

Fee APR ranks liquidity pools backwards: realized loss-versus-rebalancing and passive returns on Uniswap v3

Maksym Kovalenko

Kremenchuk Mykhailo Ostrohradskyi National University, Kremenchuk, Ukraine

ORCID: 0000-0001-6497-8217

Corresponding author: maksym@mpost.io

Abstract

Passive liquidity providers on concentrated-liquidity automated market makers are ranked almost everywhere by fee APR, the ratio of fees earned to pool value. This paper shows that the metric orders pools close to the reverse of their realized net profitability. Using per-minute swap flows valued against a block-synchronised centralised-exchange reference price, I decompose passive returns into fees, realized loss-versus-rebalancing (LVR) and gas for 62 Uniswap v3 pools on Ethereum, Arbitrum and Base over 37 weeks of 2026, with 2,232 pool-weeks entering the ranking tests. Sorting pools weekly on headline fee APR and holding for one week returns -9.39% APR, against -5.72% for holding the whole panel and -0.49% for sorting on realized LVR-net APR. The sign replicates on Arbitrum, where the top-minus-bottom spread is -4.67 pp ($t = -3.66$), and on Base, where it is -7.90 pp ($t = -4.81$); the pooled rank-space coefficient is -0.306 ($t = -4.49$, clustered by pool). Volatility is not the channel: adding realized volatility of the external pair leaves the fee-APR coefficient intact and volatility itself insignificant. I then test whether the spatial structure of liquidity improves on aggregate LVR, which was the hypothesis this study set out to confirm, and find that it does not. Two further measurements follow. Just-in-time liquidity captured 0.089% of gross fees on Ethereum, and gas is 2.13% of LVR on Ethereum against 0.14% and 0.10% on the two rollups.

Keywords: automated market makers; concentrated liquidity; loss-versus-rebalancing; adverse selection; decentralised exchange; performance measurement

JEL classification: G12; G14; G23; D47; C58

1 Introduction

Milionis et al. (2022) established loss-versus-rebalancing as the measure of adverse selection borne by an automated market maker's liquidity providers: the difference between the value of the pool's inventory and that of a continuously rebalanced portfolio holding the same exposure at an external reference price. The result is theoretical, and its empirical counterpart is demanding, because it needs an external price synchronised to the block clock together with swap-level flow data. Practical allocation therefore proceeds on fee APR, annualised fees divided by total value locked. That figure is gross of adverse selection by construction, and it is what essentially every analytics front-end displays.

That fee APR is incomplete is not in dispute. What has not been established is the sign and the size of the error it induces in a pool-selection decision. If fee APR were only noisy, an allocator following it would do no worse than one allocating at random. This paper shows that the allocator does considerably worse than that.

The arithmetic behind this is simple. Fee APR is, to a first approximation, the fee tier times the ratio of volume to pool value, times the provider share of fees. A fee-APR sort therefore selects pools where a large amount of flow runs against a small amount of capital. That is also the condition under which realized adverse selection per dollar of capital is high, and over this sample the second effect dominates the first. Section 5.3 shows the accounting behind it: conditional on realized LVR, the sign of fee APR turns positive, which is what remains of the metric once the omitted term is restored. The decomposition of flow into arbitrage and uninformed trading is not measured here, so this is an interpretation the results are consistent with rather than a mechanism this paper tests.

The paper makes four contributions.

First, a validated measurement pipeline. The reference price is not a detail. The two feeds most readily available for this task produce LVR estimates of the wrong sign, and per-minute netting is the coarsest aggregation that preserves the quantity. Both biases are quantified against a one-second ground truth.

Second, the main result. Fee APR predicts next-week realized passive-LP net APR negatively, out of sample, with inference that accounts for overlapping windows and cross-sectional

correlation, and not through a volatility channel. The magnitude is given as a portfolio return rather than as a rank correlation.

Third, independent replication on two further chains under an identical pipeline, with the predictions registered before extraction.

Fourth, a null result on spatial structure, which was the motivating hypothesis of this work and is reported as the failure it was.

Throughout, established method is separated from modelling choices, and both from claims that could not be validated, with each marked where it appears.

2 Related literature

Four strands bear on this paper.

The first is the theory of adverse selection in automated market making. Milionis et al. (2022) define LVR and show that it is the correct accounting of what an arbitrageur takes from a liquidity provider; Milionis et al. (2023) extend the argument to the case with fees. Hasbrouck et al. (2026) work in the opposite direction, showing that higher fees can raise decentralised-exchange volume because they endogenously reduce price impact. Neither line is an empirical statement about which pool an allocator should choose, and neither is tested here; they supply the quantity this paper measures.

