Quantisation Microstructure

The Tick-Race Law of Rebalancing Dust

K.R. Ryan

Independent Researcher | ORCID 0009-0004-6295-7040 | gancor@gancor.xyz

July 2026

Keywords tick quantisation · rebalancing microstructure · corrective swap · dust ledger
Aerodrome · Uniswap V3 · automated market maker

JEL C58, D47, G14, G23

∴

Abstract. This paper identifies *quantisation microstructure*, the structured, directional, heavy-tailed flows that rounding choices in automated liquidity management generate and that imitate execution economics. The instance is the law that Operator Microstructure[1] left open. The execution residual measured there is direction-conditioned, since after up-moves the leftover lands on the bought token in more than nine events in ten, while after down-moves the split is nearly even, and no recorded mechanism explained both facts. Replaying the deployed pipeline per event reproduces the measured dust credit at a median predicted-to-measured ratio of 1.000 and the leftover side in 98.4 to 99.9% of 31,340 production rebalances (Aerodrome Slipstream on Base), with no fitted parameters. The map then reduces to a *tick race* free of range width. The deployed sizing computes its target ratio at the floor tick while executing and minting at the continuous price. The floor bias pushes every mint toward the token0 side by the price's within-tick offset, the corrective swap's impact shifts the mint requirement by its tick displacement, and the leftover side is decided by the sign of their sum. One proposition states the race. With the offset empirically uniform, the race reproduces the measured binding probabilities of 0.078 and 0.524 to within a point and a half, and the displacement has an ex-ante law in a footprint ratio, position value over active liquidity. With the companion paper's placement-quantisation result, both channels of the directional dust asymmetry are floor functions interacting with the sign of price impact. All results are single-operator, but quantisation signatures are code-determined and recoverable from public data alone, and cross-operator replication of the recipe turns the artefact into a taxonomy of automated liquidity management.

1. Introduction

Operator Microstructure[1] (OM) decomposed the directional asymmetry of concentrated liquidity dust flows into a frequency channel and a value channel. The frequency channel closed. Rebalances fire by first passage to the operator's trigger lines, and with the trigger fraction measured from the data, theory slopes of one are recovered wherever the regressor carries signal. The value channel did not close. OM measured the per-event execution residual $\hat{\varepsilon}$, credit over corrective-swap notional, and found it direction-conditioned. After up-moves the leftover lands on the bought token in 93.8% of volatile–stable and 90.0% of volatile–volatile events; after down-moves the split is nearly even. The unconditional binding probability of the companion paper[2] (ME), $\mathbb{P}(\varepsilon < 0) = 0.26$ on its substantial-swap sample, is the mixture of a one-sided regime and a balanced one. OM could not separate the candidate drivers; its residual proxy conflates fee, price impact, and deliberate sizing bias, and it names the post-swap tick as the datum that would separate them. Its future-work section names the object this paper supplies, the per-event law of the direction-conditioned execution residual.

The law is simpler than the decomposition that motivated it. Three results are established on the same production dataset.

First, the per-event map is deterministic and exact (Section 4). Reconstructing the deployed pipeline's actual inputs per event reproduces the measured dust credit at a median predicted-to-measured ratio of 1.000 across all 31,340 events, and the leftover side in 98.4 to 99.9‰ of them, cell by cell, with nothing fitted. The inputs are the burn proceeds from the position manager's collect event, the recycled dust stock from an exact ledger walk, the observed corrective swap, and the post-swap mint arithmetic. The value channel is not noise attached to a swap; it is arithmetic.

Second, the map reduces to a tick race (Section 5). The deployed sizing routine computes its target composition at the price of the *floor tick*, while the sizing input, the swap execution, and the mint all see the continuous within-tick price. Because the token0 value share falls in price, flooring biases every target toward token0 by $g\,\zeta$, where g is the local share-per-tick slope and $\zeta \in [0, 1)$ is the within-tick offset. The corrective swap's own impact moves the mint requirement by $g\,\Delta$, with Δ its realised tick displacement, signed along the trigger direction in roughly nine events in ten. The shared factor g cancels, and the leftover side is $\mathrm{sign}(\zeta + \Delta)$ identically across pools whose ranges span two orders of magnitude in width. Up-moves have both terms pushing token0, so only reversed impacts flip the side; down-moves race a uniform offset against a median impact of about a third of a tick. The measured probabilities follow with no free parameters.

