

When Uniswap v3 returns more fees for passive LPs

Austin Adams

Gordon Liao*

May 2022

1 Introduction

Understanding the returns of liquidity provision on the Uniswap Protocol is important for efficient capital allocation. While different strategies of active liquidity provision exist on Uniswap v3, non-rebalancing liquidity positions remain an important source of liquidity. This study quantifies the historical fee accruals on non-actively managed liquidity positions on Uniswap v2 and v3. We show that the non-rebalancing returns on Uniswap v3 are greater than that of v2 for the majority of token pairs that exist on both v2 and v3. Specifically, we define non-rebalancing returns as those coming from full-range v3 positions and concentrated positions on range-bound stablecoin pairs that do not require position modifications by liquidity providers. On average, non-rebalancing v3 positions outperform v2 positions by around 54% in fee returns.

- For token pairs with 100-bps fee-tier in v3, full-range v3 positions outperform v2 positions by around 80% in fee returns.
- For 1-bp stablecoin-pairs, range-bound v3 positions outperform v2 positions by around 160% in fee returns.
- For 30-bps fee-tier, full-range v3 outperforms v2 by around 16% in fee returns.

*We thank Hayden Adams, Mary-Catherine Lader, Matteo Leibowitz, Will Pote, Dan Robinson, and Devin Walsh for helpful comments and suggestions.

- The 5-bps fee-tier is the only fee-tier that non-rebalancing v3 positions underperform v2 (by around 68% in fee returns on average).

Uniswap v3 has already surpassed prior versions of constant product function automated market makers in liquidity and volume, and it has deeper market depth in common ETH pairs than centralized exchanges (see here). The results from this study suggest that liquidity providers should deploy on v3 regardless of how active they are in managing their liquidity pool positions.

2 New pool creations


Users continue to deploy a greater number of new pools on Uniswap v2 relative to v3. Figure 1 shows that 7,660 new pools were deployed on v2 while only 859 new pools were deployed on v3 in 2022 Q1.

Liquidity pool creators often cite several key reasons for why they chose to deploy on Uniswap v2 instead of v3. First, providing liquidity on Uniswap v2 is perceived as easier than on v3, as rebalancing is not necessary in v2. Second, some liquidity providers believe that v3 returns without rebalancing are lower as there are more active liquidity providers that might compete with inactive liquidity. Third, deploying pools on v3 can have higher gas costs¹. In this study, we address the first two main concerns and show that v3 is in fact better for most liquidity providers.

¹Additional reasons why some users prefer v2 over v3 are:

- Fees on v2 are automatically compounded into non-rebalancing liquidity unlike in v3, therefore more convenient for long-term liquidity providers
- v2 supports fee-on-transfer tokens that have been phased out in v3.

Figure 1: Number of new liquidity pool created by Uniswap version



Note: The figure shows the monthly number of pools created on Uniswap by version. The sample period is from November 2020 to April 2022.

3 Methodology

We calculate fee returns on marginal liquidity positions in v2 and v3 using pool-level historical data. The marginal liquidity assumes that the position added is small relative to the overall pool size, and there is no effect on the behavior of other participants in the pool. In the fee calculation for v2, we subtract the implied capital returns derived from the pool's price changes from the total returns to obtain the fee returns. For v3, we calculate the value of the fees gained from global storage variables. In non-pegged token pairs, we calculate the returns on the full-range positions. These full-range v3 positions share the same divergence loss (also known as "impermanent loss") as v2 positions but differ in fee returns due to potential differences in the fee-tier and volume traded in the pools. In pegged token pairs,

e.g. stablecoin-to-stablecoin pairs, we identify a price range that contains the historical price deviations for each pair and calculate returns for the defined range².

We perform the calculations for all v2 and v3 pools that share the same underlying pairs. As there can be multiple pools (fee-tiers) for each pair of tokens in v3, we limit our sample to v3 pools with the highest average total value locked (TVL) for each token pair in the comparison. Fees generated are expressed as a percentage of the capital deployed. The fees returned are daily and non-compounding. For a more detailed explanation of the methodology, see the methodology appendix and [code](#). We also provide our data for this study [here](#).

