# The Hidden Microstructure of Shared Balance Concentrated Liquidity

## A Master Equation for the Dust Ledger and Propagation of Chaos

K.R. Ryan

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

May 2026

**Keywords** geometric residual · interacting particle systems · multi-pool rebalancing · dust ledger Uniswap V3 · DeFi mechanism design

**JEL:** C63, C65, D47, G14

∴

**Abstract.** The geometric residual identified in the prior paper[1] takes a richer form once a concentrated liquidity manager retains it in a depositor-keyed dust ledger shared across multiple pools. The residual is no longer isolated to its own rebalance, but becomes a hidden state coupling the portfolio's pools. This paper derives the closed-form law governing that state. Theorem 1 (the Master Equation) gives the exact per-event jump law on any pool, induced by the V3 mint geometry at rebalance, and recovers the prior result in the single-pool swap-free corner. Proposition 1 promotes the prior paper's Connector Rule conjecture to a sign-keyed correlation under explicit V3-mechanical assumptions, while Theorem 2 generalises zero-swap extinction to a per-token mass-conservation identity across multi-pool rebalances, with a donor/absorber partition under heterogeneous arrival. On the homogeneous hub-spoke topology, Proposition 2 establishes a conditional propagation-of-chaos limit under uniform-in-$N$ hub mixing, and Proposition 3 identifies the limiting one-spoke marginal as a binary atom-mixture whose mass and shape are determined by the slippage law of the in-pool swap correction. The half-normal mixture obtained under symmetric Gaussian slippage is a narrow corollary. The empirical anchor spans roughly 34,500 rebalance events over the V6 through V9 evolution of a PM contract family on Aerodrome Slipstream (Base) under one operator, with per-event predictions matching on-chain dust credits at median ratio 1.0000 throughout. Contract-level verification reproduces the closed form against unmodified Aerodrome Slipstream and Uniswap V3 infrastructure to bit-exact tolerance. The asymptotic factorisation and atom-mixture form are empirically consistent with V9-realistic asymmetric heavy-tailed slippage at the kernel level; the heterogeneous per-spoke V9 marginal is left as an open extension problem, and cross-operator replication remains open.

## 1. Introduction

A concentrated liquidity (CL) position holds two tokens within a sqrt-price range and must be rebalanced (withdrawn, optionally swap-corrected, and re-minted into a new range) whenever the price drifts outside that range. The V3 mint geometry fixes the input ratio at any sqrt-price inside the new range, so the supplied amounts almost never coincide exactly with what the new mint requires. Each rebalance therefore leaves a leftover on one side of the pair. On a stock `NonfungiblePositionManager` (NFPM), that leftover is simply returned to the depositor and disappears from the subsequent state of the manager. The architecture studied here departs from that default by retaining it in a depositor-keyed mapping shared across the depositor's positions, from which later mints may draw as standing input. It is this mapping that will be called the *dust ledger*.

The prior paper[1] isolated the single-pool, swap-free leftover as a closed-form per-event residual. What matters here is that, once the leftover is retained and recycled across multiple pools, it is no longer merely the by-product of an individual rebalance. It becomes a shared state variable on the portfolio, coupling one position's residual to another position's mint. The question is then no longer only what residual a given rebalance leaves behind, but what law governs the evolution of the standing balance once the architecture itself recycles that balance across the portfolio.

The empirical instance is a Position Manager built on Aerodrome Slipstream (Base), which evolved through four independently deployed versions (PM V6 through V9) over a 23-day window in March 2026. All four implement the same depositor-keyed shared dust ledger, and the per-event closed form of the prior paper holds on every one across roughly 34,500 rebalance events. That persistence across four reimplementations is what makes it natural to analyse the dust ledger as a system rather than as a curiosity of one contract, though the four contracts still share a single operator; cross-operator replication of the framework developed here remains open, with the underlying dust-circulation mechanism itself observed in independent operator data.[2]

Three questions arise in that setting. What is the per-event update law on a multi-pool dust ledger, once standing balance and contract-level swap correction are admitted? What conservation structure does such a ledger impose on a portfolio of positions through time? And what asymptotic regime emerges as the portfolio scales? The present paper answers those questions in closed form. Theorem 1 gives the per-event jump on any pool, including standing balance and swap correction, and is validated event-by-event against on-chain data. Proposition 1 promotes the prior paper's Connector Rule conjecture to a sign-keyed correlation under explicit V3-mechanical assumptions. Theorem 2 generalises the zero-swap extinction result to a per-token mass-conservation identity across multi-pool rebalances, together with a donor/absorber partition under heterogeneous arrival. On the homogeneous hub-spoke topology, Proposition 2 establishes that under uniform-in-$N$ hub mixing the stationary $m$-spoke marginals factorise asymptotically; this is the propagation-of-chaos result, and the factorisation diagnostic is consistent across slippage-law families in the empirical study of Section 7.3. Building on 2A, Proposition 3 characterises the limiting one-spoke marginal as a binary mixture of an atom at zero and a slack-side density, with mass and shape determined by the slippage law and valid under arbitrary $\nu_\varepsilon$. The half-normal mixture obtained when the slippage is symmetric mean-zero Gaussian is a narrow corollary, applicable in idealised positions but not in the full V9 production regime; the per-spoke V9 marginal variation lies outside this homogeneous closure and is named as an open extension problem. Fig. 1 summarises the resulting hierarchy.

