk of impermanent loss.[11] Impermanent loss can be defined as "the difference between the value of the current fee adjusted liquidity position in an AMM, and the value of the position that was originally contributed had it simply been held instead of providing liquidity ('HODL value')." In many instances, this definition is further narrowed down to apply to constant product AMMs.[12] The broad intuition behind impermanent loss is that "an AMM always sells the outperforming asset and buys the underperforming asset".[13] In trending markets, this results in acquiring more and more of the underperforming asset at the expense of the better performing one. The Appendix contains examples on impermanent loss and shows that liquidity providers on any market, including those operating with central limit order books (CLOB), can be prone to it.

Second Generation Approaches: Concentrated Liquidity

The key-innovation of concentrated liquidity AMMs is to allow for liquidity provision over a custom price range. In traditional AMM constant product-type platforms ('traditional AMMs') such as Uniswap v2, liquidity is distributed uniformly over the price range from 0 to infinity ('infinite price range liquidity provision'). Allocation over a custom price range allows for different strategies ranging from narrow price ranges (approximating a market maker using a CLOB) to infinite-range liquidity provision (traditional AMM model). As shown in Appendix A and outlined below, provision of the same amount of liquidity over a narrower price range results in higher capital efficiency but magnifies the risk of impermanent loss. Unlike liquidity provision on traditional AMM platforms covered in the previous sub-section, this results in strategic competition for fee income between liquidity providers of the same pool.

---

[11] Some DEXes, for example Orca, use the term 'divergence loss' instead of impermanent loss.
[12] See, for example, "Topaze.Blue: Impermanent Loss in Uniswap V3".
[13] See "Topaze.Blue: Impermanent Loss in Uniswap V3".

Concepts for Concentrated Liquidity[14]

A concentrated liquidity pool depends on several parameters, some of which can be chosen by liquidity providers. Broadly speaking, those parameters are tick size, tick spacing, price range, and fee level.

A **tick** is the smallest amount by which a token's price can move up or down when traded.[15] Defining the **tick size** is at the discretion of the protocol. For example, in both Orca and Uniswap v3 tick size is defined as a "1 basis point price movement away from its neighboring tick".[16]

Unlike traditional AMMs, concentrated liquidity allows for choosing a **price range** over which to provide liquidity. The price range itself is subdivided into equally (pre-defined) **tick spaces** (each of which is a multiple of the defined tick size). A trade taking place within a tick space (not using up all liquidity available within that tick space) still follows a constant product invariant. However, once the price moves to a new tick, that formula is updated, which imposes a cost on the person initiating the swap. Hence there is a trade-off between precision and transaction costs: more granular tick spacing allows for more precise ranges but may cause transaction costs to increase. It is worth noting that more granular tick spacing can approximate a traditional CLOB market, where market makers choose how much liquidity they provide on a tick-by-tick basis.

Tick spacing may also depend on the (protocol-specific) **fee levels** in the pools. For example, in Uniswap v3, 1% pools, 0.3% pools, and 0.05% pools have a tick spacing of 200 ticks, 60 ticks, and 10 ticks, respectively. 0.1% pools, which were added subsequently, allow for a tick spacing of 1 tick.[17] Orca currently has two fee levels in its pools, 0.01% with a default tick spacing of 1 and 0.2% with a default tick spacing of 64.[18] A fee is charged to the market participant, who initiates the swap. While on Uniswap the fee accrues entirely to the liquidity provider, Orca uses this to also cover a protocol fee.[19]

Concentrated liquidity results in a histogram-like liquidity distribution, with the bin size determined by tick spacing. Compared to constant product-type AMMs, which equally distribute liquidity on the infinite price range, concentrated liquidity results in higher capital efficiency. This efficiency gain can be very pronounced for the case of stablepairs, such as two stablecoins both of which are pegged to the USD. Unless the peg of one of the

---

[14] This sub-section draws on "Uniswap: Uniswap V3 Core". Most concentrated liquidity AMMs build on that seminal paper. See, for example, "Invariant: First AMM DEX providing concentrated liquidity on Solana". For Orca see "Orca: Architecture Overview".
[15] See "Investopedia: Tick".
[16] See "Uniswap: Uniswap V3 Core". Once a pool is initialized for a specified price, all possible ticks can be derived from that initial price.
[17] While there is no published documentation on the tick spacing for the 0.01% pool, the range of existing stablepairs, such as USDC/USDT, can be set in increments of 0.01%.
[18] In the future pool owners may set fee level and tick-spacing, or follow default parameters set by the protocol.
[19] See "Orca: Understanding Whirlpool Fees". 3% of fees are currently going into an Orca Impact Fund, while 97% accrue to liquidity providers.

stablecoins breaks, the active price range is confined close to unity, so there is no need to make liquidity available outside that close range. Figure 2.3 displays the difference between traditional constant product-type AMM liquidity (blue surface area) provision and concentrated liquidity provision (gray surface area).

