Operator Microstructure

Decomposing the Directional Asymmetry of Concentrated Liquidity Dust Flows

K.R. Ryan

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

July 2026

Keywords rebalancing microstructure · trigger geometry · first-passage law · dust ledger Aerodrome · Uniswap V3 · automated market maker

JEL C58, D47, G14, G23

∴

Abstract. On a production concentrated liquidity position manager, the USD-valued dust flows credited at rebalances are directionally asymmetric. The first companion paper[1] explains the volatile–stable case through a numéraire theorem; the second[2] closes on the observation that the asymmetry is stronger on volatile–volatile pairs, outside that theorem's domain. This paper decomposes it into a frequency channel and a value channel, measured on the V9 production dataset (126,339 rebalances, 48 pools, Aerodrome Slipstream on Base). A directional firing proposition gives the frequency channel as an exact first-passage law with no free parameters. A second proposition and its corollary give the swap-free per-event dust credit in closed form, with its sign inversion under token-order reversal, verified to machine precision against a direct V3 mint. The first-passage barriers are the operator's trigger lines, not the range bounds, recovered from the data as a discontinuity in the firing-position distribution. With barriers so placed the theory slopes match their predicted value of one wherever the regressor carries signal, intercepts are indistinguishable from zero, and a replay test attributes the one remaining null to a noise-dominated regressor. The volatile–volatile excess is attributed to asymmetric range placement, a grid-quantisation artefact, jump-driven exits and labelling artefacts ruled out. The value channel, from joining each rebalance to its corrective swap, is the near-cancellation of two large opposing execution factors, with the residual direction-conditioned. The numéraire component, exact in its domain, is bounded at production scale one to two orders below the execution factors. The directional asymmetry of production dust flows is carried by operator microstructure. All results are single-operator; the law's form and the trigger-measurement recipe are portable, and cross-operator replication is the immediate test.

1. Introduction

A concentrated liquidity position must be rebalanced when the pool price leaves the neighbourhood of its range. Each rebalance leaves a leftover on one side of the pair, and the position manager studied in the companion papers retains that leftover in a depositor-keyed dust ledger. The Geometric Siphon[1] (GS) proved that on a volatile–stable pair the USD-valued flow is directionally asymmetric, with the up-move side retaining strictly more value, through a closed-form identity in the position value; the mechanism is USD valuation of a directionally blind token flow, and reversed token orderings supply a built-in negative control. The Hidden Microstructure[2] (ME) derived the exact per-event jump law of the dust ledger and closed with the observation that motivates this paper. On the V9 production data the directional asymmetry *persists on volatile–volatile pairs*, where the numéraire theorem (GS's USD-denominated directional asymmetry) does not apply, and is stronger there than inside the theorem's domain; ME reports median directional ratios of 1.11 on volatile–volatile pairs against 1.05 on volatile–stable. Something other than the numéraire mechanism carries it.

This paper identifies what does. The observable, the ratio of up-move to down-move USD dust flows, factors into a *frequency channel* (how often rebalances fire in each direction) and a *value channel* (how much each firing credits). The two channels obey different laws.

The frequency channel obeys an exact first-passage law (Proposition 1). Modelling the pool tick between rebalances as drifted Brownian motion absorbed at directional barriers gives the up-to-down count ratio in closed form, with no free parameters once the trigger fraction is measured. Empirically, the barriers are not the range bounds. The V9 position manager fires its rebalances when the price comes within a configured fraction φ of the range width of either bound, and that fraction is recovered from the data; the distribution of the price's position within the previous range at firing shows sharp discontinuities at exactly $\varphi = 0.25$ from either edge. With barriers placed on the trigger lines, a two-regressor fit of measured drift and measured placement against the directional count ratio yields slopes of 1.27 and 1.06 on volatile–stable pairs and a placement slope of 1.02 on volatile–volatile pairs, against a theoretical value of one, with intercepts indistinguishable from zero.

The same fit explains ME's closing observation. Volatile–volatile pools place their ranges asymmetrically around the price, a rounding artefact of grid-aligned placement rather than a strategy, leaving the up-barrier nearer; the first-passage law converts that geometry into an up-exit surplus mechanically. The two alternative explanations fail directly. Exit overshoots are directionally symmetric, ruling out one-sided jumps, and the excess appears with the same sign whichever side of the pool WETH occupies, ruling out labelling artefacts of the V3 token ordering.

The value channel tests the prior theory against production microstructure. For swap-free rebalances the numéraire component is derived here in closed form (Proposition 2). The per-event dust credit admits an exact expression, and the up-to-down value ratio inverts under token-order reversal (Corollary 1), sharpening GS's negative control into an exact identification within the swap-free domain. But production rebalances are not swap-free. Joining each rebalance to its own corrective swap shows the realised value wedge to be the near-cancellation of two much larger execution factors, and the execution residual whose unconditional law ME measured is direction-conditioned. After up-moves the leftover lands on the bought token more than nine times in ten; after down-moves it splits nearly evenly. On the measured V9 range widths the numéraire component is one to two orders below the execution factors.

The conclusion is therefore not the one this study was designed to reach. The directional asymmetry of production dust flows is dominated by operator microstructure, namely where the operator places ranges and how its corrective swaps execute, and not by the numéraire mechanism, which is exact in its domain but quantitatively small at the range widths a production manager actually uses. GS's theorems are unaffected; this paper adds a production-scale bound on the numéraire contribution and a measured account of what carries the asymmetry.

All measurements are of a single operator's contract family, whose trigger fraction, placement policy, and swap-sizing behaviour are the quantities under study. What is portable is the method. The law's form is operator-independent, the trigger fraction is measurable from any operator's firing-position distribution, and the decomposition into drift, placement, and execution terms transfers unchanged. Cross-operator replication is the natural test and remains open.

Section 2 places the paper in the literature. Section 3 describes the dataset and the trigger geometry. Section 4 states the decomposition and the two channel laws. Section 5 reports the frequency-channel measurements, including the explanation of the volatile–volatile excess. Section 6 reports the value-channel decomposition. Section 7 discusses portability, the numéraire bound, and limitations. Derivations are collected in Appendix A, estimator detail in Appendix B, and notation in Appendix C.