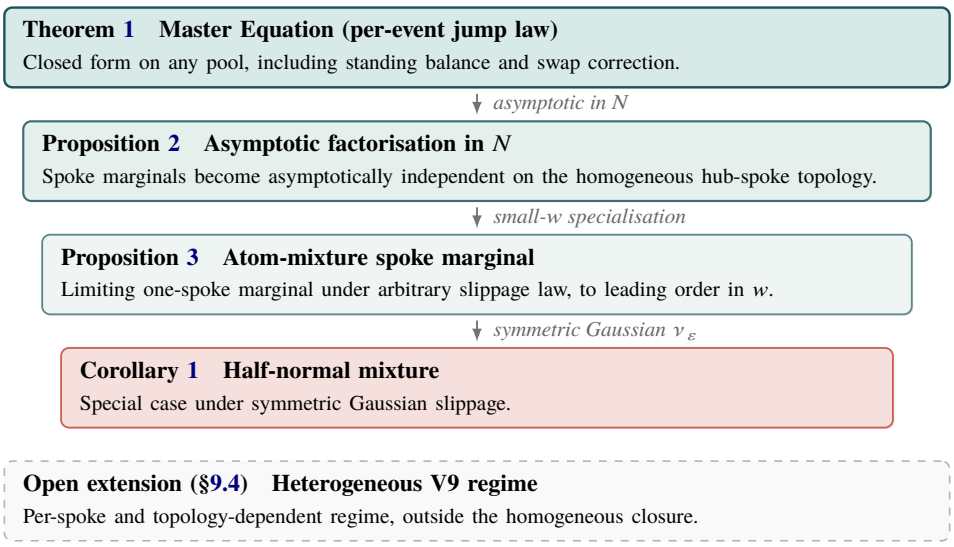


Fig. 1. Hierarchy of dust-ledger laws. Each layer specialises its predecessor; the dashed open extension lies outside the homogeneous symmetric closure (§9.4).

## 2. Related work

The geometric residual was introduced in the prior paper,[1] to which the present one is directly indebted. In particular, this paper extends that work's Theorem 1 (vanishing condition), Theorem 2 (single-pool convergence), Theorem 3 (zero-swap extinction), and the architectural-precondition argument of §7.1. The empirical hub-spoke pattern reported in §5.4 of that paper (Group B) is, in retrospect, the simplest non-trivial special case of the framework developed here.

*Convexity and Topology in Cross-Pool Dust Redistribution*[2] (2026) provides a graph-theoretic characterisation of cross-pool dust circulation, including a degree-based residence-time identity (its Proposition 2) and a directional convexity scaling (its Proposition 3). Projecting the bipartite graph $G$ onto its token vertices recovers that work's portfolio graph, and its Proposition 2 corresponds to the degree-versus-velocity scaling that reappears in the hub-spoke analysis of Section 6.1 below.

The closest existing formal frameworks for V3 LP value loss concern costs incurred while a position is open. Cartea *et al.*[4] introduce *predictable loss* (PL), a measure of the unhedgeable loss borne by strategic LPs, combining the convexity of the LP's position-value function with the opportunity cost of locked capital. Milionis *et al.*[5] introduce *loss-versus-rebalancing* (LVR) for general CFMMs, defined as the return gap between an AMM LP and a continuously rebalancing benchmark under arbitrage against stale prices. Both frameworks concern losses accumulating between rebalances. The present paper concerns a different object, the per-rebalance redistribution of value through a depositor-keyed dust ledger shared across multiple pools. The contrast is not merely verbal. The geometric residual is zero on same range rebalances regardless of price travel (Theorem 1 of the prior paper), whereas LVR is generally non-zero whenever the AMM price drifts. A multi-pool manager subject to the Master Equation is therefore also subject to PL and LVR between rebalances, but the V6 through V9 empirical anchor of Section 7.1 concerns only the jump at the rebalance itself.

Empirical work on V3 LP returns has largely focused on position-level risk and return distributions. Heimbach *et al.*[6] derive an analytical impermanent loss expression for V3 and show empirically that LP outcomes vary materially across strategies, with active management playing a substantial role. The roughly 34,500-event anchor used here extends that empirical line from aggregate position outcomes to the event-level dynamics of a shared dust ledger.

The propagation-of-chaos argument of Section 6 belongs to a classical line of work. Sznitman[7] reviews the foundational theory of mean-field limits for interacting particle systems, and uses the term *master equations* in the sense of the forward Markov equations for $N$-particle systems. The homogeneous hub-spoke chain studied here is a discrete-event instance of that general setting. Closer in spirit is Graham and Méléard,[8] who prove propagation of chaos in variation norm on the laws of sample paths for a fully connected loss network with alternate routing. The correspondence is not exact, but it is more than superficial; their alternate-routing structure is the analogue of cross-pool absorption in the dust ledger, and both processes are non-exchangeable interacting jump systems coupling through shared state.