Figure 2.3: Constant Product-Type AMM liquidity and Concentrated Liquidity Provision

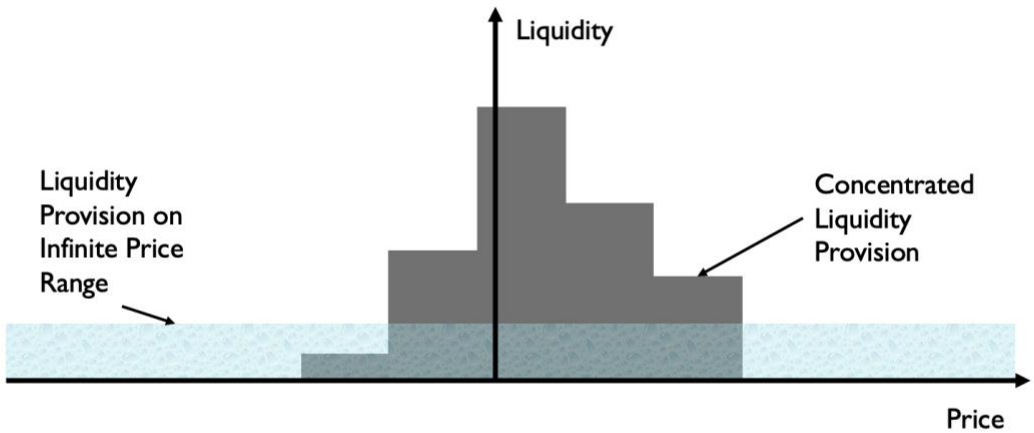

Source: Author

THE BLOCK · Research

Concentrated liquidity also allows for range orders (take profit and buy limit), while generating fees when filled. One can approximate a limit order by providing single asset liquidity within a specific range.[20] The downside is that a filled limit order can be reversed, if the price moves back before the position is closed.

Liquidity providers' gains and losses are affected by the interplay of several key factors, notably fee income from transactions carried out in their pool,[21] and transaction cost from re-balancing.[22] They are also subject to impermanent loss, which is more pronounced (relative to infinite price range liquidity provision) when providing concentrated liquidity in a more narrow price range: For the same amount of liquidity provided, a tighter price range translates into higher liquidity provision over that selected range, which then results in higher fee income (as long as the price remains within that range) but also higher risk of impermanent loss.

Figure 2.4 outlines the relation between concentrated liquidity and impermanent loss with a specific (stylized) example.[23] The horizontal axis displays the price range of WBTC in

---

[20] See "Uniswap: Range Orders".

[21] Fee generation itself is a function of the liquidity distribution over ticks, price movements, swap volumes, and the fee tier chosen.

[22] Rebalancing may be required when providing concentrated liquidity. For example, if one initially deposits 50% WBTC and 50% USDC in a concentrated liquidity pool but the price moves below the specified active price range, one is left with 100% WBTC and 0% USDC. Also see the example outlined on Figure 2.4. This requires rebalancing so that one can continue providing two-sided liquidity and receive fee income.

[23] For simplicity, everything is displayed linearly. Actual impermanent loss accrues in non-linear fashion. See Section 3 for an outline of non-linear hedges against impermanent loss.

USDC, ranging (discontinued) from 0 to infinity. The vertical axis displays the share of WBTC and USDC provided to the concentrated liquidity pool. Consider a liquidity provider, who enters the pool at a price of 30000 USDC per WBTC and sets the lower and upper bounds for liquidity provision to 29000 and 31000 USDC per WBTC, respectively, holding approximately 50% of value in each token. The red function displays the share of USDC over the selected price range in the concentrated liquidity provided. As the market price of WBTC increases, the share of WBTC in the liquidity provider's share of the pool decreases, until it consists only of USDC when the price equals 31000 USDC per WBTC or more. If the price moves below 29000, the liquidity provider's pool position consists entirely of WBTC – the relatively better performing asset gets arbitraged out of the pool. For comparison, the green dashed line displays a liquidity provider's share of USDC when choosing a price range from zero to infinity (i.e. similar to Uniswap V2-type AMMs). As the price of WBTC increases, the decline of WBTC's share in the liquidity provided is less pronounced, hence the effect of impermanent loss is smaller. In sum, the 'draining' of the better performing asset from the pool reserves results in impermanent loss (also see the examples provided in the Appendix) and becomes more pronounced if the range across which liquidity is provided is more narrow (for a given amount of capital).