2. Related work

The V3 amount equations[3] supply the mint geometry from which Proposition 2 derives. The companion papers are the immediate antecedents. GS[1] identifies the geometric residual and proves the USD-denominated directional and exit asymmetries on volatile–stable pairs, with reversed token orderings as a negative control. ME[2] derives the exact per-event jump law of the dust ledger, reduces the swap correction to the one-dimensional residual used throughout Section 6, and closes with the volatile–volatile observation this paper explains. A concurrent SSRN preprint by I. P.,[4] on an independent implementation of the same mechanism, develops a graph-theoretic circulation condition for cross-pool dust flow together with a directional-convexity proposition. The proposition's proof independently derives the swap-free per-event credits of Section 4.2; the overlap is set out at Proposition 2. The trigger geometry, the frequency law, and the execution measurement pursued here are not addressed there.

Three lines of theoretical work model costs incurred by V3 liquidity providers between rebalances. Hasbrouck *et al.*[5] model provision as a covered call sold at intrinsic value, Cartea *et al.*[6] analyse predictable loss, and Milionis *et al.*[7] introduce loss-versus-rebalancing (LVR), the return gap driven by arbitrage against stale prices. The object studied here differs in kind. The dust flow arises at the rebalance itself, and its directional structure is set by the operator's trigger geometry, placement arithmetic, and swap execution rather than by adverse selection between rebalances; Section 7 keeps the two cost layers orthogonal. On the empirical side, Heimbach *et al.*[8] measure per-position risk and return across V3 strategies, whereas the measurements here are of rebalance-event microstructure within one manager. On the timing side, Agarwal and Gobet[9] cast the withdrawal of a position as an optimal stopping problem, and Bergault, Bieber and Sánchez-Betancourt[10] characterise endogenous exit through a stochastic control problem; both work in continuous price. The firing law here is the complementary exogenous object, the first-passage behaviour of a fixed trigger policy, measured rather than optimised.

The mathematics of the frequency channel is classical; the two-barrier firing probability is a textbook first-passage computation for drifted Brownian motion,[11] and the paper claims no novelty for it. The contribution is the identification. The empirically correct barriers are the operator's trigger lines rather than the range bounds, the trigger fraction is recoverable from the firing-position distribution, and with barriers so placed the law fits production firing counts unadjusted.

3. Setting, data, and trigger geometry

3.1. Dataset

The empirical instance is the PM V9 position manager on Aerodrome Slipstream (Base), the latest of the contract family whose dust-ledger law ME validated event by event. The family is operated by the present author (Section 7). The analysis window is pinned at block 45,746,143, where the indexed lake matches ME's declared event counts; within it, pools with at least 100 rebalances across all depositors give a universe of 48 pools and 126,339 rebalance events, all but 252 of them crediting dust. Pair classes follow GS's convention, USD-pegged tokens are stable and non-USD fiat-pegged tokens are volatile, giving 22 volatile–stable and 26 volatile–volatile pools.

Per-pool tick paths between rebalances come from an era-matched swap-event stream (6.3 million swaps across 44 of the 48 pools); USD valuation uses an era snapshot of token prices. Drift $\hat{\mu}$ is estimated from the net tick displacement of each pool's stream over its rebalance era, and $\hat{\sigma}^2$ from daily-coarsened tick changes, which are near-immune to the microstructure noise that inflates swap-scale quadratic variation. Appendix B records filters, estimator variants, and reconstruction hygiene inherited from ME.[1]

[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`.

Re-deriving the headline asymmetry with the sample, estimator, and filters fixed reproduces the companion papers' observation and sharpens it (Table 1).[2] On the value-weighted estimator the volatile–volatile median exceeds the volatile–stable median (1.192 against 1.131), and the ordering holds in all twelve cells of a sensitivity grid over two windows, two filter settings, and three estimator variants. The frequency-based estimator dissociates the classes more sharply. Volatile–stable pairs show their asymmetry in value at nearly balanced event counts (0.946), while volatile–volatile pairs show it in both value and frequency (1.285). The dissociation anticipates the two-channel structure of Section 4.

Table 1. Re-derived directional anchors, paper window, per-pool medians of Φ^+/Φ^- over 22 volatile–stable and 26 volatile–volatile pools. The above-one columns count pools with the ratio above one.

	Volatile–stable		Volatile–volatile	
Estimator	Median	Above one	Median	Above one
USD value	1.131	13/22	1.192	16/26
Event count	0.946	7/22	1.285	19/26
Unit-free geometric	0.747	8/22	1.274	18/26

3.2. *Trigger geometry*

One empirical fact precedes the model. PM V9 does not rebalance when the price exits the range; it rebalances when the price comes within a configured fraction φ of the full range width of either bound. The first-passage barriers of any exit model are therefore the *trigger lines* at $\ell + \varphi W$ and $u - \varphi W$ for a range $[\ell, u]$ of tick width W, not the bounds themselves.

The trigger fraction is recoverable from on-chain data alone. For each rebalance, let u^* be the price's position within the *previous* range at firing, normalised so the bounds sit at 0 and 1 and measured from the pool's own swap stream immediately before the rebalance block. The distribution of u^* shows sharp discontinuities at 0.25 and 0.75 in both pair classes (Fig. 1); on volatile–stable pairs, 6.3% of events fire in the five-percent band just inside the trigger line against 2.1% in the band just outside, with the mirror-image step at the top of the range. The era value is $\varphi = 0.25$, and Section 5 shows the regression itself rejects the mis-specified alternative $\varphi = 0.20$.