The modern survey of the propagation-of-chaos literature is Chaintron and Diez.[9] Their discussion of mean-field jump processes and piecewise deterministic Markov processes (§2.2.3), together with their treatment of interaction-graph methods (§4.5.2), provides the most directly relevant technical backdrop for Section 6. In their sense, the dust ledger is a piecewise deterministic Markov process, constant between rebalance events with jumps at exogenous Poisson times. The bipartite-graph stochastic-flow framing of Section 3 also has an evident family resemblance to Kelly's closed-network framework.[11] Chapters 2–3 of Kelly develop closed migration processes and closed queueing networks; the per-token mass conservation of Theorem 2 is analogous to closed-network conservation of total customer count. The analogy is nevertheless limited. Rebalance arrivals here are exogenous rather than state-dependent, and the dust-creation kernel has no Kelly-style reversibility.

Mean-field game theory provides a separate, though related, lineage. Propagation-of-chaos arguments are now standard in continuous-time finance models of portfolio choice and equilibrium interaction, as in Carmona and Delarue,[12] while Cardaliaguet, Delarue, Lasry, and Lions[13] develop a different *master equation* framework on the space of probability measures, arising as the asymptotic-

$N$ limit of $N$-player Nash systems. That literature is relevant chiefly for orientation. The present paper's usage of *master equation* sits much closer to the $N$-particle Markov sense reviewed by Sznitman (Chapter 0) than to the second-order PDE on Wasserstein space studied in mean field games. To the author's knowledge, the discrete-event jump-kernel framing developed here has not previously been applied to AMM dust-ledger dynamics in the DeFi literature.

The underlying V3 mechanics are Adams *et al.*;[3] the V3 amount equations of Appendix B follow their equations 6.29–6.30.

## 3. The system

The notation follows the prior paper,[1] extending scalar quantities to vector-valued, time-indexed analogues; a glossary is in Appendix A. USD-value statements are evaluated at fixed reference token prices $(p_1, \dots, p_n)$. The framework abstracts from fees and MEV; per-event LP fees, gas, and externally induced rebalance MEV are assumed not to enter the dust ledger.

**Portfolio.** A portfolio consists of $m$ concentrated liquidity positions held by a single depositor on a single PM contract, across $n$ tokens. It induces a bipartite graph $G$ between token nodes $\{T_1, \dots, T_n\}$ and pool nodes $\{P_1, \dots, P_m\}$, with an edge between $T_i$ and $P_k$ whenever token $T_i$ belongs to pool $P_k$. Pool nodes therefore have degree 2, while the token-side degree $\deg(T_i)$ counts the pools containing $T_i$. For each pool $P_k$ with token pair $(T_a, T_b)$, we follow the V3 ordering convention throughout, with $T_a$ occupying the $T_0$ position and $T_b$ the $T_1$ position.

**Dust ledger.** The dust ledger is the standing-balance vector $D_{\text{pool}}(t) \in \mathbb{R}^n_{\geq 0}$, with one component per portfolio token. It is piecewise constant between rebalance events, with a deterministic jump whenever a pool containing that token rebalances. For a pool $P_k$ with tokens $(T_a, T_b)$, the pool-local snapshot is $(x_{\text{dust}}, y_{\text{dust}}) = (D_{\text{pool},T_a}(t), D_{\text{pool},T_b}(t))$.

**Rebalance events.** Rebalance events arrive at random times $\tau_1, \tau_2, \dots$, exogenous to the dust state. Each event records the pool $P_k$, the old range $[s_a, s_b]$, the new range $[s'_a, s'_b]$, the sqrt-price $s(\tau)$, the pre-event position value $V(\tau^-)$, and the signed pool-frame swap correction $(\sigma_x, \sigma_y) \in \mathbb{R}^2$ defined in Section 4.2.

## 4. The Master Equation

This section states the paper's central object, Theorem 1, giving the closed-form per-event jump on any pool. Sections 5 and 6 then consider its multi-pool consequences and asymptotic limit.

### *4.1. Per-event jump*

The per-event update of the dust ledger is the closed-form solution of a one-dimensional linear programme induced by the V3 mint geometry.