The second is empirical work on realized liquidity-provider outcomes. Heimbach et al. (2022) model the choices a Uniswap v3 provider faces and document how widely realized performance varies across them. Urusov et al. (2026) reconstruct position-level profit and loss from on-chain events for WETH pools on Base across four protocols, build a taxonomy of fifteen position types, and report that roughly one provider in six avoids losses in that segment. Both papers ask who earned what, at the level of the individual position. This paper asks a different question, at the level of the pool: given the metrics an allocator can observe before committing capital, what does following the standard one cost? The two are complements, and the position-level literature is the natural place to take the present result next.

The third is the measurement of specific frictions. Wan and Adams (2022) document just-in-time liquidity on Uniswap. Section 6.2 re-measures it under a detector that captures same-block multi-

transaction episodes as well as single-transaction ones, and finds it close to extinct in these pools in 2026.

The fourth is criticism of the metrics the sector reports. Luo et al. (2025) show that total value locked is inflated by double counting through wrapping and leverage, and propose total value redeemable in its place. Their objection is to how the quantity is constructed. The objection here is to how a quantity performs as a ranking signal, and the two meet in Table 3, where TVL is negatively signed as a predictor of next-week net return. Alexander (2025) studies price discovery and efficiency across Uniswap pools, treating the pool as a price-formation mechanism; the concern here is the provider's return, which is the other side of the same flow.

What is missing from all four is an evaluation of the selection metric itself. The gap is narrow and it is the one an allocator faces every week.

3 Data and measurement

3.1 Sample construction

Pools are selected on December 2025 information only, on average TVL of at least \$250,000 and average daily volume of at least \$100,000 over that month. The panel then runs from 1 January to 14 September 2026. Selecting on a pre-window month is what makes the sample survivorship-complete: pools that later faded or stopped trading stay in, and are not dropped. Six Ethereum pools in the selected set were inactive for most of the panel window and are retained.

Two exclusions are applied. Both remove pools that would otherwise flatter the results.

Rebasing and fee-on-transfer tokens are excluded, because for these the change in a pool's token balance is not the sum of its swap and liquidity events, so fee accrual and inventory accounting both break. This exclusion is a modelling choice: it is applied by symbol against a maintained list, not by comparing transfer amounts to balance deltas, which would be the rigorous test. It is an acknowledged approximation.

Pools whose two legs map to the same external reference are excluded. This trap is worth naming because it is silent. For a wrapped-BTC pair such as WBTC/cbBTC the external reference ratio is identically 1, so measured LVR is zero by construction, and such pools rise to the top of any LVR-net quality ranking. Five Ethereum pools fall into this category (two tBTC/WBTC pools at

different fee tiers, WBTC/cbBTC, tBTC/cbBTC and WBTC/LBTC) together with two on Arbitrum. They appeared at the top of my own ranking before the defect was identified.

The Ethereum universe is therefore 36 pools before this second exclusion and 31 after it. Both bases appear below, because the inference in Section 5.2 is reported on the wider one and the headline results on the narrower one; the base is stated wherever a number is given. After exclusions the panel is 31 Ethereum, 21 Arbitrum and 10 Base pools. The window spans 257 days, or 37 weeks, and carries 3.48M Ethereum pool-minutes of net flow together with 3.2M on Arbitrum and 2.1M on Base. The ranking tests in Sect. 5 use the 36 complete weeks, which is 2,232 pool-weeks across the 62 pools; the boundary week is dropped there because a partial week has no valid next-week outcome. All reported results use this basis unless stated otherwise.

3.2 The reference price is the whole measurement

LVR is defined against an external reference price. Using the pool's own price is the classic error and yields a quantity with no economic content, since the pool price is exactly what the arbitrageur moves. The practical question is which external price, sampled when.

I establish a ground truth on one pool-week, USDC/WETH at the 0.05% tier over 10 to 16 August 2026, with 22,600 swaps and \$340M of volume, by valuing every swap at second resolution against Binance ETHUSDT. Each candidate construction is then measured against it.

reference construction	realized LVR (bp of volume)	bias vs ground truth
Binance 1-second, contemporaneous (ground truth)	5.174	-
Binance 1-minute close of the containing minute	5.218	+0.043
Binance 1-minute midpoint of the two adjacent closes	2.525	-2.649
Binance 1-minute close of the previous minute	-0.167	-5.341
Warehouse minute-price feed, +2 minutes	-1.701	-6.875
Warehouse minute-price feed, same minute	-3.672	-8.846

Two things follow. The warehouse minute-price feed, which here is Dune's prices.minute table sourced from CoinPaprika and is the convenient choice most on-chain analyses reach for, is not synchronised to the block clock. It biases LVR by -8.85 bp on a 5.17 bp signal, which reverses the sign. A markout-lag study confirms the mechanism: against that feed the LVR curve flips sign between +2 and +5 minutes and has not plateaued by +15, whereas against Binance at second resolution it plateaus flat from roughly -1 s to +60 s (Fig. 1a). The correct one-minute

convention is the close of the containing minute, accurate to 0.04 bp. The previous-minute close and the midpoint are both badly biased.