Figure 2.4: Effect of Price Changes on Underlying Token Share

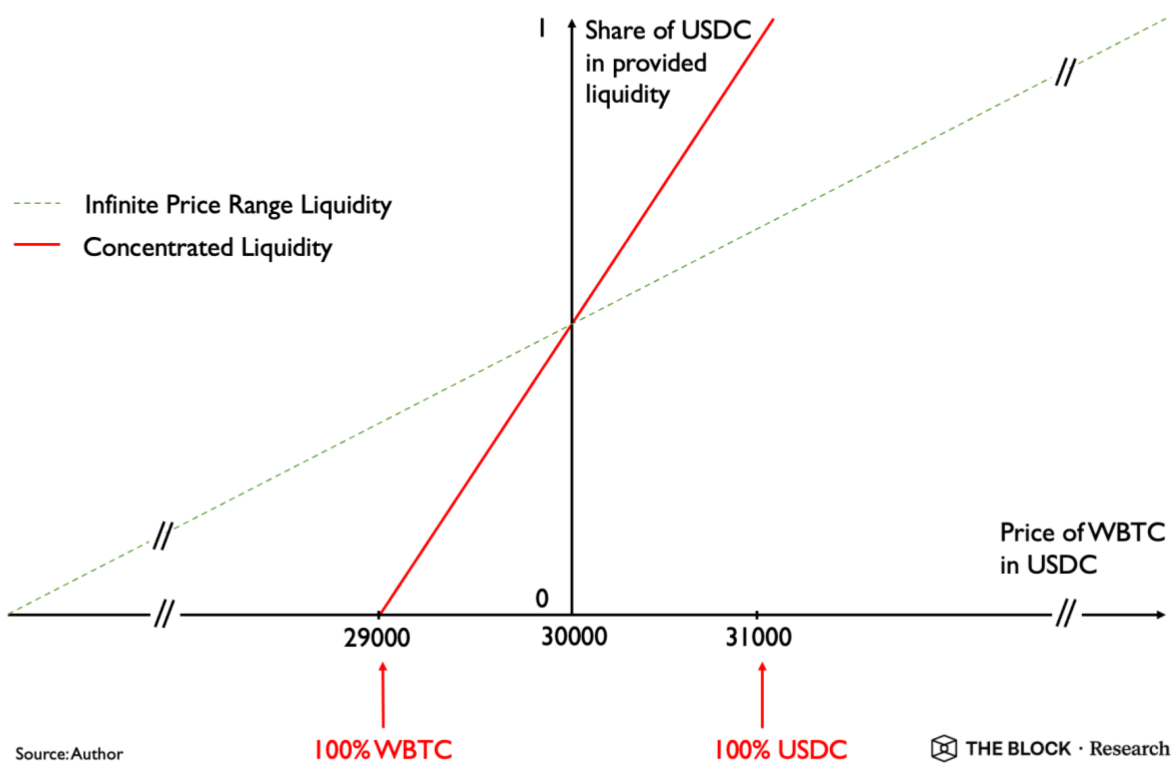

Curve v2 pools are conceptually similar to a concentrated liquidity provision, which is automatically managed based on internally set parameter values. Case Study 2.1 provides an outline of concentrated liquidity and management in Curve v2.

Case Study 2.1: Curve v2*

Curve v2 is conceptually similar to AMMs with concentrated liquidity. However, while having more degrees of freedom, which are internally set for individual liquidity pools, it algorithmically manages the liquidity position. That is, liquidity providers do not have to rebalance, or decide on a price range and liquidity distribution themselves.

Curve v2 extends its framework from Curve v1, using the hybrid invariant (see Table 2.1), to tokens, which have no stable relation to each other. This is achieved by introducing 'denomination' and 'dynamic pegging'. For a specific price point, tokens with different prices (e.g. USDC, WBTC, and ETH) are internally redenominated to have a similar price, by splitting up the 'more expensive' token into smaller units. This allows exchanging tokens based on a hybrid curve similar to Curve v1 for a given price.

If the pricing relation between the pool tokens changes, the pool's smart contract re-denominates the tokens, to have similar prices again. This however effectively rebalances the tokens and crystallizes impermanent loss. To mitigate impermanent loss, such re-denomination only takes place if the trading fees accrued since the last re-denomination exceed the crystallized impermanent loss at that new price relation. Note that this may result in the pools becoming unbalanced for some period of time. Trading fees are algorithmically increased when such pool imbalances occur.

In gist, Curve v2 provides automatic liquidity management for quasi-concentrated liquidity in its pools. However, it requires careful tuning of the parameters (by the protocol) to yield optimal results for liquidity providers and market participants carrying out transactions. Section 3 contains some insights on the performance of Curve v2 and automatic concentrated liquidity managers so far.

[*] Also see "The Block Research: Curve v2: concentrated liquidity with dynamic pegging".

The Landscape of Concentrated Liquidity AMMs

While concentrated liquidity has been firmly established with Uniswap v3 starting in 2021, there are only a few protocols, which went beyond announcing implementation. Table 2.1 displays AMM protocols,[24] ordered by TVL, which have implemented concentrated liquidity.

[24] A number of AMMs have announced to provide concentrated liquidity features in the future, however not implemented yet. These include Sushiswap, Shell Protocol, Hydraswap, Genius Yield, Acta Finance, Maladex, and Crocswap. Curve v2, while conceptually similar to concentrated liquidity, is not included. When adding all Curve v2 tricrypto pools across chains on which it is active, it ranks second with a TVL of $808mm (as of 1 June 2022).