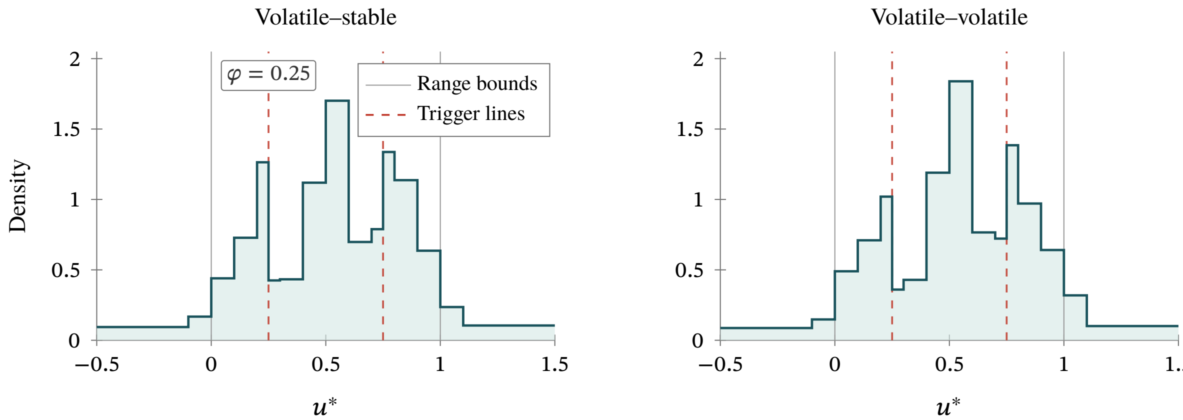

Fig. 1. Density of the price's position u^* within the previous range at firing, by pair class. The steps at $u^* = 0.25$ and 0.75 identify the trigger fraction $\varphi = 0.25$.

[2]ME quotes medians of 1.11 and 1.05 from an earlier estimator whose event filter (a 79,379-event subset) is not recoverable. The anchors in Table 1 are re-derived on the full pinned window with the estimator, universe, and filters fixed explicitly; the ordering of the two classes, which is the observation itself, is unchanged in every variant tried.

The same distribution decomposes the event population. About 54% of firings occur inside a trigger zone, as the proximity mechanism prescribes; about 11% occur beyond the range bounds, the tail of firings delayed by execution latency, rebalance cool-downs, and the manager's volatility-spike suppression; and about 34% occur mid-range, a distinct population of recalibration and width-change rebalances that fire regardless of price position. The mid-range population is directionally balanced (median log count ratio 0.016 volatile–stable, 0.041 volatile–volatile) and its direction tracks realised drift where drift is measurable. It obeys no first-passage law and is excluded from the frequency-channel tests, entering only the aggregate anchors of Table 1, which are therefore mixture statistics.

4. The two-channel decomposition

Write Φ^+ and Φ^- for the aggregate USD-valued dust flows credited at up-move and down-move rebalances of a pool (ME's D^+/D^-, renamed here; Appendix C), $N^\pm$ for the corresponding event counts, and $\bar{d}^{\,\pm} = \Phi^\pm/N^\pm$ for mean per-event credits. The observable then factors identically,

$$\frac{\Phi^+}{\Phi^-} = \frac{N^+}{N^-} \times \frac{\bar{d}^{\,+}}{\bar{d}^{\,-}}, \tag{1}$$

and the two factors are the frequency and value channels. The estimator mapping is immediate, with the count ratio isolating the frequency channel and the ratio of the value-weighted estimator to the count estimator isolating the value channel. Table 1 already displays the dissociation this predicts.

4.1. The frequency channel

Proposition 1 (Directional firing law)**.** *Let the pool tick between rebalances follow $X_t = \mu t + \sigma B_t$ from the placement point, with B_t standard Brownian motion, and let the rebalance fire at the first passage to $+w_u$ or $-w_d$ (the distances, in ticks, from the placement point to the upper and lower trigger lines). Then with $\theta = 2\mu/\sigma^2$,*

$$\frac{\mathbb{P}(\textit{up firing})}{\mathbb{P}(\textit{down firing})} = \frac{e^{\theta w_d} - 1}{1 - e^{-\theta w_u}} = \frac{w_d}{w_u}\, e^{\theta(w_u+w_d)/2}\left(1 + O(\theta^2(w_u + w_d)^2)\right), \tag{2}$$

so over successive cycles the count ratio satisfies $\log(N^+/N^-) \to \theta\,\bar{w} + \log(w_d/w_u)$ *to leading order, with* $\bar{w} = (w_u + w_d)/2$.

In words, two things decide which way a position fires more often. Drift tilts the race towards whichever trigger line the price is heading for, and asymmetric placement shortens the distance to one line and lengthens it to the other. The first term measures the drift tilt, the second the placement tilt, and at zero drift only placement survives.

Three properties carry the empirical design. The law is *exact* (the display's first equality involves no expansion). It is *numéraire-free*, since nothing in it references the quote asset, so it holds identically across pair classes; and it is *dimensionless*, since $\theta\bar{w}$ is invariant under any time change scaling drift and variance together, so the window convention used to estimate $(\hat{\mu}, \hat{\sigma}^2)$ affects noise but not the target. The two regressors separate the two mechanisms, since realised drift enters through $\theta\bar{w}$ with theoretical slope one, and placement asymmetry enters through $\log(w_d/w_u)$, also with slope one, surviving at zero drift. Throughout, placement means where the operator sets the range relative to the price; it is distinct from the displacement δ of GS, the price's offset within a given range. The proof is in Appendix A.1.

4.2. The swap-free value channel

For swap-free rebalances the value channel is fully determined by the V3 mint geometry.[3]

Proposition 2 (Swap-free dust credit)**.** *In GS notation, let a position of liquidity L on range $[s_a, s_b]$, with midpoint $\bar{s}$ and sqrt-price half-width $h = (s_b - s_a)/2$, be displaced by $\delta \in (0, h]$ and re-minted swap-free into the recentred range of equal width. The rebalance credits, exactly,*

$$\begin{aligned} d_y^+(\delta) &= \frac{L\,\delta\,(2\bar{s}+h+\delta)}{\bar{s}+h} && \text{(token1, up-move)}, \\ d_x^-(\delta) &= \frac{L\,\delta\,(2\bar{s}+h-\delta)}{s\,(\bar{s}+h)(\bar{s}+h-\delta)} && \text{(token0, down-move, } s=\bar{s}-\delta\text{)}. \end{aligned} \tag{3}$$

With the stablecoin at T_1, and writing $d^{\pm}(\delta)$ for the USD values of the credits, the per-event ratio at matched displacement is

$$\frac{d^+(\delta)}{d^-(\delta)} = \frac{(2\bar{s}+h+\delta)(\bar{s}+h-\delta)}{(\bar{s}-\delta)(2\bar{s}+h-\delta)}. \tag{4}$$

In words, an up-move leaves the withdrawn amounts token1-heavy relative to the recentred range, so token0 binds the new mint and the excess token1 is the dust credit; the down-move case mirrors it with the token roles exchanged.