A separate check fixes the netting resolution. Aggregating flows into 2-minute buckets before valuation biases LVR by -2.4 bp on a 5.2 bp signal, and 5-minute buckets reverse the sign. Per-minute netting is a requirement rather than a convenience, and it is what the panel uses.

I also exclude any pool whose tokens are priced in the warehouse from decentralised-exchange trades, since such a reference is endogenous to the quantity being measured. This removed 18 of 64 Ethereum universe tokens.

That LVR needs an external reference is established in Milionis et al. (2022). What is contributed here is the measured size of the two biases and the demonstration that the convenient feed fails.

3.3 Fee accrual

The pool mechanics follow the Uniswap v3 core specification (Adams et al., 2021). Fees are accrued per swap, as the fee tier times the gross input leg, multiplied by the provider's share of active liquidity. They are never inferred from Collect event timestamps. Collection is discretionary and lazy, so Collect-based accrual misattributes fees across time and corrupts every downstream profit-and-loss figure without leaving a trace.

The Uniswap v3 protocol-fee switch is active for every pool in the panel and was enabled in waves inside the window, twice on Ethereum (December 2025 and March 2026) and twice on each rollup (March and July 2026). Several pools switched twice, with different denominators. The provider share of gross fees is therefore a stepwise function of time per pool, recovered from SetFeeProtocol events rather than held constant. Ignoring it overstates provider fee income by 17 to 25% for part of the window.

Accrued fees are reconciled against an independent identity. Because Collect returns burned principal plus owed fees, the quantity Collect minus Burn should approximate fees actually withdrawn. Across the Ethereum panel accrued provider fees are \$52.6M against \$47.0M observed, a ratio of 0.893. That figure is struck on all 36 Ethereum pools over the full 37 weeks, which is why it exceeds the \$50.9M in Table 2: the latter is the 31-pool observable basis restricted to the 36 complete weeks used in Sect. 5. The shortfall is expected from positions still open at the window boundary, and the per-pool deviation is two-sided around 1.03, which is the

signature of collection timing rather than of estimator bias. The residual is reported rather than absorbed.

3.4 Estimators

For each pool-minute I observe net flow in both tokens, (Δ_0, Δ_1) , and value it at the contemporaneous external reference (p_0, p_1) :

Gross fee = $f \times \Sigma(\text{gross input leg} \times \text{reference price})$, with f the fee tier.

LP fee = gross fee $\times m(t)$, with $m(t)$ the stepwise provider share after protocol fees.

Mark-to-market change = $\Delta_0 \cdot p_0 + \Delta_1 \cdot p_1$.

Realized LVR = gross fee – mark-to-market change.

Net passive-LP P&L = LP fee – realized LVR – gas.

Gas is the actual cost of mint, burn and collect transactions, accounted per chain. Arbitrum folds L1 posting cost into gas used, while on Base the L1 fee is a separate field and is added.

4 What passive liquidity providers earned




Fig. 1 Measurement validation and realized outcomes. **a** the reference-price validation of Sect. 3.2, showing realized LVR against reference-price lag for two Binance sampling conventions, with the warehouse minute feed marked at its wrong-signed value. **b** fee-tier economics on the 31 Ethereum pools with an observable external basis, with provider fees, realized LVR and gas by tier and net profit and loss overlaid. **c** the tick-relative source term of Sect. 6.1, showing liquidity added at the price and removed after the price has left. **d** out-of-sample rank correlation of each predictor with next-week realized passive-LP net profit and loss, computed on the 36-pool Ethereum panel rather than the 31-pool observable basis of Table 3, so the values differ from that table by the amount reported in Sect. 5.2; the final row is model M2, aggregate features plus the spatial block

Across the three chains, passive liquidity providers earned \$119.8M in fees against \$153.9M of realized loss-versus-rebalancing on \$154B of volume, for a net of -\$35.6M.

chain	pools	volume	LP fees	realized	gas	net	gas as %	weeks
-------	-------	--------	---------	----------	-----	-----	----------	-------

				LVR			of LVR	negative
Ethereum	31	\$80.9B	\$50.9M	\$68.5M	\$1,461k	-\$19.1M	2.13%	75.9%
Arbitrum	21	\$44.1B	\$20.7M	\$27.0M	\$37k	-\$6.4M	0.14%	83.3%
Base	10	\$29.5B	\$48.3M	\$58.4M	\$56k	-\$10.1M	0.10%	84.4%

Fees recover roughly three quarters of LVR and no more. On Ethereum, 6 of 31 pools were cumulatively profitable and 75.9% of pool-weeks were negative.