4 Pairwise fee return differences

We compare pairwise return differences between non-rebalancing v3 and v2 positions and make several observations³. First, returns in v3 liquidity positions on average outperform v2 across all fee-tiers except in the 5-bps fee-tier. Second, heterogeneity in return differences can largely be explained by differences in fee-tiers (v3 has four fee-tiers ranging from 1 bp to 100 bps while v2 fees are fixed at 30 bps) and concentrated liquidity ranges for pegged pairs. Higher v3 fee-tiers generally have proportionally higher full-range fee returns. Stablecoin pairs with concentrated liquidity positions also have higher returns in v3 relative to v2. Third, the launch of the swap router that allows routing across v2 and v3 has had a significant impact on equalizing the relative fee returns between the two versions of the protocol.

Table 1 shows the summary statistics of the return differentials. The mean pairwise difference in daily returns across all fee-tiers is 4.4 bps, indicating that v3 pools outperform the corresponding v2 pools in annualized returns by around 16% (4.4 bps x 365 days) on average. For token pairs in the 100-bps fee-tier on v3, the average and median pairwise fee return differences are around 10.2 bps (36.5% annualized) and 4.9 bps (17.9% annualized), respectively. This indicates that users with non-rebalancing positions on v3 benefit from

²Defining a tick range for pegged token pairs allows us to calculate returns using only global storage variables. We exclude pegged pairs that have deviated from the peg in this calculation. This also assumes minimal divergence from the peg. The range used captures 99% of the historical price of the pegged token pairs.

³For example, for the AAVE-WETH 30 bps pool, the non-rebalancing v3 position has a mean daily return of 2.8 bps, while the AAVE-WETH v2 pool has a mean daily return of 1.8 bps. The pairwise difference would be 1.0 bps.

the higher fee-tiers and volume in Uniswap v3 despite not actively managing concentrated liquidity positions.

Table 1: Pairwise fee return difference (non-rebalancing v3 return - v2 return)

	Number of token pairs	Mean daily fee returns in bps		v3 - v2 daily returns in bps			
		v3	v2	Mean	25 th percentile	50 th percentile	75 th percentile
all fee-tiers	270	12.6	8.2	4.4	-1.1	0.2	3.9
1 bp	4	0.8	0.3	0.6	0.2	0.3	0.6
5 bps	18	0.6	1.9	-1.3	-1.3	0	0
30 bps	141	6.7	5.8	0.9	-1.7	-0.3	0.5
100 bps	107	22.9	12.7	10.2	1.7	4.9	13.1

Note: Fee return difference is calculated as the mean non-rebalancing v3 return minus the mean v2 return for each pool in the pair. For each token pair, the v3 pool with the highest average TVL is selected as a comparison against v2 when multiple v3 pools exist on a token pair with multiple fee-tiers. The sample period is from September 16th, 2021 (launch date of swap router v1) to March 28th, 2022.

Returns on stablecoin pairs in range-bound positions are also more compelling on v3 relative to v2. In the 1-bp fee-tier, range-bound v3 positions (containing 99% of the historical price movement) have returns that are 2.6x of v2 positions on the same underlying pairs. For a detailed summary of stablecoin returns, see Table A2 and A3.

The distribution of return differences is wide within each fee-tier, as indicated by the percentile ranges in Table 1 and Appendix Figure 1. These within fee-tier ranges can be explained by differences in volume and the concentration of liquidity in each pool.

5 Which v3 pools outperformed v2 in non-rebalancing returns?

Given the heterogeneity observed in return differences, we explore which pools have relative over- and under- performance of v3 versus v2 in returns. Figure 2 visualizes the fee returns for different token pairs by scatter plotting v3 non-rebalancing position returns against v2 returns. Each dot in the scatter plot indicates a single token pair, with the y-axis reflecting

the v3 fee return and x-axis indicating the v2 fee return. The dots above the diagonal line have v3 outperforming v2 in non-rebalancing return, and the dots below the diagonal have v2 outperforming v3.