The closed forms (3) coincide with the directional-convexity result of I. P.,[4] whose Proposition 3 derives the same swap-free credits independently and proves the USD inequality $d^+ > d^-$ on volatile–stable pairs. The additions made here are the exact ratio (4), the corner and reversal identities of Corollary 1, the evaluation at the trigger displacement, and the production-scale bound of Section 6.4.

Define the numéraire channel as $\kappa = \log(\bar{d}^{\,+}/\bar{d}^{\,-})$, the log ratio of mean per-event USD credits by direction. The sign structure of the swap-free channel is then exact.

Corollary 1 (Sign structure)**.** *At the range-exit corner $\delta = h$ the ratio equals s_b/s_a exactly; under token-order reversal (stablecoin at T_0) it equals s_a/s_b exactly. The swap-free value of κ at the exit corner is $\kappa = \log(s_b/s_a) = W \log(1.0001)/2$ for a range of W ticks, and κ inverts under token-order reversal while the frequency law (2) is unchanged.*

The corollary converts GS's negative control into an exact identification within the swap-free domain, since the numéraire term flips sign under reversal while the drift and placement terms do not. Two qualifications matter at production scale. First, under the trigger geometry of Section 3.2, rebalances fire at displacement $\delta \approx (1-2\varphi)h$, not at the corner; at $\varphi = 0.25$ the ratio (4) evaluates to about 0.74 of its corner value, and it is placement-sensitive there in a way it is not at the corner. Second, and decisively, the swap-free magnitudes are small on measured V9 geometry; the universe's mean half-widths run from 4 to 414 ticks, so the swap-free κ lies between 3×10^{-4} and 3×10^{-2}.

Both closed forms and both corner values are verified against a direct V3 mint computation to relative error below 2×10^{-13} (Appendix A.2). The swap-free channel is exact, verified, and quantitatively small.

4.3. The swap-corrected value channel

Production rebalances are swap-corrected. ME showed that the swap correction, referred to below as the corrective swap, collapses to a one-dimensional residual $\varepsilon \sim \nu_\varepsilon$, called the *execution residual* here to distinguish it from GS's geometric residual. The per-event leftover is proportional to $|\varepsilon|$ times a position-scale notional, and the sign of ε decides which token carries it. Conditioning on rebalance direction, which ME's unconditional treatment never required, the value channel decomposes as

$$\kappa \;=\; \underbrace{\log\frac{\mathbb{E}\,[\,\mathcal{N}^+]}{\mathbb{E}\,[\,\mathcal{N}^-]}}_{\kappa_{\rm geom}} \;+\; \underbrace{\log\frac{\mathbb{E}\,|\varepsilon^+|}{\mathbb{E}\,|\varepsilon^-|}}_{\kappa_{\rm exec}} \;+\; \underbrace{\kappa_{\rm val}}_{\text{leftover-side keying}} \;+\; \underbrace{\kappa_{\rm num}}_{\text{Proposition 2}}\;, \tag{5}$$

where $\mathcal{N}^{\pm}$ and $\varepsilon^{\pm}$ are the corrective-swap USD notionals and execution residuals conditioned on rebalance direction, $\varkappa_{\text{num}}$ is the swap-free term of Section 4.2 evaluated at the trigger displacement, and $\varkappa_{\text{val}}$ is defined as the remainder collecting the interaction of leftover-side keying with USD valuation, so that (5) holds as an identity. Every term on the right is measurable per event, because the corrective swap is recorded in the same transaction as the rebalance and can be joined to it by transaction hash. Section 6 reports the measurements.

5. The frequency channel measured

5.1. Joint regression on trigger-zone events

For each pool, direction is assigned per event by trigger zone (up if $u^* \geq 1 - \varphi$, down if $u^* \leq \varphi$, mid-range excluded). The per-event placement wedge $\log(w_d/w_u)$ is measured against the trigger lines of the newly minted range. Each pool then contributes one observation to the cross-section, namely its log count ratio, the drift regressor $\hat{\theta}\,\bar{w}$ (the pool's estimated drift-to-variance ratio times its mean barrier distance, both measured to the trigger lines), and the placement regressor $\overline{\log(w_d/w_u)}$ (the pool mean of the per-event wedge). The fit is

$$\log \frac{N^+}{N^-} \;=\; \beta_1\,\hat{\theta}\,\bar{w} \;+\; \beta_2\,\overline{\log(w_d/w_u)} \;+\; c, \tag{6}$$

with theory $\beta_1 = \beta_2 = 1$ and $c = 0$. Table 2 reports the result at the measured $\varphi = 0.25$ and at the mis-specified $\varphi = 0.20$.

Table 2. Joint fit (6) of the directional firing counts, per-pool cross-section. Theoretical values $\beta_1 = \beta_2 = 1$, $c = 0$.

φ	Class	n	β_1 (drift)	β_2 (placement)	c	R^2
0.25	Volatile–stable	19	1.27	1.06	0.03	0.53
0.25	Volatile–volatile	20	0.01	1.02	−0.06	0.58
0.20	Volatile–stable	19	0.01	1.38	0.00	0.71
0.20	Volatile–volatile	21	−0.07	0.72	0.10	0.42

At the measured trigger fraction, both slopes are consistent with their theoretical values on volatile–stable pairs and the placement slope with its theoretical value on volatile–volatile pairs, with intercepts indistinguishable from zero in both classes. Once drift and placement are measured, no further directional component is detectable in the firing counts.

The drift slope on volatile–volatile pairs is null rather than attenuated. Those pools' drift estimates are relative-price drifts extracted from noisier streams. The reading here is that volatile–volatile drift sits beneath the estimator's measurement floor rather than that the law fails; at this precision the two are indistinguishable, and no specification tried separates them.

The $\varphi = 0.20$ rows constitute the mis-specification test. Relabelling the $[0.20, 0.25)$ band as mid-range scrambles the counts and pushes both classes off theory, so the regression itself confirms the trigger fraction that the firing-position distribution identifies. Earlier fits against the range bounds instead of the trigger lines gave systematically attenuated slopes (0.76 and 0.64 on volatile–stable), which the trigger geometry explains as the affine mis-specification of the same law.