By fee tier on Ethereum (Fig. 1b), the 0.05% and 0.3% tiers lost \$19.69M between them. Both thin tiers were net positive, the 0.01% tier at +\$0.123M over 4 pools and the 1% tier at +\$0.027M over a single pool, so the low tier is not uniquely viable. The 1% figure rests on one pool, WLD/WETH at \$1.4M of TVL, and is therefore a one-observation cross-section rather than a tier result. Neither positive result is as solid as the negative ones, and the 0.01% tier's headline surplus falls from \$0.356M to \$0.123M once the structurally-zero-LVR pools of Sect. 3.1 are removed. The defensible statement is that passive providers were compensated only in the thin, low-volatility corner of the panel and lost money everywhere else.

The single largest destroyer of passive capital is USDC/WETH at the 0.05% tier: median weekly net APR of -7.3%, negative in 100% of weeks, on \$30.1B of volume.

Gas is not the explanation. It is 2.13% of realized LVR on Ethereum against 0.14% on Arbitrum and 0.10% on Base, two orders of magnitude smaller on comparable volume, and net losses are the same order of magnitude on all three chains. The rollup losses are adverse selection with the execution-cost channel effectively switched off.

5 Ranking pools

5.1 Out-of-sample design

I predict, for each pool and week, the next week's realized passive-LP net APR, using only information available at the time of the sort. The design is walk-forward with an expanding training window: models are fitted on weeks before j and evaluated on the cross-section of week j , for 18 test weeks. The evaluation statistic is the cross-sectional Spearman rank correlation within each test week, because ranking is the decision a provider actually makes.

predictor	mean ρ	naive t	Newey-West t	weeks with
-----------	-------------	---------	--------------	------------

				correct sign
(a) headline fee APR	-0.173	-3.62	-2.95	4 of 18
(b) TVL	-0.064	-1.67	-2.00	7 of 18
(c) volume / TVL	-0.068	-1.41	-1.14	8 of 18
(d) realized LVR-net APR	+0.580	+9.85	+10.12	18 of 18
M1 aggregate features (ridge)	+0.487	+11.26	+11.55	18 of 18
M2 aggregate features + spatial (ridge)	+0.484	+11.09	+11.44	18 of 18
M3 spatial only (ridge)	+0.361	+13.82	+18.17	18 of 18

Two features of this table matter. Fee APR, TVL and volume-to-TVL, the three metrics in general use, are all negatively signed. And no fitted model beats the single-feature baseline (d), so simple persistence of realized LVR-net APR is the strongest predictor tested.

5.2 Inference

The weekly ρ series has overlapping training windows and cross-sectionally correlated errors, so the naive t-statistic overstates the evidence. On the full 36-pool Ethereum panel the fee-APR series gives mean $\rho = -0.265$ with naive $t = -5.59$. The Newey-West(4) statistic (Newey & West, 1987) is -4.87, a stationary block bootstrap (Politis & Romano, 1994) with 20,000 resamples at block length 4 gives a 95% interval of $[-0.370, -0.174]$, and no resample changes the sign. On the 31-pool observable basis reported above the effect is weaker but retains its sign and its significance.

5.3 Volatility and fee-tier composition are not the channel

The natural alternative explanation is that fee APR sorts pools by volatility, and that LVR scales with σ^2 , in which case the finding would reduce to the established theory rather than say anything about the metric. I test this directly in rank space, with week fixed effects and standard errors clustered by pool (31 clusters), using realized volatility of the external pair price computed from the reference series.

specification	fee APR rank	external σ rank	realized LVR rank
A	-0.270 ($t = -2.72$)	-	-

B	-0.291 (t = -3.77)	+0.035 (t = +0.32)	-
C	+0.276 (t = +2.87)	+0.083 (t = +1.08)	-0.714 (t = -7.17)
D	-	-	+0.607 (t = +13.93), LVR-net rank alone

Adding volatility leaves the fee-APR coefficient intact and volatility itself insignificant, so the effect does not run through σ . Specification C is the accounting identity becoming visible: conditional on realized LVR, more fee income is of course better, and the coefficient turns positive.