Three main observations are apparent. First, the pairs that have v3 outperforming are typically either in the 100-bps fee-tier or 1-bp fee-tier on stablecoin pairs. 100-bps token pairs are usually less liquid, more recently deployed, and more volatile. For the 1-bp fee tier, token pair prices are less volatile, but Uniswap v3 has substantially more volume than v2. The concentrated liquidity on the 1-bp pool enables high returns that surpass those of v2. Second, there is a strong positive correlation between v2 and v3 fee returns with many pairs falling close to the diagonal. This indicates that returns between v2 and v3 are somewhat kept in equilibrium. Third, a number of pairs have close to zero v2 returns but high v3 returns (dots falling close to the y-axis). These are mainly pools with relatively small v2 TVL and volume⁴.

We label some representative pairs (top 8 highest TVL pairs) for pools that have v3 outperforming v2 and vice versa in Figure 2 panel B. Most of the large-cap ETH-stablecoin pairs tend to have higher returns on v2 relative to inactive v3 positions. This is likely a reflection of the higher concentrated liquidity and professional market making in those pools. In contrast, long-tail token pairs perform well in v3 full-range liquidity positions. Stablecoin pairs also perform well in v3 returns as rebalancing is not needed for a narrow range position.

⁴The appendix provides detailed methodology on how we handle small pools.


Figure 2: Fee returns on v2 vs. non-rebalancing v3 positions

Panel A: Full sample



Note: This figure shows the mean of daily returns for each pair of pools. The sample is from September 16th, 2021 to March 28th, 2022. The color of the point indicates the fee-tier of the Uniswap v3 pool. 3 pairs are not shown for clarity.

Panel B: Partial sample of pools with average daily fee returns less than 5 bps




Note: This figure shows the mean of daily returns for each pair of pools. The sample is from September 16th, 2021 to March 28th, 2022. The color of the point indicates the fee-tier of the Uniswap v3 pool. Only pairs with average daily fee return less than 5 bps are shown for clarity.

6 Time series of returns

Figure 3 shows the comparison of non-rebalancing v3 versus v2 fee returns over time starting from when v3 was launched in May 2021. In the early months of v3 launch, fee returns on v3 were generally lower than v2 as liquidity and volume were building up. As market participants integrated v3, fee returns on v3 increased relative to v2 over the summer of 2021.

The launch of Auto Router v1, which routed swaps strictly within each version of Uniswap protocol had little initial effect on the relative return. The second version of the router, launched on December 16, 2021, reduced the returns of v3 versus v2 as it routed some volume back to v2.

Figure 3: Median daily pairwise returns on v2 vs. non-rebalancing v3 positions



Note: The figure shows the 7-day rolling average of median daily returns across pools in bps. The sample period is from May 13th, 2021 to March 28th, 2022. The color of the line represents the fee-tiers of v3 pools.

7 Future Research and Conclusion

Our study provides a few practical takeaways. First, long-tail token-pair liquidity providers should consider using v3 as fee returns are generally higher on the 100-bps fee-tier without the additional need for rebalancing in full-range positions. Second, pool creators should take into account the higher average fee collection and capital efficiency of v3 when deciding on the protocol version for pool deployment. Third, liquidity provision on v3 for stablecoin and pegged pairs can generate competitive returns in concentrated liquidity positions absent large depegging events.


In this study, we focused on liquidity positions at two ends of the liquidity concentration spectrum - full-range positions and ultra-concentrated positions on stablecoin pairs. Future research is needed to recover fee returns on wide-range positions in v3 that historically did not require rebalancing, e.g. a position with a range of 500 to 10,000 on ETH/USDC. These wide-range positions might provide a middle ground to obtain concentrated liquidity and higher fee collections without the need for frequent active management. Furthermore, to quantify total returns for liquidity positions, the fee returns derived here need to be jointly considered in the context of divergence loss on volatile token pairs as well as depegging probability and loss-given-depeg for stablecoins. We aim to address these questions in future research.