Third, the impact leg has an ex-ante law (Section 6). The corrective notional is a trigger-geometry functional of the old and new ranges, verified at rank correlation 0.995 against the observed swaps, and the tick displacement is that notional over the pool's active liquidity. The one operator-scale quantity in the law is a footprint ratio, position value over active liquidity. Everything else is geometry.

The paper also settles what the ledger is not (Section 7). The position manager recycles a depositor-keyed dust stock into every rebalance, and that stock is real, token0-heavy, and comparable in size to the credits. It is also absorbed by the sizing at every event and never reaches the mint. A deviation regression puts its passthrough coefficient at 0.01 against the 1.0 of a passthrough mechanism. Correlations between the stock and the credits are the persistence of the credit sequence, not a mechanism, and an earlier causal reading from stock-zeroing events collapses under within-pool controls from a 19-point effect to five points.

OM attributed its frequency-channel excess on volatile–volatile pairs to asymmetric range placement, a grid-quantisation artefact of the operator's placement rounding. The value channel now traces to a second quantisation choice, the floor-tick target ratio. Both channels of the directional dust asymmetry on this operator are tick-quantisation microstructure. Traditional microstructure learned that apparent anomalies can be artefacts of discreteness in the trading mechanism; automated liquidity management reproduces the situation with the artefact written in verifiable code, so the attribution can be made exact rather than argued. The concluding section states what that means for inference from on-chain aggregates, for the companion papers' assumption blocks, and for any operator wishing to remove the bias.

2. Related work

The companion papers are the immediate antecedents. GS[3] proved the USD-denominated directional asymmetry of the dust flow on volatile–stable pairs; ME[2] derived the per-event jump law of the dust ledger and measured the unconditional execution residual whose direction-conditioning OM[1] later discovered; OM closed the frequency channel and named the object supplied here. The V3 amount equations[4] underlie the mint arithmetic throughout.

The theoretical literature on V3 liquidity provision models costs incurred between rebalances. Hasbrouck *et al.*[5] model provision as a covered call, Cartea *et al.*[6] analyse predictable loss, and Milionis *et al.*[7] introduce loss-versus-rebalancing, the return gap driven by arbitrage against stale prices; Heimbach *et al.*[8] measure per-position risk and return empirically. The object here differs in kind and in the scale of its cause. The dust flow arises inside the rebalance transaction itself, and its directional law turns out to be set neither by adverse selection nor by execution quality but by the interaction of a rounding choice with the sign of the correction's own impact.

The leftover itself is well documented as an accounting object. Elsts[9] documents the minimum-liquidity mint rule, which strands the excess on one side; Uniswap's own tooling names a recipient for "leftover dust from swap";[10] production vaults either redeploy the leftover as a single-sided order[11] or pay unused balances out pro rata;[12] and the sizing problem, including the corrective swap's own impact on the target ratio, is solved in closed form to a zero-remainder target by practitioner libraries.[13] Throughout that corpus the dust is a nuisance to be absorbed or optimised away. None of those systems accumulate the residual into a persistent ledger recycled across positions. That structure is the studied manager's own design; the companion papers' conservation and flow results concern the ledger, and this paper derives the law of the per-event credit that feeds it. The nearest formal treatment of tick-grid arithmetic verifies pool invariants under rounding as a bounded, one-sided perturbation,[14] with no rebalancing content, and in traditional markets discreteness of the price grid is a documented source of apparent anomalies, from price clustering[15] to biased moment estimators under grid rounding.[16] What is new here is the derivation. No prior work, to the author's knowledge, derives the direction, magnitude, or binding probabilities of rebalancing dust from tick quantisation, or from anything else.

3. Setting and data

The setting, dataset, and notation follow OM, and only what is needed here is restated. The operator runs an off-chain manager over a whitelisted on-chain position manager (PM) on Aerodrome Slipstream pools on Base. A rebalance burns the old position, collects proceeds and fees, optionally swaps in-pool to rebalance composition, mints into the new range, and credits any leftover to a depositor-keyed dust ledger, emitting per-token credit events. The paper window is pinned at block 45,746,143 with 126,339 rebalances across 48 pools with at least one hundred events each. Joining each rebalance to its corrective swap by transaction hash, restricting to trigger-zone events at the measured trigger fraction $\varphi = 0.25$, and excluding AERO-token pools whose reward accounting contaminates balance reconstruction gives the 31,340-event panel used throughout. The cut from the full window is the trigger-zone share of the event mix (roughly 54%) times the corrective-swap join coverage (52.5%), less the AERO pools. The retained panel's share of total dust value is not separately computed. The join and every headline number of OM's Table 3 were reproduced exactly before any extension.[1]