**Theorem 1** (Master Equation jump)**.** *Let $P_k$ be a pool with tokens $T_a$ at the $T_0$ position and $T_b$ at the $T_1$ position, sqrt-price $s$, old range $[s_a, s_b]$, and new range $[s'_a, s'_b]$, with $s \in (s'_a, s'_b)$, i.e. in the interior of the new range. Let the position's old liquidity be $L$, with token amounts*

$$x_w = L\, a_x(s, s_a, s_b), \qquad y_w = L\, a_y(s, s_a, s_b), \tag{1}$$

*where $a_x, a_y$ are the V3 amount functions (Appendix B). Let the standing balance be $(x_{dust}, y_{dust})$ and the signed swap correction be $(\sigma_x, \sigma_y)$, so that the mint inputs are*

$$\hat{x} = x_w + x_{dust} - \sigma_x, \qquad \hat{y} = y_w + y_{dust} - \sigma_y. \tag{2}$$

*Write $g_x = a_x(s, s'_a, s'_b)$ and $g_y = a_y(s, s'_a, s'_b)$ for the per-unit-L token amounts at the new range. Then the new liquidity is*

$$L_{new} = \min\left(\frac{\hat{x}}{g_x}, \frac{\hat{y}}{g_y}\right), \tag{3}$$

*and the new dust-ledger entries for $T_a, T_b$ are*

$$x'_{dust} = \hat{x} - L_{new}\, g_x, \qquad y'_{dust} = \hat{y} - L_{new}\, g_y, \tag{4}$$

*both non-negative. In the swap-free regime ($\sigma_x = \sigma_y = 0$), at least one of $x'_{dust}, y'_{dust}$ is zero by the LP-min structure, and exactly one is zero off the measure-zero ratio-preserving locus $\{\hat{x}/\hat{y} = g_x/g_y\}$ in input-ratio space; the non-zero leftover is $|\Delta R|$ in the notation of §3.2 of the prior paper,[1] namely the surplus on the slack-side token.*

*Proof sketch.* The mint maximises $L_{\text{new}}$ subject to the two non-negativity constraints $L_{\text{new}}\, a_x(s, s'_a, s'_b) \le \hat{x}$ and $L_{\text{new}}\, a_y(s, s'_a, s'_b) \le \hat{y}$. Since the objective is increasing in $L_{\text{new}}$, the optimum lies where one constraint binds, or where both bind simultaneously. If both bind, the leftover is zero on both tokens; this is the ratio-preserving locus of Theorem 1 of the prior paper, of measure zero on the displacement axis. Otherwise exactly one constraint binds, and the surplus on the slack side appears as dust. The full derivation, including the swap-corrected case, is given in Appendix C. ■

Each rebalance withdraws the old position, recycles the pool's standing dust into the mint inputs, optionally applies a swap correction (Section 4.2), and credits the LP leftover to the dust mapping for the pool's two tokens. Standing balances for tokens outside the pool are unchanged:

$$D_{\text{pool},i}(\tau_N^+) = \begin{cases} x'_{\text{dust}} & \text{if } T_i = T_a, \\ y'_{\text{dust}} & \text{if } T_i = T_b, \\ D_{\text{pool},i}(\tau_N^-) & \text{otherwise.} \end{cases} \tag{5}$$

All terms in (5) are observable per event from on-chain data and contract state. The single-pool swap-free corner ($m = 1$, $\sigma_x = \sigma_y = 0$) recovers Theorems 1 and 3 of the prior paper. As a worked example, the canonical Foundry test of that paper's §6 (1 WETH plus 2,500 USDC minted at tick 73,135, displaced by +200 ticks, and rebalanced into a ±1,000-tick range) yields a 2,061 USDC residual on the bound side from (4) alone, matching the value reported there. Fig. 2 summarises the jump structure.

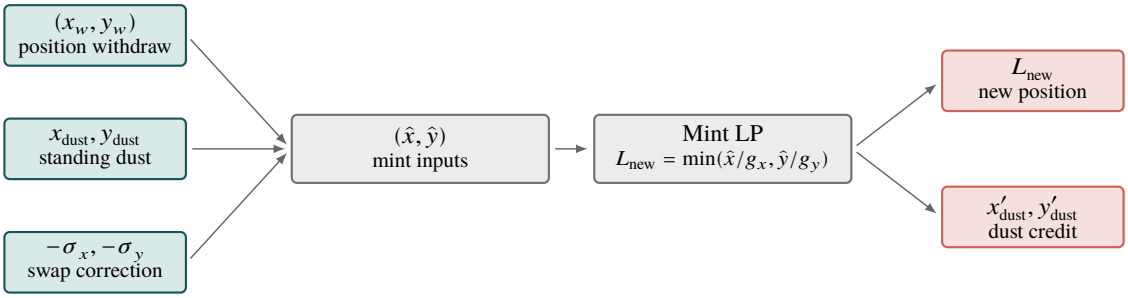


Fig. 2. Per-event jump of the Master Equation.

### 4.2. Swap correction

PM V9 supports an atomic in-pool swap[1] as part of the rebalance transaction, executed when needed to align the supplied tokens with the ratio required by the new range. The swap routes through the rebalance pool itself and uses the contract's own balance, exchanging part of the over-supplied token for the under-supplied one. This is the modal production configuration, with an in-pool swap firing on 74–96% of V9 rebalances across the active portfolios examined (Appendix G).

On chain, the swap appears as a Pool.Swap event in the same transaction as the rebalance, with signed amount fields $(\alpha_0, \alpha_1)$ in the pool's frame. In the contract's frame, $(\sigma_x, \sigma_y) = (\alpha_0, \alpha_1)$ enter the mint-inputs equation (2) directly. Theorem 1 covers the swap-free and swap-present cases in a single form, with $\hat{x}, \hat{y}$ shifting linearly by $(-\sigma_x, -\sigma_y)$ and the LP structure unchanged. Under in-pool routing the closed form uses the pre-swap tick for the withdraw amounts and the post-swap tick for the geometry of the new range (Appendix G).