This study shows that Uniswap v3 outperforms v2 in full-range and concentrated stablecoin pair returns on average. As Uniswap v3 has both an advantage for returns and liquidity, we hope these results will encourage projects and users to migrate from Uniswap v2 to v3.

8 Appendix


8.1 Appendix Figures

Figure A1: Distribution of pairwise returns



Note: The figure shows the density of mean pairwise differences between non-rebalancing v3 positions and v2 positions. The bandwidth of the density plot has been slightly increased to slightly smooth the chart. The sample is from September 16th, 2021 to March 28th, 2022. The color of the density line represents the fee-tier of the v3 pool.

Figure A2: Stablecoin/pegged assets return by pair



Note: This figure shows the average daily fee return for each token pair on Uniswap v2 and v3. For each token pair, the v3 pool with the highest average TVL is selected as a comparison against v2 when multiple v3 pools exist on a token pair with multiple fee-tiers. The sample period is from September 16th, 2021 (launch date of swap router v1) to March 28th, 2022.

Table A1: Pairwise fee return difference: full-range v3 return - v2 return

	Number of token pairs	Mean daily fee returns in bps		v3 (full-range) - v2 daily returns in bps			
		v3 (full-range)	v2	Mean	25 th percentile	50 th percentile	75 th percentile
all fee-tiers	270	12.6	8.2	4.4	-1.2	0.1	3.9
1 bp	4	0.01<	0.3	-0.3	-0.4	-0.1	-0.1
5 bps	18	0.2	1.9	-1.7	-1.5	0.0	0.0
30 bps	141	6.7	5.8	0.9	-1.7	-0.3	0.5
100 bps	107	22.9	12.7	10.2	1.7	4.9	13.1

Note: Fee return difference is calculated as the mean full-range v3 return minus the mean v2 return for each pool in the pair. For each token pair, the v3 pool with the highest average TVL is selected as a comparison against v2 when multiple v3 pools exist on a token pair with multiple fee-tiers. The sample period is from September 16th, 2021 (launch date of swap router v1) to March 28th, 2022.

Table A2: Pairwise fee return difference: stablecoin/pegged v3 return - v2 return

	Number of token pairs	Mean daily fee returns in bps		v3 (range-bound return) - v2 daily returns in bps			
		v3 (range-bound)	v2	Mean	25 th percentile	50 th percentile	75 th percentile
all fee-tiers	9	1.1	1.2	0	-0.7	0.2	1.2
1 bp	4	0.8	0.3	0.6	0.2	0.3	0.6
5 bps	5	1.4	1.9	-0.5	-2.5	-0.7	1.2

Note: Fee return difference is calculated as the mean range-bound v3 return minus the mean v2 return for each pool in the pair. For each token pair, the v3 pool with the highest average TVL is selected as a comparison against v2 when multiple v3 pools exist on a token pair with multiple fee-tiers. The sample period is from September 16th, 2021 (launch date of swap router v1) to March 28th, 2022.

Table A3: Range position returns on stablecoin pairs

Pool Information			Mean daily fee returns in bps			Position range (99% of historical price distribution)	
Token 0 - Token 1	Fee-tier (in bps)	Average TVL (\$mm)	v3 (stable)	v2	v3 - v2	Lower Price	Upper Price
DAI-USDC	1	184.7	1.8	0.2	1.6	0.9996	1.0003
USDC-USDT	1	182.8	1.1	0.8	0.2	0.9989	1.0003
FRAX-USDC	5	177.7	0.4	2.8	-2.5	0.9944	1.0034
FEI-USDC	1	50.7	0.1	0.1	-0.6	0.9993	1.0004
BUSD-USDC	5	26.9	1.7	4.6	-2.9	0.9978	1.0016
DAI-USDT	5	18.5	2.8	0	2.3	0.9981	1.0021
DAI-FEI	5	18.3	0.6	0.5	0.6	0.9943	1.0073
DAI-PAX	1	10	0	0	0	0.9995	1.0007
BUSD-USDT	5	0.8	0.8	1.5	-1.2	0.9975	1.0025