Table 2.1: AMMs with Concentrated Liquidity

DEX	Blockchain	TVL (CL)* in MM USD
UniswapV3	Arbitrum, Ethereum, Polygon, Optimism	4200
Orca	Solana	18
Crema	Solana	11
Lifinity	Solana	5
Sheep Dex	BSC	3
Cykura	Solana	1
Algebra Finance	Polygon	1
Invariant	Solana	0

* Figures as of 1 June 2022, TVL for concentrated liquidity only

Source: Collected from the mentioned individual protocol pages THE BLOCK · Research

The three top exchanges for concentrated liquidity by TVL are Uniswap v3, Orca, and Crema. Uniswap v3, which is by far the largest in terms of TVL ($4.2 bn), operates on Ethereum, and other layer 2 chains (Arbitrum, Optimism) as well as a sidechain (Polygon). Given its firm foothold in the EVM (compatible) landscape and first mover advantage, it continues to dwarf all other concentrated liquidity AMMs.

Orca is second in terms of TVL ($18 mm). It went live with concentrated liquidity on Solana Mainnet for all users in April 2022. Its concentrated liquidity framework allows liquidity providers to choose a preset range for liquidity provision of stable pairs, for example, a narrow or a wider range. More experienced users can set this range however also completely free. In addition, Orca also integrated liquidity farming into its framework, which allows for subsidizing liquidity providers' yields with its own and other protocols' tokens. For example, at the time of writing the Whirlpool for SOL and stSOL (Lido's liquid staked SOL) is subsidized with both Orca and Lido tokens. Crema Finance is third in terms of TVL ($11 mm), went live on Solana Mainnet on 1 January 2022, and also allows for subsidizing yields with its own token.

Note that the cost (minting and rebalancing) for providing concentrated liquidity on different chains can be considerably different and therefore impacts profitability of liquidity provision. Table 2.2 compares the cost to initiate or re-balance a concentrated liquidity pool on Uniswap v3 on Ethereum and Polygon.

Table 2.2

DEX	Blockchain	Median Cost to Mint a Position	Median Cost to Change a Position (Native)
	Ethereum	$ 29.017	$ 16.112
	Polygon	$ 0.013	$ 0.007

Numbers are for Uniswap v3 based on 1 June 2022 for the median gwei cost of the last 30 days (for a position in ETH/USDC) and a price of 1998 USD/ETH or 0.66 USD/MATIC.

THE BLOCK · Research

While the cost on Ethereum can be prohibitively high for small positions, blockchains with cheaper transaction cost, such as Polygon and Solana, are much more amenable for active liquidity management, which requires rebalancing. However, further elements, such as network security or decentralization, may also be considered by liquidity providers, in particular if the value of the liquidity position is high.

Summing up, concentrated liquidity is a relatively new but versatile concept, which has quickly gained traction. It can be effectively used for stable pairs as well as those, which feature a wider price range. While it is more capital efficient than infinite range liquidity provision, it features a higher risk of impermanent loss. Concentrated liquidity providers have considerably more degrees of freedom to choose their strategies compared to infinite range liquidity provision, and they compete within the same liquidity pool. Uniswap v3 is by far the biggest AMM for concentrated liquidity, however, several AMMs (such as Orca on Solana) have joined the race and many other AMMs (such as Sushiswap) have announced plans to introduce concentrated liquidity features in the future. The next section looks at concentrated liquidity provision strategies and empirical results.

Section 3: Concentrated Liquidity Provision and Empirical Evidence

This section outlines strategies to provide concentrated liquidity and gives an overview on automatic concentrated liquidity management. It then sheds light on the empirical evidence for concentrated liquidity provision.

Strategies for Concentrated Liquidity Provision

As shown in Section 2, compared to infinite range liquidity provision, concentrated liquidity increases the risk-reward trade-off between fee income and impermanent loss. One can distinguish between simple approaches and more sophisticated approaches. **Simple approaches** consist of passive or semi-passive strategies. Such strategies include–

i. –providing liquidity across multiple price ranges for pairs, which can be expected to revert in price towards each other. For such token pairs, for example mSOL and SOL, it may make sense to provide the bulk of liquidity in a tight price range around unity (or a ratio one expects reversion to), and a lower amount of liquidity below and above that main range to still generate fee income for the case that the price temporarily moves out of the specified tight range.

ii. –providing liquidity across a wider price range, to achieve a lower-risk trade-off between impermanent loss, rebalancing[25] and capital efficiency (fee income). The trade-off between capital efficiency and rebalancing cost is affected by the pool fee selected. All else equal, lower fees normally result in a higher volume, which may or may not make up for less fee income per unit of liquidity provided. Many liquidity providers choose to provide liquidity among highly correlated tokens such as WBTC, ETH and SOL. Using a relatively wide price range prevents frequently realizing impermanent loss, and incurring transaction fees from rebalancing when operating in a very tight price range.

iii. –'buy low, sell high'-type strategies. For example, providing liquidity for the USDC-WBTC pair over a wide range from 20000 to 75000, say, results in gradually accumulating more WBTC while prices are decreasing, and slowly cashing out of it while prices are increasing, generating fee income in the process. Depending on the situation and the investor's expectation, the position may be exited at any point (e.g. when the lower bound is reached because the investor expects that WBTC has reached its floor).