## 5. Multi-pool consequences

Beyond the per-event jump, the Master Equation has two consequences for multi-pool portfolios, summarised alongside the asymptotic limit in Table 1.

Table 1. Stages of the Master Equation framework.

| Stage | Mechanism | Result | Empirical anchor |
|---|---|---|---|
| Per-event | LP solve at the V3 mint | Theorem 1 | §7.1, §7.2, §7.2 |
| Cross-pool | Connector-side correlation | Proposition 1 | §7.3 |
| Portfolio | Donor/absorber under $S = 0$ | Theorem 2 | §7.3 |
| Asymptotic | Factorisation in $N$ | Proposition 2 | §7.3 |
| Asymptotic | Atom-mixture spoke marginal | Proposition 3 | §6.3, §7.3 |
| | (half-normal mixture under | Corollary 1 | §6.3 |
| | symmetric Gaussian) | | |

### 5.1. Connector Rule

The Connector Rule of §3.5 of the prior paper[1] conjectures that per-rebalance dust-flow direction is dominated by the higher-connectivity token

$$T^* = \arg\max\{\deg(T_a), \deg(T_b)\},$$

supported by 10 per-pool Spearman correlations in $[-0.98, -0.56]$. The Master Equation makes the mechanism explicit. By (4), the binding side at each rebalance is determined by the supply ratio, and the displacement of a connector shared across pools contains a common component that co-moves with the slack-side dust. The sign of that co-movement depends on which side of the pool the connector occupies (positive at $T_0$, negative at $T_1$), a prediction the prior paper's sample of $T_1$-only connectors could not test.

**Proposition 1** (Connector Rule under V3 mechanics)**.** *Let $G$ be a portfolio bipartite graph with a multi-pool connector token $T^*$ such that $\deg(T^*) \geq 2$ and $\deg(T^*) > \deg(T)$ for every token $T \neq T^*$ appearing in any pool containing the connector. Assume:*

*(i) the connector token's marginal sqrt-price evolution against the rest of the portfolio is stationary over the rebalance arrival window;*

[1] Position Manager configuration `AUTO_OPTIMAL`.

*(ii) the per-pool sqrt-price displacement at the rebalance moment decomposes additively into a shared connector-token component and a pool-specific noise term;*

*(iii) the mint-time sqrt-price lies in the interior of the new range with probability one.*

*Define the* signed connector-side dust quantity *at a rebalance as the dust credited to the non-connector side minus the dust credited to the connector's side, so that it is positive when the non-connector side leaks. Under (i)–(iii), across rebalances of pools containing $T^*$, the per-pool Spearman rank correlation between signed sqrt-price displacement and signed connector-side dust quantity has a deterministic sign keyed to the pool position of $T^*$: positive when $T^*$ occupies the $T_0$ position, and negative when it occupies the $T_1$ position. Its magnitude is governed by the variance share of the connector-shared component in the displacement decomposition; under the homogeneity assumption in (ii), that share gives a lower bound.*

Proof in Appendix D; V9 verification in Section 7.3.

### *5.2. Donor/absorber decomposition*

Under shared depositor balances, residuals left by one pool may be absorbed into the mint inputs of another. Theorem 3 of the prior paper[1] shows that a single position under $S = 0$ decays geometrically. Once depositor-keyed dust is shared across multiple pools, that strict extinction is replaced by a per-token mass-conservation identity holding across rebalances at the contemporaneous sqrt-price.

**Theorem 2** (Multi-pool mass conservation under $S = 0$)**.** *Let a portfolio have $m \geq 2$ positions sharing a depositor-keyed dust mapping, and let $a_{k,i}$ denote position $k$'s implied amount of token $T_i$ at sqrt-price $s$ (equal to $L^{(k)} a_x(s, s_a^{(k)}, s_b^{(k)})$ when $T_i = T_0$ in pool $P_k$, analogously for $T_1$, and zero when $T_i \notin P_k$). Under $S = 0$ rebalances, and in the absence of LP fees and external token transfers, the per-token total*

$$M_i(t) \;:=\; \sum_k a_{k,i}(s(\tau)) \;+\; D_{pool,i}(t) \tag{6}$$

*is invariant across rebalance events, evaluated at the contemporaneous sqrt-price $s(\tau)$ of each event, modulo integer tick-rounding remainders. The fixed-price USD aggregate*

$$V_{total}(t) \;:=\; \sum_k \sum_i a_{k,i}(s(\tau))\, p_i \;+\; \sum_i D_{pool,i}(t)\, p_i$$

*inherits the same event-level invariance. The single-pool corner $m = 1$ recovers the swap-free single-pool regime of the prior paper. Under $S = 0$, the lone position decays geometrically by Theorem 3 there; under non-zero S, the $m = 1$ corner is the convergence-in-expectation regime of Theorem 2 there.*