Three inputs beyond OM's panel are used. The position manager's collect events, fetched for all 66,190 join transactions, give the exact burn-plus-fees proceeds per event. An exact ledger walk over the credit and rebalance event stream reconstructs the recycled dust stock per depositor, portfolio, and token at every event; the walk's in-call ordering was validated against the contract's event layout, and a live on-chain spot check of its terminal state agrees. The per-swap rows of the pool event stream supply the pre- and post-swap sqrt prices and the pre-swap active liquidity.

Throughout, $\omega_0(s; s_a, s_b)$ denotes the token0 value share of a mint at sqrt-price s in the range $[s_a, s_b]$, V the position value at the event in token1 units, N_1 the corrective-swap input value in the same units, L_p the pool's active liquidity, and $q^{\pm}$ the probability that the leftover lands opposite the bought token after up- and down-moves respectively. Table 3, in Section 9, summarises the notation.

[1] Estimator scripts, artefacts, and the note series from which the paper was assembled live in the paper's working notes; a consolidated reproduction harness is at https://github.com/g4nc0r/operator-microstructure. Prices enter through the same era snapshot as OM; the model side works in raw token units, so the snapshot affects only the USD framing and the side inference on near-tie events.

4. The per-event map is exact

The deployed sizing path is the PM's automatic mode, used by 125,768 of the 126,339 window rebalances. Its implementation is short. It reads the pool state once, pre-swap; computes the new range's target token0 share at the current (floored) tick; computes the current share from available balances at the continuous sqrt price; sizes an exact-input swap of the surplus token to close the gap at the pre-swap mid, with no fee or impact adjustment; and caps or skips only in rare regimes (a 200 basis-point dead-band, a 90% input cap, a 1% sqrt-price limit; together about 3% of events). The mint then consumes at the post-swap price, and the refund becomes the dust credit. Sizing against the pre-impact ratio is therefore deployed fact, read from the implementation rather than assumed, and the whole pipeline is available to replay.

Replaying it with observed inputs closes the value channel at the event level. Available balances are the collect proceeds plus the recycled stock. Post-swap balances subtract the observed pool-frame swap amounts, so no execution model is needed. The mint at the observed post-swap price then determines the leftover pair by the standard binding arithmetic. Table 1 reports the comparison against the measured credits.

Table 1. Exact replay against measured dust credits. Rank correlations are of log credits.

	Volatile–stable		Volatile–volatile	
	Up	Down	Up	Down
Events	9,665	8,937	6,876	5,862
Median predicted/measured credit	1.000	1.000	1.000	1.000
Side agreement (%)	99.9	98.4	99.9	99.6
Rank correlation	0.997	0.992	0.999	0.995

The replay is a closure test rather than a prediction. It establishes that the credits contain no unobserved input, and its residual side disagreements sit in near-tie events, where the era price snapshot decides the labelling of a wei-scale split. The direction-conditioned binding probabilities become outputs of the same arithmetic. Measured, $q^+ = 0.078$ and $q^- = 0.524$; replayed, 0.078 and 0.534. The heavy tail OM measured (mean $|\hat{\varepsilon}|$ up to six times the median) is reproduced with the rest. What remains of the value channel is not measurement but explanation. The replay is exact and opaque; the law is the structure that generates its outputs.

5. The law is a tick race

5.1. The stock is absorbed

One candidate structure is eliminated first. The recycled dust stock is comparable in size to the credits (its median is 86 to 95% of the median credit) and token0-heavy at three events in four, and its correlation with the credits suggested a mechanism in which the stock rides through the mint and decides close calls. It does not. Regressing the post-swap deviation from the sizing target on the stock's share of position value gives coefficients of 0.006 to 0.019 in magnitude against the unit coefficient of a passthrough mechanism, with a median absolute deviation near one basis point against stock shares of 130 to 170 basis points. The sizing includes the stock in its balances (predicted swap size matches the observed input at a median relative error of 2×10^{-4} with the stock included and 9×10^{-2} without), and re-sells it every event. The stock's statistics are those of the lagged credit sequence, and the apparent causal effect of stock-zeroing events shrinks from nineteen points to five under within-pool controls. The side law must be stock-free.