[25] Rebalancing may consist of adjusting the active price range, changing the proportion of assets in the pool, and/or adding range orders.

iv. –providing liquidity in a single narrow range, which results in high capital efficiency but high risk of impermanent loss and frequent rebalancing if the price moves out of range. Such strategies, when supposed to be passive, make most sense for established stablecoin pairs such as USDT and USDC because their price usually remains aligned closely around unity.

v. –providing single-sided liquidity, which is very similar to a limit order in traditional finance. For example, one may deposit a single token in a price range above or below the current price. Once the price moves into the specified range, this results in selling the specified token for the desired one while earning swap fees. This approach combines a limit order with elements of liquidity provision.

Given the multitude of factors and strategic approaches affecting concentrated liquidity provision, several automatic liquidity managers have emerged. Active liquidity management protocols, which usually follow more **sophisticated approaches** (while charging fees), automatically rebalance their pools under management and re-invest proceeds.

Table 3.1: Overview of Automatic Concentrated Liquidity Managers

Liquidity Manager	Strategy Outline
	Aloe Capital manages 5 pairs on Uniswap v3 and aims to outperform passive liquidity provision on AMMs such as Uniswap v2. Its strategy consist of entering a concentrated liquidity position on Uniswap v3 that mimicks fee and impermanent loss accrual as if liquidity were provided for the pair on an (infinite range) constant product AMM. The leverage that comes with concentrated liquidity allows to employ unused capital to be deployed on other dapps such as compound or yearn to generate additional yield. If the market price moves out of range, that capital is used to rebalance the position.
arrakis	Arrakis Finance (recently rebranded from Sorbet Finance) provides automatic liquidity via Gelato's G-Uni pools, which auto-compound via re-investing fees back into the pool. It currently features 47 managed pairs for Uniswap v3. Liquidity is provided for an efficient range around the current price, with bots automatically rebalancing the position if the price moves out of the current bounds. The methodology framework is undisclosed.
	Charm Finance features liquidity management for 3 Uniswap v3 pairs. The strategy consists of always maintaining two orders. First, a base order taken symmetrically around the current market price, with the range determined by the pair's volatility/market situation. Second, a limit order above/below the current price (for the token, which is currently held less of) to increase the probability that the position maintains a balanced token ratio in case the limit order gets filled.
	Gamma Finance currently manages 5 pairs on Uniswap v3 and follows a similar strategy framework as Charm Finance with a base and a limit order. It also automatically manages price range, rebalancing and re-investment of fee proceeds.

Source: Collected from the mentioned individual protocol pages

THE BLOCK · Research

Strategies are mostly based on backtesting and machine learning. For example, Gamma Strategies uses a statistical model to set center and width of the liquidity range based on recent market conditions.[26] Automatic liquidity managers have the advantage that their unit cost of rebalancing/reinvesting decreases as their assets under management increase. Table 3.1 outlines the main automatic (rebalancing) concentrated liquidity managers, which are currently live.[27]

Broadly speaking, Curve v2 also features elements of automatic concentrated liquidity management. In particular, it automatically chooses the optimal liquidity distribution for a given pool based on a set of parameters. However, it is not included in Table 3.1 because it is rather a DEX than an automatic liquidity manager for concentrated liquidity. For more details on the algorithmic management strategy in Curve v2, see case study 2.1.

Since impermanent loss is a major risk for concentrated liquidity providers, they may try to hedge this exposure. Three main approaches exist. First, taking the USD as unit of account, for pairs that include a stablecoin, such as USDC, one can use a long position in the volatile token to mitigate impermanent loss. Note that, given the non-linear nature of impermanent loss accrual, this is only an imperfect hedge.[28] Second, to take into account non-linearities, one may also add options to use straddles to cover impermanent loss risk.[29] This however incurs the fix cost of the option premium, which needs to be offset by fee income. Third, for the case of ETH as the volatile token, instead of options to cover the non-linear element of impermanent loss, one can use squeeth, which is a perpetual contract indexed to a power function of ether (power of two). Similar to using options, squeeth comes with a cost in the form of a funding rate, which can be considerable.[30]

Empirical Evidence

Most empirical investigations for concentrated liquidity are based on Uniswap v3 because it has been live for longest and is also largest in terms of TVL.[31] These analyses can be broadly categorized into those, which carry out assessments based on i) backtesting specific strategies using past market data or ii) counterfactuals. The latter compare results from market participants' investment choices with alternatives such as simply holding tokens ('HODLING') instead of providing them as liquidity on a decentralized exchange.