**Remark (Donor/absorber decomposition).** Let the per-position USD-value change at a rebalance of pool $P_k$ be

$$\Delta V_k = L_{\text{new}}^{(k)}\, v_k^{\text{new}}(s) - L^{(k)}\, v_k^{\text{old}}(s),$$

with $v_k^{\text{old}}$ and $v_k^{\text{new}}$ the per-unit-$L$ position values at sqrt-price $s$ over the old and new ranges respectively. Cumulating $\Delta V_k$ over a window gives the position's net cumulative dust impact; a position with negative cumulative impact is a *net donor* over the window, and one with positive cumulative impact a *net absorber*. By (6), the sum of net impacts across positions plus the change in $\sum_i D_{\text{pool},i} p_i$ is zero over any window. A non-trivial donor/absorber partition (at least one strictly positive and at

least one strictly negative cumulative impact) follows under heterogeneous rebalance arrival rates and ranges across positions, though it is not forced by (6) alone. Lemma 1 (Appendix E) gives a sufficient condition, and the V9 evidence of Section 7.3 verifies the partition empirically in eight of eight multi-pool portfolios.

Proof in Appendix E; V9 verification in Section 7.3.

## 6. Propagation of chaos

**Scope.** The Master Equation of Section 4 (Theorem 1) and the multi-pool consequences of Section 5 (Proposition 1, Theorem 2) are topology-agnostic: they apply to any portfolio graph formed by the depositor's positions, including triangles, double hubs, ladders, and the mixed topologies observed in the V9 dataset of Section 7. The closed-form per-event jump, together with the mass-conservation and sign-keyed-correlation results above, does not depend on any particular wiring of the portfolio. This section turns to a stronger but topology-specific statement. In the homogeneous hub-spoke case, the stationary spoke marginals of the dust ledger admit a closed-form asymptotic law as $N$ grows, via propagation of chaos. This is a distributional refinement rather than a precondition. The empirical fits on V9 portfolios (Section 7) and the Foundry verification on cross-pool topologies (Section 8) rest on the topology-agnostic Master Equation; the present theorem concerns the additional structure that appears when the portfolio graph is itself a star.

The framework invoked here belongs to the classical propagation-of-chaos line of Sznitman[7] and its mean-field descendants.[12,13] Its more familiar applications in finance concern continuous-time interacting-agent models with latent aggregate state, where the mean-field coupling is part of the modelling apparatus and must be justified at scale. The setting here is rather different. The aggregate state (the dust ledger $D_{\text{pool}}(t)$) is the contract's own depositor-keyed dust mapping; the coupling is supplied by the architecture rather than posited by the model; and the per-event jump kernel is the closed-form Theorem 1 rather than a fitted mechanism. In that sense, the kernel is exact and on-chain observable at every event. Only the asymptotic-$N$ decoupling is a limit statement.

Under the swap routing of Section 4.2, the perfect-routing target preserves position liquidity $L$ exactly. Under slippage, $L$ drifts on a slower scale, and the per-event change in $L$ is mean-zero conditional on the pre-event state. The remainder of this section works under the following assumption block.

**Assumptions.**

(A1′) Slippage is i.i.d. from a law $\nu_\varepsilon$ on $[-K_{\min}, K_{\max}]$ with $K_{\min} < 1$, $K_{\max} < \infty$, finite second moment $\sigma^2_{\text{slip}} := \text{Var}(\varepsilon)$, and density bounded away from zero on a neighbourhood of the origin. The Pool.Swap routing conservation $\sigma_y^{\text{raw}} = -s^2\,\sigma_x^{\text{raw}}$ collapses the swap-correction noise to this single one-dimensional residual; no symmetry, zero-mean, or Gaussianity is assumed.

(A2) Boundary trigger $\delta \sim \text{Unif}(\{\pm w/2\})$ at each event, independent of past state and slippage.

(A3) Homogeneous exchangeable spokes: parameters $(L, w, \lambda)$ and $\nu_\varepsilon$ shared. The heterogeneous case is treated as Conjecture 1.

(A4) There exist $C, \rho \in (0,1)$ independent of $N$ such that the hub-marginal chain in stationary spoke environment mixes in total variation at rate $C\rho^n$ after $n$ hub events.

(A5) The kernel induced by Theorem 1 is TV-Lipschitz in $(H, B_k)$ with constant $\Lambda = K_{\max}(1 + s_h^2)/2$, with $s_h^2$ the swap-routing slope in the human-scaled token frame.

Of these, (A1′)–(A3) and (A5) are model definitions or directly verifiable from (7) below; (A4) is the load-bearing limit-theory assumption. The chain on $D_{\text{pool}}(t)$ admits a unique stationary law $\pi$ and is positive recurrent under (A1′); rigorous Foster-function constructions are sketched in Appendix F. The V9 evidence of Section 7.3 is consistent with these assumptions at production scale.