5.2. Two terms of the same scale

With the stock absorbed, the token0 surplus at the mint, in value-share units and positive when the leftover lands on token0, is

$$\Sigma_0 = \underbrace{\left[\omega_0(s_{\lfloor t \rfloor}) - \omega_0(s_{\mathrm{pre}})\right]}_{b\,=\,g\zeta\,\geq\,0} + \underbrace{\left[\omega_0(s_{\mathrm{pre}}) - \omega_0(s_{\mathrm{post}})\right]}_{g\Delta} + O(\eta N_1/V), \tag{1}$$

where t is the continuous pre-swap tick coordinate and $s_{\lfloor t \rfloor}$ the sqrt price of its floor tick, g is the local magnitude of $\partial\omega_0/\partial\,\mathrm{tick}$, $\zeta = t - \lfloor t \rfloor \in [0, 1)$ is the within-tick offset of the continuous pre-swap price, Δ is the corrective swap's realised displacement in ticks (positive when the price moves upward), and the error term collects the fee shortfall on the bought leg (η the pool fee, N_1/V the swap-to-value fraction) together with mint rounding, both second-order at the dominant one-basis-point fee tier.

The first term is the discovery. The deployed target-ratio routine evaluates at the floor tick's sqrt price while every other stage of the pipeline, the current-ratio computation, the swap, and the mint, sees the continuous price. Since ω_0 falls in price, flooring biases every target toward token0 by the within-tick offset, scaled by the same local slope g that scales the impact term. The bias is not small relative to what it races; both terms are of order g, and g cancels in the sign.

Proposition 1 (Tick-race law)**.** *To the stated order, the leftover side is* $\mathrm{sign}(\zeta + \Delta)$ *and the credit is* $g\,|\zeta + \Delta| \cdot V$. *In particular, the side condition contains no range width, no fee tier, and no notional; conditional on direction,*

$$q^{+} = \mathbb{P}\big(\zeta + \Delta < 0 \,|\, \mathrm{up}\big), \qquad q^{-} = \mathbb{P}\big(\zeta + \Delta > 0 \,|\, \mathrm{down}\big). \tag{2}$$

In words, the leftover token is decided by a race. The mint's target is computed at the floored tick, which sits a sub-tick offset ζ below the live price, while the corrective swap moves the price by Δ. Whichever of the two is larger in the relevant direction decides which token is left over, and neither range width nor fee tier nor notional enters that comparison.

The empirical inputs behave as the proposition needs (Fig. 1). The within-tick offset is uniform to good approximation (quartiles within 0.02 of the uniform 0.25 and 0.75, medians 0.47 and 0.52 by direction), as equidistribution of a diffusing price within a tick would suggest, and is uncorrelated with the displacement (rank correlations within ± 0.04 by direction; the joint and independence-product forms of q^{-} agree within half a point).

The impact is signed along the trigger direction in 87.9% of up-moves and 94.0% of down-moves, with a median displacement along that direction of about a third of a tick. Up-moves therefore have both terms pushing token0, and only an impact reversal larger than the offset flips the side; down-moves oppose the terms, and on the nine in ten whose impact is along the trigger direction the side is the race $\zeta > |\Delta|$, a uniform draw against a sub-tick impact. The up-side reversal rate, 12.1% of up-moves, is thus an input; the law's content there is the conversion (only reversals larger than the offset flip), which takes 0.121 to the predicted 0.071.

Evaluated on the realised (ζ, Δ) pairs, equation (2) gives $q^{+} = 0.071$ and $q^{-} = 0.510$ against the measured 0.078 and 0.524; the full two-term form of equation (1), with the fee tie-breaker included, gives 0.080 and 0.528 and matches the measured side in 96.1 to 98.3% of events across the four class-direction cells, against the exact replay's 98.4 to 99.9%. Sampling errors at these counts are ± 0.002 on q^{+} and ± 0.004 on q^{-}, so the two-term values sit within one standard error of the measured probabilities while the bare race sits three to four below them, short by exactly the tie-breaker terms it omits.