A replay check supplies the sampling yardstick. Generating Bernoulli firings from the exact law (2) at each pool's measured regressors and re-running the identical fit over 200 replications recovers slopes of 0.96 ± 0.16 (drift) and 1.01 ± 0.04 (placement) with intercepts centred on zero (± 0.02 volatile–stable, ± 0.03 volatile–volatile; both empirical intercepts sit within two of these spreads), so the pipeline is unbiased under the law's own data-generating process. Replayed cross-sections also reproduce the observed scatter; the per-pool correlation between simulated and empirical count ratios is 0.58 ± 0.02, in line with the fit's R^2 of 0.53 to 0.58, so the cross-sectional variance the fit leaves unexplained is at the sampling scale of the law itself.

Every empirical coefficient in Table 2 sits within two standard deviations of theory except the volatile–volatile drift slope, whose replay value is 1.02 ± 0.07 against the measured 0.01. A slope near zero is the errors-in-variables limit of a noise-dominated regressor, and the placement slope on the same pools sits at theory. The fifteen-standard-deviation discrepancy is therefore best read as regressor noise rather than law failure; the drift term on these pairs is left untested rather than confirmed.

5.2. The volatile–volatile excess explained

On the frequency side, ME's closing observation reduces to an up-firing excess on volatile–volatile pools whose median log exit-count ratio of +0.334 is largely unexplained by measured drift. Three candidate mechanisms were tested.

Asymmetric placement explains it. Volatile–volatile pools place ranges with the price systematically above the midpoint. The per-pool median placement wedge $\log(w_d/w_u)$ is 0.405 measured against the pre-rebalance price and 0.332 against the post-rebalance price, compared with 0.095 on volatile–stable pools. The class-level medians agree closely with the exit asymmetry (0.332 against 0.334), and pool by pool the placement wedge predicts the exit-count asymmetry with slope 0.81, correlation 0.65, and intercept 0.05 even before the trigger-line correction; after it, the placement slope is 1.02 (Table 2; Fig. 2, left). The nearer up-barrier produces more up-firings by first passage; no new mechanism is required.

Jump-driven exits do not. If one-sided heavy-tailed moves drove the excess, the overshoot beyond the crossed bound would be directionally asymmetric. It is not; on volatile–volatile pools the median overshoot is 0.250 of the previous half-width upward against 0.268 downward, with ninetieth percentiles 1.016 against 1.000; volatile–stable pools are equally symmetric.

Labelling structure does not. Since "up" means token0 appreciating against token1, an artefact of the V3 token ordering would flip sign between pools with WETH at T_0 and WETH at T_1. The excess is positive on both subsets (median +0.113 with WETH at T_0, +0.159 at T_1, +0.249 with no WETH), so the effect is structural, not notational.

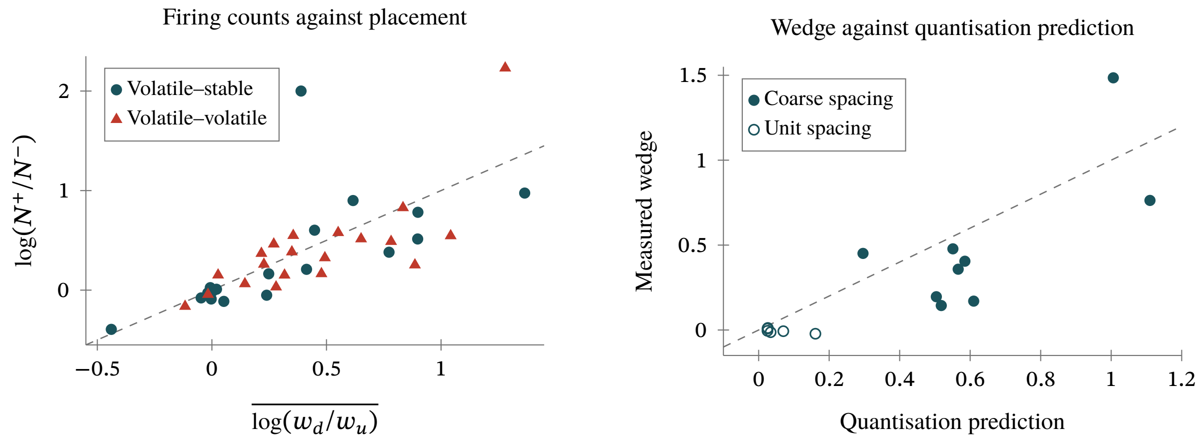

Fig. 2. Left: per-pool log firing-count ratio against the placement regressor. Right: measured placement wedge against the grid-quantisation prediction; six further pools, whose mean half-width is at most one spacing so the prediction is undefined, are absent. Dashed lines mark the identity.

The wedge originates in grid arithmetic rather than in strategy or market structure. The manager centres each new range by rounding the current tick down to the pool's tick-spacing grid, which drops the placed midpoint below the price by a uniform fraction of one spacing. The per-pool wedge matches this quantisation prediction with correlation 0.82 and slope 0.94, and the six pools on unit tick spacing, where the quantisation is negligible, carry no measurable wedge (median −0.006; Fig. 2, right); the class difference is composition, since the volatile–volatile universe is dominated by coarse-spacing pools with narrow ranges. The first-passage law then converts a rounding convention into the frequency excess that closed the companion paper, with slope one. The frequency channel is thereby accounted for within the resolution of the replay check.

6. The value channel measured

6.1. Per-event measurement via the corrective-swap join

Each rebalance's corrective swap appears in the pool's swap stream under the rebalance's transaction hash. Joining the two event streams yields, per event, the corrective notional $\mathcal{N}$ (USD value of the swap's input side), the implied execution residual proxy $|\hat{\varepsilon}| =$ credit USD$/\mathcal{N}$, and the leftover side (which token the credit lands on). The join covers 52.5% of the event lake (the swap stream's pool coverage), giving 31,340 events classified by trigger zone. The join mechanics validate directly; up-zone events sell token1 in 90.4% of cases on volatile–stable pairs and 84.3% on volatile–volatile, and down-zone events sell token0 in 94.9% and 92.5%, confirming both the direction classification and the surplus-side prediction of the mint geometry. The residual mismatch, largest on volatile–volatile up-zone events, is not characterised further here.