Neuder et al.[32] investigate the optimal solution across three investigated strategies, which consist of

---

[26] See "ConsenSys: Gamma Strategies".
[27] Inactive or announced projects, which have not yet started liquidity management with automatic rebalancing for concentrated liquidity, such as Izumi Finance, Lixir Finance, Mellow Protocol, Popsicle Finance and Steer Finance are not covered. Also see "The Block Research: Charm Finance and automated Uniswap v3 LP strategies".
[28] See, for example "Altonomy: Uniswap V3".
[29] See, for eample, "Kim: Uniswap V3 + Options Hedge Explainer".
[30] The funding rate can reach several 100% (annualised). See "Opyn: Squeeth".
[31] A notable exception is Au, who investigates the historical performance of Curve v2's tricryto2 pool, which features elements that are conceptually similar to concentrated liquidity and consists of USDT, WBTC, and ETH, between July and November 2021. Interestingly, while it outperforms a hypothetical constant mean market maker with 0% fee, it underperforms a simple "buy and hold" strategy.
[32] See "Neuder, Rao, Moroz, and Parkes: Strategic Liquidity Provision in Uniswap v3".

Liquidity provision–

1. –on a fixed price range,

2. –that is uniform over a price range, but resets to a new range if the price moves out of the current range, and

3. –that is distributed over a price range, with the liquidity provided in sub-ranges of that price range weighted by the probability that the price will move into that sub-range. It is reset to a new distribution if the price moves into the current liquidity distribution's tail.

Using a **backtesting** approach with ether data ranging from March 2018 to April 2020 they find that, depending on the liquidity providers' degree of risk aversion, **strategies, which automatically reset the price range, can yield up to 200x the income that would be obtained with infinite range liquidity provision** (e.g. on Uniswap v2). It is important to note that the authors did not take into account the effect of impermanent loss for these results, however.

When taking into account impermanent loss for a set of possible strategies including i) HODLING (no liquidity provision, but simply buy and hold), ii) infinite range, iii) fixed, and iv) active rebalancing, Fritsch,[33] using data from January to August 2021, finds that liquidity providers can achieve **considerably higher returns when using concentrated liquidity** (relative to the other approaches) across a tight range **for stablecoin pairs**. However, in his analyses, active liquidity providers' fee income tends to be lower than used to be the case for infinite price range liquidity provision in pre-Uniswap v3 times. This is due to much of the liquidity moving from DEXes with infinite price range liquidity provision at a relatively high fixed transaction fee (e.g. 0.3% at Uniswap v2) to Uniswap v3 and then, due to competition between liquidity providers, being provided in lower-fee pools. While, similar to Neuder et al, he finds fee generation is significantly higher **for volatile pairs**, there seems to be **no strategy that consistently outperforms infinite range liquidity when also taking into account impermanent loss**. Overall, according to the author, **the main beneficiaries of concentrated liquidity seem to be traders** profiting from much lower slippage. This finding is corroborated by Liao and Robinson,[34] who find that "Uniswap Protocol v3 has deeper liquidity in ETH/USD, ETH/BTC and other ETH pairs than leading centralized exchanges." This shows that AMM market structure "can surpass order-book exchanges and transform traditional financial market structure to be more liquid, stable, and secure".

Loesch et al[35] carry out **counterfactual** analyses. Based on liquidity pools covering 43% of the total value locked on Uniswap v3 (**excluding stable pairs**) they find that over the period May 5th 2021 to September 20th 2021, in **aggregate liquidity providers would have been better off by $60.8m had they HODLED** (counterfactual) instead of

[33] See "Fritsch: Concentrated Liquidity in Automated Market Makers".
[34] See "Liao and Robinson: The Dominance of Uniswap v3 Liquidity".
[35] See "Loesch, Hindman, Welch, Richardson: Impermanent Loss in Uniswap V3".

carrying out liquidity provision. Except for flash liquidity providers, which provide intra-block liquidity, no liquidity provision strategy was consistently (that is, statistically significant) profitable relative to HODLING.

Building on Loesch et al., fbifemboy carries out a four-part posting, which investigates profitability in aggregate but also breaks it down to a per-wallet basis to analyse profitability in the face of impermanent loss and identify successful strategies. The author broadly confirms the findings of Loesch et al. **Providing liquidity on Uniswap for pairs that incur impermanent loss results in worse outcomes than simply holding**.

Overall, available empirical evidence points towards an uneven accrual of the benefits of higher capital efficiency from concentrated liquidity, with swappers achieving much more favorable conditions due to less slippage, and liquidity providers, which would be better off 'HODLING'. The major drivers behind unfavorable results for concentrated liquidity providers on Uniswap v3 may be driven by i) high transaction costs, which may ii) prevent re-balancing/moving a position back into an active price range, iii) complexity of active management, and iv) intra-block 'just-in-time' liquidity provision.[36] These issues may be improved upon by operating on chains with lower transaction costs (for rebalancing), liquidity providers taking more informed decisions (for example based on tailored analytical tools such as the Uniswap v3 fee calculator from Flipside Crypto), or vaults that carry out active liquidity management on behalf of the users. Note that so far, the above results also seem to apply to automatic concentrated liquidity managers.[37]

In summary, concentrated liquidity allows liquidity providers to choose from a broad spectrum of strategies as well as automatic liquidity management approaches. However, the empirical evidence available so far indicates that concentrated liquidity providers are struggling to be profitable.