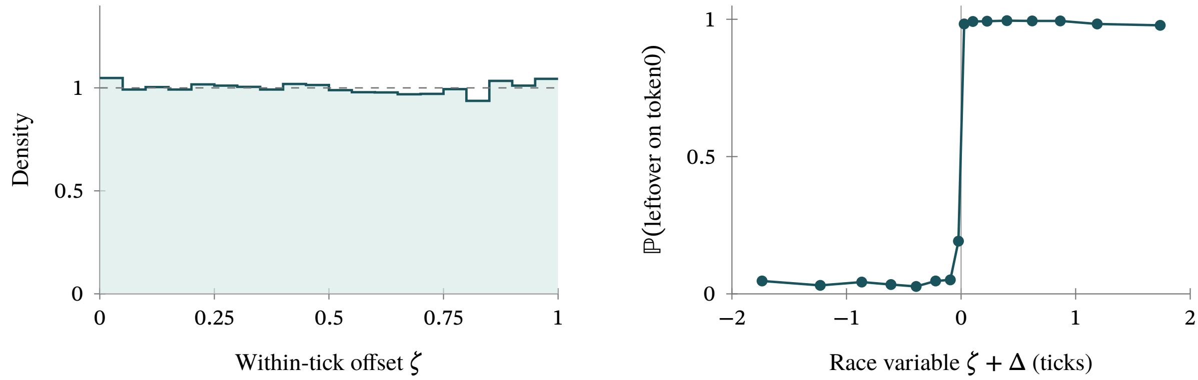

Fig. 1. The tick race, 31,340 events. Left: density of the within-tick offset ζ against the uniform reference (dashed). Right: probability that the leftover lands on token0 by the race variable $\zeta + \Delta$; bins beyond ± 2 are not shown.

Two splits probe the law's stability. Across range-width quartiles spanning under 34 to over 100 ticks, the race matches the measured side in 100.0, 99.7, 98.7, and 94.9% of events. The measured q^- moves from 0.83 to 0.34, and the race tracks it within 1.5 points everywhere but the widest quartile, where fee-scale ties bite; the mapping is width-free even as its inputs are not. Across a temporal split at the window's median block, the unconditional probabilities drift substantially (q^+ from 0.109 to 0.048, q^- from 0.404 to 0.648, the reversal rate from 0.166 to 0.077) while the race tracks each half within 1.5 points at hit rates near 98%. The inputs move with the operator's conditions; the mapping does not.

The unconditional binding probability follows as the firing-frequency mixture of the two regimes under OM's frequency law; on the joined panel the mixture is 0.289, and ME's measured $\mathbb{P}(\varepsilon < 0) = 0.26$, taken on its substantial-swap sample, sits where the sample-composition shift puts it. The sum-based versus median-based estimator dissociation that motivated OM's measurement is the tail of $g|\zeta + \Delta| \cdot V$; the same race that sets the side sets the magnitude, and g scales inversely with range width, so narrow ranges throw heavy tails.

6. The impact leg, ex ante

Proposition 1 takes the realised Δ as input. It has an ex-ante law. For an exact-input swap of token1-value N_1 against active liquidity L_p at sqrt price s, constant-liquidity arithmetic gives the displacement

$$\Delta \;\approx\; \pm\,\frac{2\,N_1}{s\,L_p\;\ln 1.0001}, \tag{3}$$

with the upper sign for a token1-input swap, which moves the price upward, and the notional is itself geometric; at a trigger-line firing the imbalance is the value-share gap between the old range's composition and the new range's target at the trigger price, so $N_1 \approx |\omega_0(s; s_a, s_b) - \omega_0(s; s'_a, s'_b)| \cdot V$, primes denoting the new range as in GS. Substituting,

$$\Delta \;\approx\; \frac{2\,|\omega_0(s; s_a, s_b) - \omega_0(s; s'_a, s'_b)|}{\ln 1.0001} \cdot \frac{V}{s\,L_p}, \tag{4}$$

a pure trigger-geometry factor times the operator's *footprint ratio*, position value over the token1-denominated active-liquidity scale sL_p. Every quantity on the right is known before the swap executes.

Table 2 and Fig. 2 report the verification chain. The one-step form (3) matches the realised displacement with perfect sign agreement; it is exact below half a tick and attenuates to a rank correlation of 0.92 beyond two ticks, where crossed initialised ticks violate constant liquidity in the expected direction, and the attenuation concentrates in the same multi-tick tail that carries the value-weighted flow, so the ex-ante law is weakest where the credits are largest. The geometric notional matches the observed input more tightly still, and the fully ex-ante form (4) loses little of the one-step correlation. Substituting the ex-ante magnitudes into equation (2) gives q^- of 0.469 to 0.491 against the measured 0.524 and, with the realised impact signs, q^+ of 0.068 against 0.078. The gap is carried by the constant-liquidity attenuation on the multi-tick tail and, on the up side, by the impact-reversal events that the trigger direction alone cannot predict.