Note: Fee return difference is calculated as the mean v3 return minus the mean v2 return for each pool in the pair. For each token pair, the v3 pool with the highest average TVL is selected as a comparison against v2 when multiple v3 pools exist on a token pair with multiple fee-tiers. The sample period is from September 16th, 2021 (launch date of swap router v1) to March 28th, 2022.

8.2 Fee returns methodology

8.2.1 v2 fee returns

Let x_t , y_t be the amount of asset 1 and asset 2 reserves in a constant product function market making (CPFMM) liquidity pool at time t . Let p_t be the marginal price of token y in units of token x , $p_t = \frac{y_t}{x_t}$. Let k_t denote the liquidity that exists in the pool at time t . The constant product formula, $x_t y_t = k_t$ is held in every period. That is, $x_t = k_t / y_t = \sqrt{k_t / p_t}$ and $y_t = \sqrt{k_t p_t}$. k_t can only be changed via accrual of return fees absent of addition and subtraction of reserves through liquidity pool mints and burns.

Let $v_t \equiv v(k_t, p_t)$ be the portfolio value of the liquidity pool expressed in asset 1 as the numeraire at time t . When expressed in terms of the amount of reserve asset 1 and 2, the portfolio value is $v(k_t, p_t) = x_t + y_t p_t^{-1}$. Applying the CPFMM formula, the liquidity pool

value can be expressed as

$$v_t \equiv v(k_t, p_t) = \sqrt{\frac{k_t}{p_t}} + \sqrt{k_t p_t} p_t^{-1}.$$

Absent of mints and burns, the simple return on a liquidity pool due to fee accrual is

$$r_{t+1}^{fee} = \frac{v_{t+1} - v(k_t, p_{t+1})}{v_t}.$$

Note that $\frac{v(k_t, p_{t+1})}{v_t}$ is the gross return of the liquidity pool portfolio value without fee accrual. Minting and burning of liquidity pool positions can be accounted for as follows. The evolution of liquidity value follows $k_{t+1} = k_t + \kappa_{t+1} + \phi_{t+1}$, where κ_{t+1} is the net mints minus burns that occurs in between t and $t+1$ and ϕ_{t+1} is the fee accrual due to transactions in this period, i.e. $\phi_{t+1} = \sum_i \lambda |s_i| \forall$ swap s_i that occurs between t and $t+1$ and fee tier λ . Let $k'_{t+1} \equiv k_{t+1} - \kappa_{t+1}$. Thus, we have the mint-burn adjusted simple fee return on a liquidity pool

$$r_{t+1}^{fee} = \frac{v(k'_{t+1}, p_{t+1}) - v(k_t, p_{t+1})}{v(k_t, p_t)}.$$

8.2.2 v3 returns

Uniswap v3 added several pool, tick, and positional indexed state values. The following methodology follows the descriptions in the v3 white paper.

Position- and tick- indexed state

Each position has three state values, two associated with upper and lower tick and a liquidity value l . The liquidity value denotes the virtual liquidity held by the position. The two ticks are the lower tick - i_l - and the upper tick - i_u . These ticks both have two values associated with them - feeGrowthOutside0X128 - $f_{o,0}$ - and feeGrowthOutside1X128 - $f_{o,1}$.

The ticks are determined by the user when they create a position. The tick values $f_{o,0}$ and $f_{o,1}$ track how many fees were accumulated within a certain range.

Pool-indexed state

Pools also track feeGrowthGlobal0X128, $f_{g,0}$, and feeGrowthGlobal1X128 - $f_{g,1}$. These two values track the total amount of fees collected per unit of virtual liquidity l .

General fee calculations

Unclaimed fees are calculate as follows. We drop the subscript indicating token 1 and 2 since the formula for the two are the same.