[36] Just-in-time liquidity provision is a sophisticated strategy executed by automated bots. It allows liquidity providers to add and remove liquidity within one block, diluting fee income of other liquidity providers in the pool.

[37] While automatic liquidity management protocols' advertised fee-based APYs usually look very enticing, they do not include absolute returns including impermanent loss in those. Chances are that such total returns would be displayed if they were positive.

Conclusion and Outlook

Concentrated liquidity has quickly gained traction as indicated by its strongly increasing market share among DEXes. While it is more capital efficient, the flipside is an increasing risk of impermanent loss for liquidity providers. In the good old days of traditional infinite range constant function AMMs, providing liquidity was easy and everybody could (passively) hop on the gravy train. In the brave new world of concentrated liquidity AMMs, which can approximate CLOBs,[38] risks and rewards are higher, and empirical evidence so far shows that more sophisticated (automatic?) liquidity management may be required to not get one's feathers plucked.

Liquidity providers are given many more degrees of freedom when providing concentrated liquidity. This leads to a number of different strategy setups, some of which may require sophisticated active liquidity management. For example, to efficiently generate liquidity for new stablecoin pairs or setting up exotic stablecoin pairs, which are not covered by Curve or other efficient stableswap protocols, automatic liquidity management may not be necessary given there is little need for rebalancing.

However, for more volatile pairs, which require rebalancing to be profitable, sophisticated (automatic) liquidity management (protocols) may have advantages. Liquidity providers are well-advised to study their market maker brethren's liquidity management strategies from traditional finance. So far the main advantage from higher capital efficiency seems to accrue to swappers, who get much less slippage at lower transaction cost.[39] It remains to be seen whether liquidity providers' total return (including protocol fees, impermanent loss, and 'compounded' contract risk) offers sufficient value relative to HODLING or infinite price range liquidity provision.

Going forward, this situation is unlikely to persist. The mediocre returns found for liquidity providers so far will likely improve over time, because it does not make economic sense for them to provide liquidity under those circumstances. Such an adjustment may take place either (i) because more successful liquidity management strategies are employed, or (ii) liquidity providers reduce liquidity (or choose higher fee tiers) until they receive at least a 'crypto risk free rate[40] plus a risk premium, which compensates them for the risk taken.

[38] See "Chitra, Angeris, Evans: How Liveness Separates CFMMs and Order Books".
[39] See "Liao and Robinson: The Dominance of Uniswap v3 Liquidity".
[40] See, for example, "Serenity Funds: DeFi Risk Free Rate". The article discusses use of the ether staking yield and USD-based yields obtained from low-risk Aave, Compound and Curve pools as a proxy for a DeFi risk-free rate.

Appendix: Impermanent Loss. Does it Only Exist on Constant Product AMMs?

Narrowing down impermanent loss to exclusively exist on constant product AMM may be problematic. It is to some extent understandable, because the term was introduced when the landscape of AMMs consisted overwhelmingly of those with constant product invariants, which, due to their design, did not require active liquidity management. However, the narrow definition may be an issue because economic losses similar to impermanent loss can arise outside the scope of constant product AMMs, including those with central limit order books (CLOBs).

To fix ideas, consider two examples, with the first for impermanent loss on a constant product AMM:[41] An infinite range constant product AMM pool (x*y=k) for the pair ETH to USD is initialized with the following values: x[USD]=200, y[ETH]=2, hence k=400.[42]

The current pool price, which may be different from the market price, is p=x/y=100 USD/ETH. An increase in the market price of ETH (for example at central exchanges), results in arbitrageurs exchanging USD for ETH to realize a profit. If the pool price rises to p=110 USD/ETH, then, abstracting from fees and transaction cost, an arbitrageur removes ~0.093 ETH[43] from the pool in exchange for ~9.76$.[44] In other words, the arbitrageur obtained ETH at a price of ~104.9 USD/ETH, while the market price (and pool price after his transaction) equals 110 USD/ETH.

At the new price, the market value of the liquidity provider equals ~419.51 USD,[45] while the value of his initial portfolio (200 USD and 2 ETH) would be worth 420 USD had he followed a buy and hold strategy instead of providing liquidity. Therefore, the liquidity provider has an impermanent loss of ~0.49 USD. This loss is only realized, if the liquidity provider exits his position, it may revert if the price moves back to the initial value. Hence this type of loss is termed 'impermanent'.