Table 2. The Δ verification chain (upper panel) and the realised displacement by pool class (lower panel), 31,340 events. Rank correlations are of log magnitudes.

	Rank correlation		Median ratio	
	Up	Down	Up	Down
One-step $\hat{\Delta}$ vs realised Δ	0.977	0.979	0.994	0.991
Geometric notional vs observed input	0.995	0.995	1.020	1.034
Ex-ante $\hat{\Delta}$ vs realised Δ	0.967	0.972	–	–
	Volatile–stable		Volatile–volatile	
	Up	Down	Up	Down
Median $\lvert\Delta\rvert$ (ticks)	0.64	0.76	0.39	0.22
$\mathbb{P}(\lvert\Delta\rvert > 1)$	0.44	0.46	0.36	0.31

The footprint distribution is the one operator-scale object in the law. Its median puts the typical corrective swap below a tick, and its tail (a 95th percentile of sixteen to twenty ticks) is what pulls q^- from the constant-impact value of about 0.66 down to one half. Cross-operator, the footprint ratio V/sL_p is the single quantity to measure per manager; the rest of the law transfers as geometry.

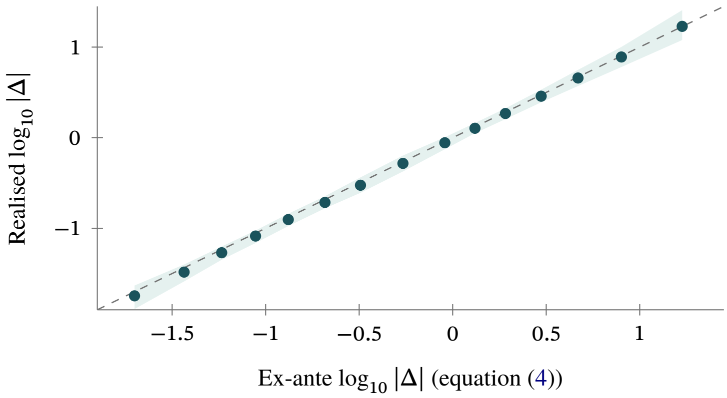

Fig. 2. The ex-ante impact law. Median realised tick displacement against the footprint-ratio prediction of equation (4) in quantile bins, with the interquartile band (shaded) and the diagonal (dashed).

7. What the ledger is, and is not

The dust ledger deserves its own accounting, because a natural reading of the data, and one I adopted at first, assigns it a causal role it does not have. The stock exists, recycles into every rebalance, and is token0-heavy for the same reason the credits are. Its size sits at the credit scale, so the naive decomposition "credit equals carried stock plus fresh residual" fits the levels well. Conditioning the down-move side on the carried composition also produces large, monotone differences. All of that is persistence.

The stock at an event is the credit sequence of the recent past. Surplus magnitudes and signs persist within a pool because g, the footprint, and the impact sign distribution persist, and every correlation between stock and outcome survives exactly as long as the pool fixed effect that generates both. The interventions that looked causal, stock zeroings by other pools' rebalances in the same portfolio, lose three quarters of their effect under within-pool controls. The deviation regression of Section 5 bounds the mint's sensitivity to the stock at the basis-point level, and the residual five points of the zeroing contrast are consistent with regime clustering of cross-pool activity and are not further attributed here.

For ME, the consequence is a revision of the independence discussion rather than of any unconditional result; the execution residual is direction-conditioned through the sign of Δ, and its serial dependence is pool-level persistence of $(g, V/sL_p)$, not ledger memory. For OM, the published κ_{exec} remains an estimator-true statement about composite credits, but its "execution error" gloss should now be read in light of this paper's decomposition.