The second composition story is that fee APR is a proxy for fee tier, and that the result only says high-tier pools do badly. It is not that either. Within the 0.05% tier the rank correlation is -0.576 (t = -12.5) and within the 0.3% tier it is -0.246 (t = -4.0), so the metric ranks pools backwards inside a single tier as well as across tiers. The one tier where fee APR ranks correctly is 0.01%, at +0.265 (t = +2.2), which holds the stable and closely correlated pairs where adverse selection is small; that is also the only corner of the panel where passive provision made money, as Sect. 4 reports.

What the results are consistent with instead can be read off the same table together with the ranking results. Fee APR is close to the fee tier times the volume-to-TVL ratio, and volume-to-TVL is itself negatively signed in Sect. 5.1, with a group spread of -7.89 pp in Sect. 5.5. The sort works as intended: it picks the pools generating the most fee income per dollar of capital. Those are also the pools where the most is being extracted per dollar of capital, and over this sample the second quantity is the larger one. External volatility does not capture this because it is a property of the price process rather than of how much flow a pool sees per dollar of depth at a given volatility. I state this as an interpretation, not a tested channel: separating arbitrage from uninformed flow at the swap level would be the test, and it is not performed here.

5.4 Ranks and levels

The result is a rank phenomenon, and it does not survive a move to levels. The same regressions with the variables in levels rather than ranks give a fee-APR coefficient of +0.090 (t = +0.74) alone and +0.150 (t = +1.24) with volatility, so there is no reliable linear partial effect.

Three considerations reconcile the two, and the first two are why ranks are reported as primary.

The decision under study is a ranking. A provider with a fixed budget chooses among pools; the object that determines the outcome is the ordering of the candidates, not the conditional mean of net APR given a one-unit change in fee APR. Evaluating a selection metric in rank space measures the thing the metric is used for.

Net APR is heavy-tailed in this panel, so a small number of pool-weeks with extreme values carry most of the weight in a least-squares fit, while a rank statistic gives them the same weight as any other observation. The portfolio result of Sect. 5.5 survives dropping the largest pool in the panel, at -4.87 pp ($t = -9.87$), which shows the ranking evidence does not rest on the one pool that dominates by size. It does not by itself establish that the levels null is caused by the tails, and no such test is offered.

A null in levels and a monotone negative relation in ranks are not in contradiction. Ordinary least squares estimates a conditional mean under a linear functional form. A relation that is monotone and strongly nonlinear can leave that coefficient indistinguishable from zero while the ordering is preserved.

The boundary this places on the claim should travel with the result. The evidence supports the statement that a fee-APR sort selects worse pools than no sort at all. It does not support a statement of the form "one percentage point of fee APR costs b percentage points of net APR", and no such coefficient is offered.

5.5 The cost in portfolio terms

Rank correlations are hard to price. I therefore sort pools each week on a signal known at that time, hold for one week, and compute the TVL-weighted realized net APR of each group. No model is fitted, so all 36 weeks are usable.

sorting signal	top tercile	bottom tercile	all pools	top – bottom	Newey-West t
headline fee APR	-9.39%	-2.92%	-5.72%	-6.46 pp	-3.19
volume / TVL	-9.64%	-1.75%	-5.72%	-7.89 pp	-3.21
external volatility	-8.97%	-2.36%	-5.72%	-6.61 pp	-3.29
realized LVR-net APR	-0.49%	-9.95%	-5.72%	+9.47 pp	+4.69

A provider allocating by headline fee APR earned -9.39% APR against -5.72% for holding the entire panel. The metric costs 3.66 percentage points of APR relative to not choosing at all, and 6.46 points relative to inverting it. Allocating on realized LVR-net APR earned -0.49%, within half a point of break-even.

The sign survives equal weighting (-2.90 pp, $t = -2.64$), dropping the largest pool in the panel (-4.87 pp, $t = -9.87$), and restriction to the 0.3% fee tier alone (-5.34 pp, $t = -7.49$).

5.6 Replication on Arbitrum and Base

The prediction registered before extraction was that the fee-APR group spread would be negative on both chains, that the LVR-net group spread would be positive, and that the LVR-net top group would sit closer to break-even than the panel. The pipeline is identical in every respect: same selection rule, same reference construction, same per-minute netting, same stepwise protocol-fee treatment, same window.