### 6.1. *Hub-spoke topology and statement of the limit*

Consider the canonical hub-spoke topology, with one degree-$N$ hub token $T_0$ shared across $N$ pools and each pool contributing one degree-1 spoke token $T_k$ (Fig. 3). Write $H := D_{\text{pool},T_0}$ and $B_k := D_{\text{pool},T_k}$ for the hub and spoke-$k$ components of the dust ledger throughout this section and Appendix F. The hub is updated at rate $N\lambda$ and each spoke at rate $\lambda$, so the inter-event-rate ratio is $\tau_{\text{hub}}/\tau_{\text{spoke}} = 1/N$. This rate-ratio identity is consistent with Proposition 2 of[2] (residence time inversely proportional to degree) on homogeneous-rate hub-spoke graphs; the Group B hub-spoke pattern of[1] §5.4 is the same special case at finite $N$.

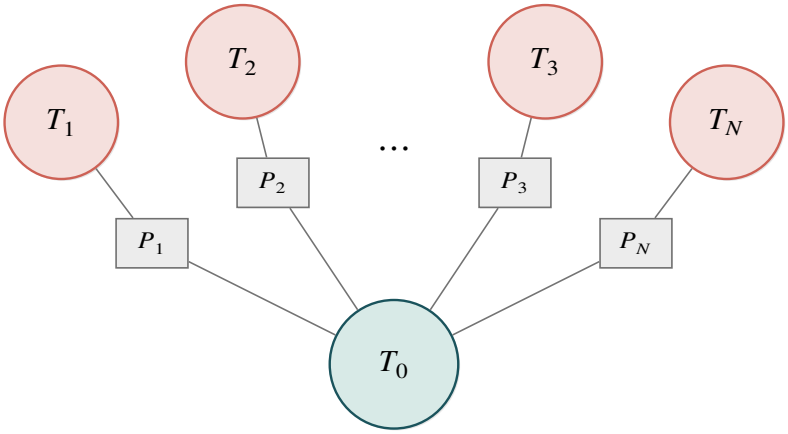


Fig. 3. Canonical hub-spoke topology with $N$ spokes.

In this topology the hub is touched by every rebalance while each spoke is touched only by its own. As $N$ grows, the hub mixes faster than any single spoke sees individual updates, and the marginal kernel governing each spoke decouples from the others; this is the standard mean-field setup. The kernel induced by Theorem 1 on this topology is, to leading order in $w$,

$$\widetilde{\mathcal{K}}_w((H,B_k),\cdot) = \begin{cases} B_k' = 0,\ H' = |Lw + B_k - H|\,|\varepsilon|\,(1+s_h^2)/2 & \text{if } (Lw + B_k - H)\,\varepsilon < 0, \\ H' = 0,\ B_k' = |Lw + B_k - H|\,|\varepsilon|\,(1+s_h^2)/2 & \text{otherwise.} \end{cases} \tag{7}$$

The factor $s_h^2$ is the swap-routing slope in the human-scaled token frame, equal to 1 to leading order in the homogeneous symmetric regime. The asymptotic-$N$ limit of this kernel admits two complementary statements: a structural factorisation result (Proposition 2) and a closed-form one-spoke marginal (Proposition 3), with the half-normal mixture obtained for symmetric Gaussian slippage as a corollary.

**Proposition 2** (Conditional propagation of chaos)**.** *Under (A1′)–(A5), let $\pi^N$ denote the unique stationary law of $X^N = (H, B_1, \dots, B_N)$ on $\mathbb{R}_{\geq 0}^{N+1}$, and $\pi_{(m)}^N$ its $m$-spoke marginal for any fixed $m \geq 1$. Then*

$$\pi_{(m)}^N \xrightarrow[N\to\infty]{w} \mu^{\otimes m},$$

*where $\mu$ is the stationary law of the limiting one-spoke chain $\bar{B}$ driven by i.i.d. hub samples $\bar{H} \sim \pi_H$, and $(\mu, \pi_H)$ is the unique fixed point of the McKean–Vlasov system associated with the kernel (7).*

**Proposition 3** (Closed-form one-spoke marginal)**.** *Under (A1′)–(A5) and the small-$w$ regime in which typical $|H|, |B| \ll Lw$ in stationarity (which holds whenever $\sigma_{slip} \ll 1$, since $H, B = O(Lw\,\sigma_{slip})$), the limit $\mu$ from Proposition 2 admits the binary atom-mixture form*

$$\mu(db) = q\,\delta_0(db) + (1-q)\,\rho_+(db), \tag{8}$$

*where* $q := \mathbb{P}_{\nu_\varepsilon}(\varepsilon < 0)$ *is the spoke-binding probability in the typical-state regime, and* $\rho_+$ *is the law of* $\frac{Lw}{2}|\varepsilon|(1+s_h^2)$ *conditional on* $\varepsilon \geq 0$*, both to leading order in* $w$*.*