8. Discussion

8.1. Quantisation microstructure

OM resolved its frequency-channel anomaly into asymmetric range placement produced by grid-aligned rounding, and left the value channel open. The value channel now resolves into a second rounding choice, a target ratio computed at the floor tick. Neither is a strategy; both are one-line implementation choices whose interaction with direction produces the asymmetries this research line set out to explain. On volatile–stable pools token0 is typically the volatile leg, so the floor bias points the ledger's accumulation at the volatile asset in both directions, which is the mechanical content of the earlier observation that the ledger accumulates the just-appreciated asset. The lesson for inference from on-chain aggregates is that the dominant microstructure of this operator's dust flows is not execution quality, adverse selection, or valuation, but the placement of two floors. The same fact has a use. Quantisation choices are code-determined, stable between deployments, and recoverable from public data alone, so they fingerprint the operator that made them; replicating the recipe of this paper across managers turns the artefact into a taxonomy.

8.2. Portability and the operator's remedy

The law's form is portable to any V3-style manager that sizes at a quantised price and corrects in-transaction; the recipe is the same join, the same walk, and the same race, and the footprint ratio is the one number to measure per operator. The scope condition is checkable in source, because the mechanism is implementation-dependent rather than universal. The canonical SDK snaps ticks to nearest rather than flooring,[17] so floor arithmetic is a manager's choice, and at least one production vault computes its ratios at the continuous pool price in its on-chain arithmetic and carries no floor bias there.[12] How the deployed population divides is an open empirical question, and it is the natural first measurement of the cross-operator taxonomy.[2] The remedy is equally portable. Computing the target at the continuous sqrt price, or minting against the sizing price, removes the b term and collapses q^- toward q^+, eliminating most of the typical credit. The value-weighted tail survives that change; it is impact-driven, and its removal takes an impact-aware sizing. The constant-liquidity closed form[13] is exact until the swap crosses initialised ticks, so the multi-tick swaps that carry the tail need a sizing that walks the tick ladder. Because the studied manager is live and modifiable, the statistics reported here are era-specific; a configuration change that alters the sizing or the trigger

[2] The cross-operator programme has since begun. A companion paper[18] applies OM's trigger-fraction recovery unilaterally, from public chain data alone, to 81 independently operated production managers across 102 operator-pool cells. The sizing-arithmetic division asked about here, floor against nearest, is not among its fingerprint axes and remains unmeasured.

geometry would render them historical, and the measurement recipe rather than the point estimates is what transfers.

8.3. Limitations

All results are one operator's configuration on one chain era, with era-snapshot prices entering the USD framing and the side inference on near-tie events. The constant-liquidity form of the impact law attenuates on multi-tick swaps, and the impact-reversal rate (six to twelve percent by direction, from same-transaction flows) is measured, not modelled. The law was identified and evaluated on the same panel; the identification is carried by the code derivation rather than by fit, and cross-operator replication is the out-of-sample test. The volatile–volatile drift cell of OM's frequency law remains untested here; a cross-pool factor instrument was built, but at seven pools it lacks the power to decide the cell. I do not claim the uniformity of the within-tick offset as a property of the price process; it is an empirical regularity consistent with equidistribution, verified only on this panel.

9. Notation

Hats denote model-implied values, as in the predicted displacement $\hat{\Delta}$ of Table 2. The trigger fraction φ is OM's and is distinct from GS's position-value function ϕ, which does not appear in this paper.

Table 3. Notation used throughout the paper.

Symbol	Meaning
The tick race.	
$\omega_0(s; s_a, s_b)$	token0 value share of a mint at sqrt-price s in range $[s_a, s_b]$
s'_a, s'_b	sqrt-price bounds of the new range after a rebalance (GS)
Σ_0	token0 surplus at the mint, in value-share units; positive when the leftover lands on token0
s; $s_{\text{pre}}, s_{\text{post}}, s_{\lfloor t \rfloor}$	sqrt price (V3 / GS), with pre-swap, post-swap, and floor-tick values; t the continuous pre-swap tick
ζ	within-tick offset of the continuous pre-swap price, $\zeta = t - \lfloor t \rfloor \in [0, 1)$
Δ	corrective swap's realised price displacement in ticks, positive upward
g	local magnitude of $\partial\omega_0/\partial$ tick
b	tick-floor target bias $b = g\,\zeta$, in value-share units
Scales and probabilities.	
V	position value at the event, token1 units (GS)
N_1	corrective-swap input value, token1 units
L_p	in-range pool liquidity at the pre-swap tick (GS)
η	pool fee rate
q^+, q^-	probability the leftover lands opposite the bought token, by direction
ε	per-event execution residual, $\mathbb{P}(\varepsilon < 0)$ its unconditional binding probability (ME)
$\hat{\varepsilon}$	execution residual, credit over corrective-swap notional (OM)
$\varkappa_{\text{exec}}$	execution factor of OM's value-channel decomposition (OM)
φ	trigger fraction, 0.25 in the paper window (OM)

References

[1] Ryan, K.R. "Operator Microstructure: Decomposing the Directional Asymmetry of Concentrated Liquidity Dust Flows." *SSRN preprint 7164459* (2026) https://ssrn.com/abstract=7164459.