Fig. 2 Replication of the ranking result on Arbitrum and Base. **a** top-minus-bottom group spread in next-week realized net APR, by chain and sorting signal. **b** realized net APR of the selected portfolio against the hold-everything benchmark. **c** fees, realized LVR and gas by chain, with gas shown as a percentage of LVR

chain	pools	pool-weeks	split	fee-APR spread	t	LVR-net spread	t
Ethereum	31	1,116	terciles	-6.45 pp	-3.18	+9.47 pp	+4.69
Arbitrum	21	756	terciles	-4.67 pp	-3.66	+7.46 pp	+5.67
Base	10	360	median	-7.90 pp	-4.81	+9.78 pp	+6.84

All three predictions held on both chains. In rank-space regressions with week fixed effects and pool-clustered errors the fee-APR coefficient is -0.270 on Ethereum, -0.327 on Arbitrum and -0.387 on Base, pooling to -0.306 ($t = -4.49$) across 62 pools and 2,232 pool-weeks. The LVR-net coefficient is +0.607, +0.683 and +0.655, pooling to +0.641 ($t = +20.00$).

Two caveats on the replication, both of which cut against the result. Base contributes only 10 pools, so its cluster-robust t of -1.96 rests on 10 clusters and is unreliable at that count; the statistic to read there is its portfolio spread, computed over 36 weekly observations. More substantively, the thinness of Base is not incidental. Most of that chain's volume sits in tokens with no centralised-exchange listing, among them ZORA, AERO, BRETT, DEGEN, TOSHI and CLANKER, which cannot be LVR-measured at all under this method. The Base result therefore speaks to listed pairs only, and those are a minority of the chain's activity. This is a selection boundary on the entire method, not on Base alone.

6 Spatial structure: a null result

The motivating hypothesis of this work was that where liquidity sits relative to price, and how it moves, carries information about adverse selection that aggregate LVR does not. It does not, and that is reported here as the outcome.

6.1 The source term, and what it is not

A note on formulation, because the fluid analogy is commonly overstated. Liquidity does not advect between adjacent ticks. It is added and removed pointwise by mint and burn, so the honest statement is that the time derivative of liquidity density equals a pointwise source term:

$$\partial L(i,t) / \partial t = S(i,t)$$

with no flux divergence. I construct $S(i,t)$ in coordinates relative to the pool tick at the moment of each event, recovered by an as-of join to the last swap at or before the event block, at (block, log-index) precision and with complete coverage of 327,647 event-blocks. No $\nabla \cdot F$ term is written on the tick lattice, because on this lattice it would not mean what it means in a fluid.

The resulting map (Fig. 1c) shows the expected asymmetry cleanly. Mints land a median 3.5 tick-spacings from the pool price and are in range 74.4% of the time; burns land 9.6 spacings away

and are in range 61.8% of the time. Liquidity is supplied at the price and withdrawn after the price has left it, which is the range-exit signature and is visible in every pool.

The meaningful transport is on the position-identity graph rather than the tick lattice: a provider burns at range A and mints at range B. I construct these edges by successive burn-then-mint events of the same identity, taking the owner, or the transaction sender where the position manager is the nominal owner. Within 24 hours there are 109,856 such edges, of which 72,418 (65.9%) move the provider to a different range. Across those range-changing edges the median shift is 5.9 tick-spacings and 68.6% move toward the price. Migration is concentrated: the top five identities account for 93% of migrated liquidity, so this is a handful of automated managers rather than the provider population.

6.2 Just-in-time liquidity is nearly extinct

I detect just-in-time episodes as mint and burn of the same (pool, range, owner) key within a single block, fully closed, bracketing a swap. Two findings follow.

The dominant pattern is same-block multi-transaction bundling rather than the single-transaction sandwich: 31,628 fully-closed same-block episodes against only 626 same-transaction ones. Any detector keyed on single-transaction mint-and-burn misses roughly 98% of episodes, which may explain divergence from earlier estimates.

The magnitude is also small. When a just-in-time position is present it holds a median 80.8% of active liquidity at the targeted swap, but episodes are rare and target small swaps. Total capture is \$59k of \$65.7M of gross fees, or 0.089%, and episode counts fall month by month from 5,088 in January to 2,600 in August 2026. On this evidence just-in-time liquidity is not a material drain on passive providers in these pools in 2026. Whether the activity migrated to Uniswap v4 hook-enabled pools was not tested, and that is the obvious next question.

The same-transaction channel admits a second check that does not rely on fee attribution at all, and I ran it on USDC/WETH at the 0.05% tier over 1 August to 20 September 2026 with an independent pipeline: an integer-faithful reimplementation of the Uniswap v3 swap mathematics, validated against 314,125 realized swaps, of which 97.26% reproduce the on-chain output to the wei at a notional-weighted mean absolute residual of $1.05e-08$ basis points. Within that window 442 transactions mint and burn inside a single transaction, but only 26 of them also swap this

pool. Re-simulating those swaps with the same-transaction mint removed gives the benefit the provider actually extracted: at most 0.0018 basis points on any single order, 0.00056 on a notional-weighted basis, across \$106k of affected volume. The just-in-time position never exceeds 0.25% of active liquidity. Two methods with nothing in common beyond the pool therefore agree that the single-transaction sandwich is immaterial here. This check speaks only to that channel, which Sect. 6.2 shows is the smaller of the two, and not to the same-block bundled pattern.