Cumulative fees per share $f_r(t)$ in the range between two ticks i_l and i_u is

$$f_r(t) = f_{g,t} - f_{b,t}(i_l) - f_{a,t}(i_u)$$

where i_c is the current tick state of the pool, $f_a(i)$ and $f_b(i)$ are defined as

$$f_a(i) = \begin{cases} f_g - f_o(i) & i_c \geq i \\ f_o(i) & i_c < i \end{cases}$$

$$f_b(i) = \begin{cases} f_o(i) & i_c \geq i \\ f_g - f_o(i) & i_c < i \end{cases}$$

The uncollected fees can be defined as

$$f_u(t_1, t_0) = l(f_r(t_1) - f_r(t_0))$$

where t_1 is the current time and t_0 is the time the position was opened. $f_r(t_0)$ can either be calculated from the pool state or from feeGrowthInside value stored in the position state.

8.2.3 Full-range fee Calculations

For full-range positions, i_c is always between i_l and i_u , i.e. $i_l \leq i_c \leq i_u \forall i_c$

The fee calculation simplifies to

$$f_r(t) = f_{g,t} - f_{o,t}(i_l) - f_{o,t}(i_u)$$

However, $f_{o,t}(i_l)$ and $f_{o,t}(i_u)$ are static values, i.e., $f_{o,t}(i_l) = 0$ and $f_{o,t}(i_u) = c$. The constant c is the value of $f_r(t)$ when tick i_u is initialized and drops out. Thus uncollected full range fee is

$$f_u(t_1, t_0) = l(f_{g,t_1} - f_{g,t_0}).$$

$f_{g,t}$ is the fee growth of one unit of liquidity l since time 0 and $f_{g,t_1} - f_{g,t_0}$ gives the fee growth of one unit of liquidity from t_0 to t_1 . The fee return on a full-range $v3$ positions is

$$r_{t+1}^{fee} = l \frac{f_{g,t_1+1} - f_{g,t}}{v_t}$$

where v_t is the portfolio value of the full-range liquidity position at time t defined similarly as that in $v2$.⁵

8.3 Data Methodology

We use data from Uniswap $v2$ and $v3$ subgraphs. The data sample is constructed as follows. First, we pull the top 500 pools by volume from Uniswap $v3$ and remove any pool that has no Uniswap $v2$ counterpart in the sample. Since multiple $v3$ pools may use the same $v2$ pool as Uniswap $v3$ implemented multiple fee-tiers, we choose the $v3$ pool with the highest average TVL over the sample period if there are multiple $v3$ pools (choosing the pool with the highest volume yields a similar sample).

To calculate return on pegged token pairs, we empirically assess the tick range that bounds minor price fluctuations in the pegged asset pairs. Specifically, we winsorize historical peg token prices at the 0.5% threshold symmetrically (capturing 99% of the historical prices). With the observed historical price range, we create synthetic positions on the pegged pairs. This methodology takes a conservative approach in creating liquidity positions that do not have the need to rebalance.

To avoid extreme outliers as a result of small pool size on any particular date, we impute the returns with 0 if the total value locked for a given pool on an observation date is less than \$1000. We chose this data cleaning procedure for two reasons. First, large returns for small TVL pools could not be realized. As we assume that additional liquidity deployed into the pool is marginal. For pools with TVL less than \$1000, absolute returns are generally small (e.g. fraction of a cent) despite non-zero returns when expressed as a percentage. Additionally, for small TVL pools, large capital deployment would impact the behavior of other participants in the pool and thus contradict our marginal liquidity assumption.

⁵Equation 2.2 in the $v3$ white paper has the real reserve of a position defined by the curve $\left(x + \frac{L}{\sqrt{p_b}}\right)(y + L\sqrt{p_a}) = L^2$. In full-range positions, $p_a \rightarrow 0$ and $p_b \rightarrow \infty$, this indifference curve becomes identical to $xy = L^2 \equiv k$.