[2] Ryan, K.R. "The Hidden Microstructure of Shared Balance Concentrated Liquidity: A Master Equation for the Dust Ledger and Propagation of Chaos." *SSRN preprint 6745218* (May 2026, revised July 2026) https://ssrn.com/abstract=6745218.

[3] Ryan, K.R. "The Geometric Siphon: Existence, Equilibrium, and Directional Properties of the Residual in Concentrated Liquidity Portfolios." *SSRN preprint 6686798* (2026) https://ssrn.com/abstract=6686798.

[4] Adams, H., Zinsmeister, N., Salem, M., Keefer, R., Robinson, D. "Uniswap v3 Core." *Uniswap Labs Technical Whitepaper* (2021) https://uniswap.org/whitepaper-v3.pdf.

[5] Hasbrouck, J., Rivera, T.J., Saleh, F. "An Economic Model of a Decentralized Exchange with Concentrated Liquidity." *Management Science* **72.5** 3666–3683 (2026) https://doi.org/10.1287/mnsc.2024.04510.

[6] Cartea, Á., Drissi, F., Monga, M. "Decentralized Finance and Automated Market Making: Predictable Loss and Optimal Liquidity Provision." *SIAM Journal on Financial Mathematics* **15.3** 931–959 (2024) https://doi.org/10.1137/23M1602103.

[7] Milionis, J., Moallemi, C.C., Roughgarden, T., Zhang, A.L. "Automated Market Making and Loss-Versus-Rebalancing." *arXiv:2208.06046* (2022) https://arxiv.org/abs/2208.06046.

[8] Heimbach, L., Schertenleib, E., Wattenhofer, R. "Risks and Returns of Uniswap V3 Liquidity Providers." In *Proceedings of the 4th ACM Conference on Advances in Financial Technologies (AFT 2022)*, pp. 89–101 (2022) https://doi.org/10.1145/3558535.3559772.

[9] Elsts, A. "Liquidity Math in Uniswap v3." *Technical note* (2021) https://atiselsts.github.io/pdfs/uniswap-v3-liquidity-math.pdf.

[10] Uniswap Labs. "Swap and Add Liquidity." *Uniswap v3 SDK guides*, accessed July 2026, https://developers.uniswap.org/docs/sdks/v3/guides/managing-liquidity/swap-and-add.

[11] Charm Finance. "Alpha Vaults Strategy." *Documentation*, accessed July 2026, https://learn.charm.fi/charm-finance/alpha-vaults/strategy.

[12] Gamma Strategies. "Hypervisor." *Source repository*, accessed July 2026, https://github.com/GammaStrategies/hypervisor.

[13] Aperture Finance. "UniV3Automan (OptimalSwap library)." *Source repository*, accessed July 2026, https://github.com/Aperture-Finance/uniswap-v3-automan.

[14] Tranquilli, J., Gupta, N. "Formal State-Machine Models for Uniswap v3 Concentrated-Liquidity AMMs: Priced Timed Automata, Finite-State Transducers, and Provable Rounding Bounds." *arXiv:2512.06203* (2025) https://arxiv.org/abs/2512.06203.

[15] Harris, L. "Stock Price Clustering and Discreteness." *Review of Financial Studies* **4.3** 389–415 (1991) https://doi.org/10.1093/rfs/4.3.389.

[16] Gottlieb, G., Kalay, A. "Implications of the Discreteness of Observed Stock Prices." *Journal of Finance* **40.1** 135–153 (1985) https://doi.org/10.1111/j.1540-6261.1985.tb04941.x.

[17] Uniswap Labs. "v3-sdk." *Source repository*, accessed July 2026, https://github.com/Uniswap/v3-sdk.

[18] Ryan, K.R. "Operator Fingerprinting: Recovering and Defending the Hidden Rebalancing Policy of Concentrated Liquidity Managers." *SSRN preprint 7202340* (2026) https://ssrn.com/abstract=7202340.