6.3 The incremental test

I form a twelve-feature spatial block, comprising mint and burn mass concentration within 2 and 5 tick-spacings of price, mint-minus-burn asymmetry, in-range fractions at mint and burn, median range width, churn, migration rate, migration directionality, just-in-time fee share and distinct provider identities, and ask whether it improves on the aggregate-LVR baseline out of sample.

model	mean ρ	Newey-West t	Δ vs baseline
(d) realized LVR-net APR alone	+0.580	+10.12	-
(d) + spatial block, ridge	+0.526	+10.10	-0.055
(d) + spatial block, gradient boosting	+0.549	+11.72	-0.032
(d) + aggregate features + spatial, ridge	+0.516	+12.96	-0.064
spatial block alone	+0.361	+18.17	-0.219

Adding spatial features degrades the baseline. The paired weekly difference is -0.051 with $t = -2.88$ on the 36-pool panel, where the baseline is +0.616 and the spatial-augmented model +0.565; on the 31-pool observable basis reported in the table the difference of means is -0.055 and the ordering is unchanged. Gradient boosting narrows the gap without closing it. The spatial block carries real signal on its own at +0.361, but that signal is already contained in realized LVR.

The boundary of this null should be stated plainly. It establishes that these spatial features add nothing. The one spatial quantity most likely to matter, tick-local realized LVR concentration, never entered the out-of-sample test, because the compressed panel encoding required to fit the extraction budget carries net flow but not the pool tick, leaving that measure available for only one month of the window. I therefore cannot claim that spatial information is uninformative, only that this feature set is. Closing that gap is a re-extraction rather than a research problem.

7 Range placement and a falsified assumption

A standard analytic treatment writes the net return of a position of relative half-width x as $\text{net}(x) = A/x - (\sigma^2/x^2) \cdot g/C$, with the first term the fee-minus-LVR margin scaling as concentration, the second the amortised cost of rebalancing at gas g on capital C under a driftless first-exit frequency σ^2/x^2 , and an interior optimum at $x^* = 2g\sigma^2/(A \cdot C)$. I derived that policy with rebalancing cost as an explicit input and measured gas, then tested its leading assumption, that fee-minus-LVR margin per unit liquidity is invariant to range width, against the position-level panel.

The assumption is false, and it fails in the direction that reverses the conclusion. On USDC/WETH at the 0.05% tier over 10 to 16 August 2026, with 5,458 identified positions, LVR per unit of fee rises with range width.

range width (tick-spacings)	positions	LVR per unit fee
under 2	82	1.088
2-5	48	1.231
5-10	49	1.327
10-25	109	1.155
25-50	159	1.366
50-100	341	1.677
100-250	1,618	1.465
over 250	3,052	1.452

The Spearman correlation across width bins is +0.833. Concentrating a range therefore does not simply multiply a fixed margin, and the analytic optimum should not be acted on. It is reported as conditional and marked unvalidated.

What survives is simpler and more useful. Where realized margin is negative, which is the case for 25 of the 36 Ethereum pools, no width is profitable, so the decision facing a provider is pool selection rather than range placement.

8 Limitations

The method can only see pools whose reference price exists off-chain on a centralised venue. This excludes long-tail and chain-native assets entirely, and on Base it excludes the majority of activity. Results generalise to listed pairs, not to concentrated liquidity as a whole.

The spatial null is incomplete in the way set out in Sect. 6.3.

Position identity is partial. 83.1% of liquidity events are owned by the position manager, so pool-level (range, owner) keys aggregate multiple token IDs. I matched pool events to IncreaseLiquidity and DecreaseLiquidity by (transaction, liquidity) to resolve 244 token IDs in the focal week; 5,205 ownership transfers were extracted but not applied, because token-ID coverage is too thin to build an ownership timeline. Panel-level identity uses the transaction sender.

Legacy liquidity is not distributed. Positions minted before the window supply a median 34% of active liquidity and 49% of fees in the focal pool-week. I use the observed active liquidity from Swap events as the attribution denominator and report this as a residual.

Cross-pool and cross-chain attribution is bounded rather than clean. Same-pair fee tiers have a median daily net-flow Spearman of +0.865 and disagree in sign on 9.7% of days, so they share common directional flow. That is a daily-resolution observation and not a test of block-level cross-tier or cross-chain arbitrage, which would require simultaneous per-block flow across all tiers and was not extracted. Single-pool LVR attribution should be read as an upper bound.