One null result removes an obvious alternative route. The 571 events the manager flags as swap-free,[3] which would have been the one subset where Proposition 2 applies exactly, turn out to credit wei-scale rounding residue (10^{-18} to 10^{-16} tokens). They mark already-at-ratio events, not geometric residuals, and production credits are in effect always post-swap leftovers.

6.2. The execution residual is direction-conditioned

ME's slippage law ν_ε was measured unconditionally, with binding probability $q = \mathbb{P}(\varepsilon < 0) =$ 0.26 on its substantial-swap sample. Conditioning on direction reveals structure. After up-moves the leftover lands on the bought token (token0) in 93.8% of volatile–stable and 90.0% of volatile–volatile events; after down-moves the split is nearly even (46.2% and 61.9% on token0). The unconditional q is the mixture of a strongly one-sided up-move regime ($q^+ \approx 0.06$ to 0.10) and a balanced down-move regime ($q^- \approx 0.5$); up-move corrections systematically overshoot towards the bought token and down-move corrections do not. This refines ME's assumption block, which treats ν_ε as independent of the event; ME's unconditional results are unaffected, but anything direction-resolved must carry the conditioning.

6.3. Two large factors nearly cancel

Table 3 reports the factor decomposition (5) at event level. The two leading factors are large and opposed. Up-moves carry larger relative execution error, with median $|\hat{\varepsilon}|$ of 0.122 against 0.085 on volatile–stable pairs, giving $\kappa_{\text{exec}} = +0.363$ at the median. Down-moves carry larger corrective notionals, giving $\kappa_{\text{geom}} = -0.334$; this is consistent with the placement asymmetry surfacing again, since the down trigger line sits farther from the price, so down-firings arrive with more travel, more imbalance, and a larger correction. Because the medians of the individual factors need not sum to the median of the total, the rows of Table 3 are not additive; the identity (5) holds at the level of the underlying means.

The residual is heavy-tailed at event level. On volatile–stable pairs the mean $|\hat{\varepsilon}|$ runs five to six times the median and the ninety-fifth percentile exceeds one, meaning the credit occasionally exceeds the corrective notional; volatile–volatile tails are milder, with means around twice the median. The

[3] Events recorded with `swap_mode` $= 0$ in the V9 rebalance event.

value wedge is consequently tail-driven, and value-weighted (sum-based) estimators dissociate from median-based ones, which explains the estimator inconsistencies that motivated this measurement.

Table 3. Value-channel factor decomposition, event-pooled by class. The $\varkappa$ rows are in log terms; medians unless stated.

	Volatile–stable		Volatile–volatile	
$\varkappa_{\text{exec}}$ (median)	+0.363		+0.191	
$\varkappa_{\text{geom}}$ (median)	−0.334		−0.214	
$\varkappa$ total (median)	+0.153		−0.084	
$\varkappa$ total (mean-based)	+0.044		+0.225	
	Up	Down	Up	Down
Median $\lvert\hat{\varepsilon}\rvert$	0.122	0.085	0.151	0.125
Leftover on token0 (%)	93.8	46.2	90.0	61.9

6.4. The numéraire component bounded

Corollary 1 predicts a sign flip of the value wedge under token-order reversal. It does not appear at any median-level estimator; pools with the stablecoin at T_1 show a median wedge of +0.168 and pools with the stablecoin at T_0 show +0.108, the same sign. The decomposition explains why the test has no power. On the measured V9 range geometry, the swap-free numéraire term at the trigger displacement lies between 3×10^{-4} and 3×10^{-2} across the universe, while the execution and geometry factors it rides on are of order 0.2 to 0.4 and are themselves direction-keyed. The numéraire mechanism is exact where GS proved it, and its sign structure is exact by Corollary 1, but at production range widths it is one to two orders subleading in the realised value wedge. The numéraire component is bounded rather than confirmed at production scale.

7. Discussion

7.1. What generalises and what is operator-specific

The quantities measured in this paper, the trigger fraction φ, the placement wedge, and the direction-conditioned execution residual, are all properties of one operator's configuration, and changing the operator changes each of them. What does not change is the structure. The firing law (2) is exact for any barrier placement. The trigger fraction of any manager whose proximity trigger is sharp enough to imprint a discontinuity is recoverable from its firing-position distribution the same way, with the regression's mis-specification sensitivity as a cross-check. And the decomposition (5) is defined by the mint geometry and the corrective-swap join, both of which exist on any V3-style manager that corrects in-transaction.

The companion papers' single-operator caveat transfers in sharpened form. The results are one operator's configuration, but the measurement recipe is operator-independent, and cross-operator replication reduces to re-measuring those three quantities and re-fitting two slopes whose theoretical values are one.

Two operational observations follow for the operator itself. A manager's placement policy sets its directional firing mix through the first-passage law, so a top-heavy placement convention yields systematically more up-firings, with whatever fee and gas consequences follow. The direction-conditioning of the execution residual, overshooting towards the bought token on up-moves but not on down-moves, is either a deliberate asymmetry or an unnoticed one; the data cannot distinguish intent, but it can now measure the behaviour.

7.2. Relation to the companion papers

Nothing here weakens GS's theorems. Proposition 2 extends the directional identity from position values to dust flows within the same swap-free domain and recovers GS's boundary magnitude Lw exactly at the exit corner. What changes is the production-scale attribution. The asymmetry ME observed at the close of its empirical study is real and holds in every estimator variant; its frequency component is now accounted for, and its value component is decomposed and bounded. Its carriers are placement and execution rather than the numéraire mechanism, and the volatile–volatile excess that appeared anomalous is the placement wedge converted through an exact law. For ME, the direction-conditioning of ν_ε (Section 6.2) is a refinement its independence assumptions can absorb unchanged at the unconditional level. The loss mechanisms of the LVR literature[7] remain an orthogonal cost layer throughout, exactly as both companion papers treat them.

7.3. Limitations

Five limitations bound the claims. All results are single-operator, as above, and the operator is the present author. The measurements are of the author's own production configuration; the paper therefore offers the recipe's portability, rather than the parameter values, for replication.