Liquidity providers on central exchanges, which use CLOBs, are usually more sophisticated than those on decentralized exchanges. To distinguish them, in the following they are referred to by the term 'market makers'. Market makers on CLOBs follow complex strategies consisting of simultaneously posting bids and asks around the market price, with the spread being determined by their expectation of market volatility and their inventory policy. In gist, market makers aim to buy a token slightly below the market price in the hope that they can sell it to another market participant for their posted asking price, before the

[41] See "Topaze.Blue: Impermanent Loss in Uniswap V3" for a similar example. For simplicity transaction costs and fees are not included in the calculations.
[42] Note that k remains fixed at 400 for the following calculations.
[43] Obtained by taking the difference between the original quantity of ETH in the pool, and the new quantity at the new price, y=(400/p_new)^0.5.
[44] Calculated as the difference between the original quantity of USD in the pool, and the new quantity at the new price, x=p_new*y_new.
[45] Calculated as 1.9069*110+209.76.

market price moves against them. However, if the market price moves consistently against them, they face inventory risk.[46] In a situation, where market participants are selling only, the market maker is buying into a falling market, or, in other words, the market maker is buying the underperforming asset.

To fix ideas consider the following example of a market maker in a strongly directional market situation. Table A.1 displays the market maker's liquidity offered in the direction of price movement.

Table A.1: Example for Liquidity Provision

Liquidity provided at the following ticks	Liquidity offered (ETH for USD)	USD obtained
101	0.2	20.2
102	0.2	20.4
103	0.2	20.6
…	…	…
109	0.2	21.8
110	0.2	22
Sum	**2**	**211**

THE BLOCK · Research

The current market price is 100 USD/ETH, and the market maker has posted 0.2 ETH at each of the indicated ticks (101 to 110). After the FED suddenly announces that it plans to 'look through the current bout of inflation', market participants immediately start selling USD for ETH all the way up to a price of 110 USD/ETH. She sells all ETH for 211 USD. Let's apply the same before-after calculation to the market maker's portfolio (assuming she also posted bids worth 200 USD below a price of 100 USD/ETH, to match the AMM example above):

At the current market price of 110 USD/ETH, the USD value of the market maker's inventory is 411 USD (200 USD+ 211 USD). Had the market maker just held to the initial inventory without providing liquidity, its value would be 420 USD (200 USD + 110 USD/ETH * 2 ETH). The difference between these two positions, which is termed impermanent loss in the crypto space, is 9 USD.

The (arguably highly stylized) example offers two insights: First, comparing the extent of impermanent loss (or temporary 'economic loss' due to inventory risk for the case of the CLOB) between the examples, it is considerably larger for the case in which the market

[46] For an outline see "Hummingbot: Inventory Risk".

maker has 'concentrated' the liquidity in a tight range around the market price (2 ETH liquidity were provided up to 110 USD/ETH). In traditional constant product AMMs, the liquidity is uniformly distributed between a price range of zero to infinity, which, compared to liquidity provision on a tighter range on CLOBs (and, similarly, on a concentrated liquidity AMM), is less capital efficient but also less risky in terms of impermanent loss/temporary economic losses.[47]

Second, and more important, the answer to the above question "Does Impermanent Loss Only Exist on Constant Product AMMs?", is 'no' from a theorist perspective and 'well, somewhat' from a practitioner's perspective. No, because the example proves that by the definition of impermanent loss, market makers on CLOB exchanges can face risks that are economically similar to those called impermanent loss on decentralized exchanges. The fact that these risks may be referred to by a different term on other markets (e.g. inventory risk) does not make them disappear. However, and to be fair by taking a practitioner's view, the example assumed a market maker, who follows a very passive ('dumb') liquidity provision approach. In reality, market makers update their bid and ask prices, spreads and liquidity on a tick-by-tick basis and do not provide liquidity uniformly across ticks, thereby employing strategies to mitigate their inventory risk. As a corollary to this it is important that liquidity providers are aware that the risk-reward ratio is higher when providing concentrated liquidity. This indicates a need for sophisticated strategies to be profitable.

[47] It can be shown that i) concentrated liquidity AMMs allow approximating CLOB, and ii) that providing the same amount of liquidity on a tighter price range results in more capital efficiency at the risk of higher impermanent loss. See "Uniswap: FAQ".

Disclosures

This report is commissioned by Orca. The content of this report contains views and opinions expressed by The Block's analysts which are solely their own opinions, and do not necessarily reflect the opinions of The Block or the organization that commissioned the report.

The Block's analysts may have taken positions in the assets discussed in this report and this statement is to disclose any perceived conflict of interest. Please refer to The Block's Financial Disclosures *page for author holdings. This report is for informational purposes only and should not be relied upon as a basis for investment decisions, nor is it offered or intended to be used as legal, tax, investment, financial or other advice. You should conduct your own research and consult independent counsel on the matters discussed within this report. Past performance of any asset is not indicative of future results.*

© 2022 The Block Crypto, Inc. All Rights Reserved.