Uniswap v4 and hook-enabled pools are out of scope and are treated as non-comparable by construction.

Statistical power is limited: 18 out-of-sample test weeks, and Base contributes 10 pools.

Assumption diagnostics beyond those reported were not assessed. The window covers 2026 only, so the result is not tested across a change in volatility regime, and extension of the panel to earlier years is the first robustness check a reader should want.

9 Conclusion

Passive liquidity provision in the pools studied here was unprofitable over 2026 to date, at \$119.8M of fees against \$153.9M of realized loss-versus-rebalancing across three chains. That much is a measurement, and it confirms theory.

The result that is not a confirmation concerns the metric. Fee APR does not merely omit adverse selection. It ranks pools in approximately the reverse of their realized profitability, and an allocator following it does measurably worse than one who does not choose at all. The correction is available and cheap once the reference price is right: persistence of realized LVR-net APR is

the strongest predictor tested, on every chain, and it brings a passive allocator to within half a point of break-even in a set of pools that lost 5.7% on average.

The spatial hypothesis that motivated this work failed. Where liquidity sits relative to price, how it migrates, and how much of it is just-in-time are all measurable, all behave as theory would suggest, and none of them improve on the aggregate number once you have it.

Statements and declarations

Funding. No funding was received for conducting this study.

Competing interests. The author declares no competing financial or non-financial interests related to this work.

Data availability. All primary data are public: Uniswap v3 on-chain events read through a public blockchain data warehouse, and Binance spot trade history for the external reference price. The derived panels supporting every table and figure are deposited at [TODO Zenodo DOI] under a CC BY 4.0 licence. The deposit holds the pool universe with its selection variables, pool-week and pool-day decompositions of fees, realized LVR and gas for all 62 pools, the reference-price validation series, the tick-relative liquidity-event panel, the just-in-time episode set, the migration edge set, the position-level panel for the focal pool-week, the estimation outputs behind each reported table, the swap engine and the per-order counterfactual behind the check in Sect. 6.2, and a file-to-table map. The code is available from the author on request.

Use of generative AI. Generative AI tools were used for code generation for the extraction and estimation pipeline. The author designed the study, specified and verified every analysis, checked all reported figures against the underlying data, and takes full responsibility for the content.

References

- Adams, H., Zinsmeister, N., Salem, M., Keefer, R., & Robinson, D. (2021). *Uniswap v3 core*. Uniswap Labs.
- Alexander, C. (2025). Price discovery and efficiency in Uniswap liquidity pools. *Journal of Futures Markets*. <https://doi.org/10.1002/fut.22593>

Hasbrouck, J., Rivera, T. J., & Saleh, F. (2026). The need for fees at a DEX: How increases in fees can increase DEX trading volume. *Management Science*.

<https://doi.org/10.1287/mnsc.2023.00726>

Heimbach, L., Schertenleib, E., & Wattenhofer, R. (2022). Risks and returns of Uniswap V3 liquidity providers. In *Proceedings of the 4th ACM Conference on Advances in Financial Technologies*. ACM. <https://doi.org/10.1145/3558535.3559772>

Milionis, J., Moallemi, C. C., Roughgarden, T., & Zhang, A. L. (2022). *Automated market making and loss-versus-rebalancing*. arXiv. <https://arxiv.org/abs/2208.06046>

Luo, Y., Feng, Y., Xu, J., & Tasca, P. (2025). *Piercing the veil of TVL: DeFi reappraised*. arXiv. <https://arxiv.org/abs/2404.11745>

Milionis, J., Moallemi, C. C., & Roughgarden, T. (2023). *Automated market making and arbitrage profits in the presence of fees*. arXiv. <https://arxiv.org/abs/2305.14604>

Newey, W. K., & West, K. D. (1987). A simple, positive semi-definite, heteroskedasticity and autocorrelation consistent covariance matrix. *Econometrica*, 55(3), 703-708.

<https://doi.org/10.2307/1913610>

Politis, D. N., & Romano, J. P. (1994). The stationary bootstrap. *Journal of the American Statistical Association*, 89(428), 1303-1313. <https://doi.org/10.1080/01621459.1994.10476870>

Urusov, A., Berezovskiy, R., Krestenko, A., & Kornilov, A. (2026). *Liquidity provision in CLMMs: Evidence from transactions data*. arXiv. <https://arxiv.org/abs/2604.22069>

Wan, X., & Adams, A. (2022). *Just-in-time liquidity on the Uniswap protocol*. Uniswap Labs working paper. <https://doi.org/10.2139/ssrn.4382303>