Two limitations concern valuation and coverage. USD valuation uses a static era price snapshot, so a cross term rides on price movement between event and snapshot, worst on volatile–volatile pairs. The corrective-swap join covers 52.5% of the event lake, limited by the swap stream's pool coverage.

Two further limitations concern the estimators. The execution-residual proxy conflates fee, price impact, and deliberate sizing bias; separating those needs the post-swap tick, which the stream contains but this paper does not exploit. The drift regressor has no power on volatile–volatile pairs at this estimator's precision (the replay check of Section 5.1 places the face-value volatile–volatile drift estimates fifteen standard deviations from the observed counts), so the drift slope's agreement with theory is established on volatile–stable pairs only.

7.4. Future work

Cross-operator replication is the immediate test; the preceding subsection reduces it to re-measuring three operator quantities and re-fitting two slopes.[4] Two estimator improvements would sharpen the present results on any dataset. Exploiting the post-swap tick would separate fee, price impact, and sizing bias within the execution residual, and a drift estimator with a lower noise floor would allow the drift slope to be tested on volatile–volatile pairs, closing the one null that the replay check attributes to the regressor. The swap-corrected derivation of the value channel's per-event law, explaining both the sum-based wedge and the median-level null from the direction-conditioned residual, remains the open analytical front.[5]

8. Conclusion

The directional asymmetry of production dust flows decomposes into a frequency channel and a value channel. The frequency channel obeys an exact, parameter-free first-passage law once the barriers are placed on the operator's trigger lines, whose position the data discloses. On the V9 dataset, drift and placement enter with slopes consistent with their theoretical value of one wherever

[4]This test has since been run. A companion paper[12] applies the trigger-fraction recovery unilaterally, from public chain data alone, to independently operated production managers, recovering per-pool fractions spanning 0.10 to 0.42 across 81 operators and 102 operator-pool cells, and pairs the recovery with a quantified jitter defence.

[5]This front has since closed. A companion paper[13] derives the per-event law by exact replay of the deployed pipeline, entering through the first estimator improvement above; the residual reduces to a race between the within-tick offset of the price and the corrective swap's own tick displacement, and the direction-conditioned binding probabilities follow with nothing fitted. The execution attribution of Section 6.3 accordingly refines to tick quantisation interacting with the sign of price impact.

the regressors carry signal, and intercepts sit at zero. The volatile–volatile excess that motivated the study is asymmetric range placement, itself a grid-quantisation artefact, converted through that law.

The value channel is decomposed and bounded. It is the near-cancellation of two large, direction-keyed execution factors, with the execution residual direction-conditioned. The numéraire mechanism of the prior theory is derived exactly for swap-free dust flows, sign inversion included, and bounded one to two orders below the execution factors at production range widths.

For the operator observed here, the asymmetry at production scale is carried by operator microstructure, with market drift entering only through the operator's trigger geometry; the numéraire mechanism does not carry it. The law and the measurement recipe are portable; running them on a second operator is the next test.

References

[1] 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.

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

[4] P., I. "Convexity and Topology in Cross-Pool Dust Redistribution: Structural Edge in Active CL Management." *SSRN preprint 6509938* (2026) https://ssrn.com/abstract=6509938.

[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] Agarwal, A., Gobet, E. "Optimal Exit from Uniswap v3 and Best Expected Return for a Liquidity Provider." *HAL working paper hal-05231633* (2025) https://hal.science/hal-05231633.

[10] Bergault, P., Bieber, S., Sánchez-Betancourt, L. "Optimal Exit Time for Liquidity Providers in Automated Market Makers." *arXiv:2509.06510* (2025) https://arxiv.org/abs/2509.06510.

[11] Redner, S. *A Guide to First-Passage Processes*. Cambridge University Press, Cambridge (2001).

[12] 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.

[13] Ryan, K.R. "Quantisation Microstructure: The Tick-Race Law of Rebalancing Dust." *SSRN preprint 7166739* (2026) https://ssrn.com/abstract=7166739.

A. Derivations

A.1. Proof of Proposition 1 (Directional firing law)

For $X_t = \mu t + \sigma B_t$ from 0 with absorption at $+w_u$ and $-w_d$, the scale function $S(x) = (1 - e^{-\theta x})/\theta$ with $\theta = 2\mu/\sigma^2$ gives

$$\mathbb{P}(\text{up first}) = \frac{S(0) - S(-w_d)}{S(w_u) - S(-w_d)} = \frac{e^{\theta w_d} - 1}{e^{\theta w_d} - e^{-\theta w_u}}, \tag{7}$$

and the complementary probability has numerator $1 - e^{-\theta w_u}$, giving the ratio in (2). Expanding the numerator and denominator to second order in θ yields $(w_d/w_u)(1+\theta(w_u+w_d)/2)+O(\theta^2)$, hence the stated leading-order form; at $\theta = 0$ the ratio is w_d/w_u by the standard gambler's-ruin computation. Under era-constant drift and variance, successive inter-rebalance cycles are independent in firing direction, since each rebalance re-places the range and the tick process is Markov. The empirical count ratio therefore converges to the probability ratio regardless of the direction-dependent cycle durations. Invariance under time change follows because θ is a ratio of drift to variance rates and w_u, w_d carry no time dimension.

A.2. *Proofs of Proposition 2 and Corollary 1*

In the up-move case, $s = \bar{s} + \delta$. The withdrawn amounts on $[s_a, s_b] = [\bar{s} - h, \bar{s} + h]$ are $x = L(1/s - 1/s_b) = L(h-\delta)/(s(\bar{s}+h))$ and $y = L(s - s_a) = L(\delta + h)$. The recentred range $[s-h, s+h]$ requires per-unit-L amounts $g_x = h/(s(s+h))$ and $g_y = h$. The up-move leaves the old position token1-heavy relative to the new requirement, so token0 binds, $L_{\text{new}} = x/g_x$, and the token1 leftover is

$$
\begin{aligned}
d_y^+ = y - \frac{x}{g_x}\,g_y = L(\delta+h) - L\,\frac{(h-\delta)(s+h)}{\bar{s}+h} \\
= \frac{L\left[(\delta+h)(\bar{s}+h) - (h-\delta)(\bar{s}+\delta+h)\right]}{\bar{s}+h} = \frac{L\,\delta\,(2\bar{s}+h+\delta)}{\bar{s}+h}. \quad (8)
\end{aligned}
$$

The down-move case is symmetric with token1 binding, giving the down-move form of (3). With the stablecoin at T_1, the token0 credit is valued at s^2 with $s = \bar{s} - \delta$, and the ratio (4) follows by division. At $\delta = h$ the up credit is $2Lh = Lw$, recovering GS's boundary magnitude, the down credit valued at the displaced price is $2Lh(\bar{s}-h)/(\bar{s}+h)$, and the ratio is $(\bar{s}+h)/(\bar{s}-h) = s_b/s_a$. Under token-order reversal the token1 credit is valued at $1/s^2$ and the same computation gives s_a/s_b at the corner. All closed forms, both corner values, and the reversal identity are verified against a direct evaluation of the V3 amount functions and the LP-min mint over a grid spanning $\bar{s} \in [0.5, 2000]$, $h/\bar{s} \in [0.005, 0.1]$, and $\delta/h \in [0.1, 1]$, with maximum relative error below 2×10^{-13}. The placement sensitivity away from the corner is quantified on the same grid; perturbing the recentred range by a 10% placement skew moves the log ratio by 0.196 at $\delta = 0.5h$ and by 0.002 at $\delta = 0.99h$.

B. Data and estimators

Lake and window. Rebalance and dust-credit events come from the PM V9 index, with the analysis window pinned at block 45,746,143 and a universe of pools with at least 100 rebalances across all depositors (48 pools; 126,339 events, of which 126,087 credit dust). Tick paths come from the era-matched V9 swap-event stream (6.3 million swaps, 44 pools); events are excluded when the nearest preceding swap is more than 450 blocks stale. The corrupted-print filter drops swap-to-swap tick changes exceeding 5,000; it never binds on range-midpoint deltas. Reconstruction hygiene follows ME; flow statistics are safe against the credits-only blindness that affects standing-balance walks, and no standing-balance statistic is used here.

Direction and trigger classification. Per event, the price position u^* within the previous range is computed from the last swap strictly before the rebalance block, avoiding contamination by the rebalance's own corrective swap. Events with $u^* \geq 1 - \varphi$ ($\leq \varphi$) are up (down) firings; mid-range events are excluded from frequency-law tests and reported as a separate population. Placement per event is measured on the newly minted range against the same pre-event price, with a post-rebalance variant (first swap after the block) used as a cross-check; the two agree to within the reported class differences. A block-tercile split of the window confirms the trigger fraction is stable; the firing-position discontinuity sits at $\varphi = 0.25$ in each tercile, with no step at any other fraction.

Drift and variance. $\hat{\mu}$ is the net tick displacement divided by elapsed time over the pool's rebalance era; $\hat{\sigma}^2$ is the variance of consecutive daily last-tick changes divided by the day length. Daily coarsening is essential; swap-scale quadratic variation is inflated by microstructure noise and attenuates the drift regressor towards zero.

Replay check. Synthetic firings are drawn per pool from the exact law (2) at the measured regressor pair $(\hat{\theta}\,\bar{w}, \overline{\log(w_d/w_u)})$, which jointly determine the up-firing probability with no further inputs, and the cross-section is re-fitted with the identical pipeline over 200 replications; the recovered slope distributions supply the spreads quoted in Section 5.1.

Value estimators. Credits are valued at an era snapshot of token prices; per-event medians by direction give the median-based wedges, with means reported alongside because the residual is heavy-tailed. The unit-free geometric estimator of Table 1 takes the geometric mean of per-token raw-unit credit ratios and requires no price data. The corrective-swap join aggregates signed swap amounts per transaction and pool; the notional is the USD value of the input side; the leftover side is the USD-majority credit token.

C. Notation

Hats denote estimates $(\hat{\mu}, \hat{\sigma}^2, \hat{\theta}, \hat{\varepsilon})$; overbars denote per-pool averages of per-event quantities where indicated; the exception is $\bar{s}$, GS's range midpoint. The trigger fraction φ is distinct from GS's position-value function $\phi(s, s_a, s_b)$; the two never co-occur in a display.

Table 4. Notation used throughout the paper.

Symbol	Meaning
Price process and trigger geometry.	
s, s_a, s_b	sqrt price and range bounds (V3 / GS)
$\bar{s}, h$	range midpoint and half-width $(s_b - s_a)/2$ in sqrt-price space
w	GS's full range width in sqrt-price space, $s_b - s_a = 2h$
δ	displacement from the range midpoint (GS)
μ, σ^2	tick drift and variance rates
θ	drift-to-variance ratio $2\mu/\sigma^2$ (tick units)
W	range width in ticks
φ	trigger fraction; the firing zone starts φW inside each bound
w_u, w_d	distances, in ticks, from the placement point to the upper/lower trigger lines
$\bar{w}$	per-pool mean barrier distance $(w_u + w_d)/2$ over trigger-zone events
ℓ, u	lower/upper range bounds in ticks (Section 3.2)
u^*	price position within the previous range at firing
Flows and channels.	
T_0, T_1	token slots of the V3 pair ordering (V3 / GS)
L	position liquidity (V3 / GS)
Φ^+, Φ^-	directional USD-valued dust flows
N^+, N^-	directional firing counts
$\bar{d}^{\pm}$	mean per-event USD credit by direction
$d^{\pm}_{y,x}(\delta)$	swap-free per-event dust credits, eq. (3)
κ	value channel, $\log(\bar{d}^+/\bar{d}^-)$; subscripts geom/exec/val/num per eq. (5)
$\mathcal{N}$	corrective-swap USD notional
$\varepsilon, \nu_\varepsilon$	execution residual and its law (ME); $q = \mathbb{P}(\varepsilon < 0)$, $q^{\pm}$ by direction
$\overline{\log(w_d/w_u)}$	pool mean of the per-event placement wedge over trigger-zone events, eq. (6)
$\lvert\hat{\varepsilon}\rvert$	execution residual proxy, credit USD over corrective notional $\mathcal{N}$