**Corollary 1** (Half-normal mixture, symmetric Gaussian case)**.** *If* $\nu_\varepsilon$ *is symmetric mean-zero Gaussian with standard deviation* $\sigma_{slip}$*, then* $q = 1/2$ *and* $\rho_+ = \mathrm{HN}(\beta)$ *is half-normal with scale*

$$\beta \;=\; \frac{Lw\,\sigma_{slip}\,(1+s_h^2)}{2} \;\xrightarrow{s\to 1,\ d_0=d_1}\; Lw\,\sigma_{slip},$$

*giving the closed form* $\mu(db) = \frac{1}{2}\delta_0(db) + \frac{1}{2}\,\mathrm{HN}(b;\beta)\,db$ *with explicit moments*

$$\mathbb{P}(D_{pool,T_k} = 0) = \tfrac{1}{2}, \qquad \mathbb{E}[D_{pool,T_k} \mid D_{pool,T_k} > 0] = \beta\sqrt{\tfrac{2}{\pi}}.$$

The decoupling argument is a leading-order mean-field calculation; the rate of convergence in $N$ at fixed $w$ is conditional on (A4). The $O(w)$ correction to the leading-order kernel preserves the atom mass and shifts the slack-side scale by $O(w)$. Numerical verification of Corollary 1 across $N \in \{10, 50, 100, 200, 500\}$ is in Section 6.3; verification of Proposition 3 under V9-realistic asymmetric heavy-tailed slippage is in Section 7.3.

The hub does not converge to the same limit as the spokes. It retains autocorrelation through its per-rebalance updates and exhibits a heavier right tail. The spokes, not the hub, are the propagation-of-chaos limit objects, since between consecutive spoke events the hub mixes through $\sim N$ other updates, so each spoke effectively sees a fresh hub sample. Proof in Appendix F.

### *6.2. The dust ledger as a loss-network analogue*

The dust ledger is a discovered, rather than designed, instance of a mathematical structure more familiar in queueing theory and statistical mechanics. The closest existing technical analogue is the propagation-of-chaos analysis of Graham and Méléard[8] for fully connected loss networks with alternate routing. The point of contact is summarised in Table 2.

Table 2. Correspondence between Graham-Méléard's loss network with alternate routing and the dust-ledger framework of this paper.

| Graham-Méléard[8] | Dust ledger (this paper) |
|---|---|
| Node | Token $T_i$ |
| Link (edge between two nodes) | Pool $P_k$ (pair $(T_a, T_b)$) |
| Call arrival on link, Poisson rate $\nu$ | Rebalance event on pool, exogenous Poisson rate $\lambda_k$ |
| Capacity check (link full or not) | LP-min binding side (Theorem 1) |
| Alternate routing to a third node when link saturates | Cross-pool absorption via shared depositor-keyed dust ledger (eq. (5)) |
| Non-exchangeable in pair structure (couples sharing a node ≠ disjoint couples) | Non-exchangeable in pair structure (positions sharing a token ≠ disjoint positions) |
| Fully connected $n$-node graph ⇒ $\binom{n}{2}$ links | Bipartite portfolio graph $G$, projected onto tokens |

The correspondence is nevertheless only partial. Graham-Méléard's calls have exponential durations after which channels are released; the dust ledger has no analogue of call expiry, and dust persists until consumed by a subsequent mint. Their state space is integer-valued, with channel counts in $\{0, 1, \ldots, C\}$, whereas the dust ledger evolves on the continuous state space $\mathbb{R}^n_{\geq 0}$. Lifting Graham-Méléard's variation-norm convergence theorem to the present kernel remains an open technical

step, as noted in Appendix F. Modern propagation-of-chaos methods for jump systems with quantitative rates[9,10] provide the more natural contemporary toolkit, subject to the corresponding Lipschitz conditions on the kernel.

### *6.3. Numerical verification*

The Master Equation simulator described in Appendix G implements the homogeneous hub-spoke dynamics directly. Tables 3–4 verify Corollary 1 (the symmetric-Gaussian special case) at $\sigma_{\text{slip}} = 0.01$. At $N = 200$, $L = 1$, $w = 0.01$, a run of $5 \times 10^6$ rebalance events confirms the closed form across all moments and quantile points to within 1–4%. At these parameters,

$$\beta = \frac{0.01}{\sqrt{2}} \approx 7.07 \times 10^{-3}$$

in the simulator's homogeneous symmetric convention, so the closed form gives

$$\mathbb{E}[b \mid b > 0] = \beta\sqrt{\tfrac{2}{\pi}} \approx 5.64 \times 10^{-5},$$

while the simulator returns $5.71 \times 10^{-5}$. The general atom-mixture form (Proposition 3) under V9-realistic asymmetric heavy-tailed slippage is verified separately in Section 7.3; the simulator's $\mathbb{P}(b = 0)$ matches the conditional $\mathbb{P}_{\nu_\varepsilon}(\varepsilon < 0)$ prediction to within 0.001 at the kernel level under each conditional bootstrap of V9 events.

Table 3. Verification of Corollary 1 for the spoke stationary distribution at $N = 200$ ($5 \times 10^6$ samples), under symmetric Gaussian slippage of standard deviation $\sigma_{\text{slip}} = 0.01$.

| Quantity | Empirical | Closed form | Ratio |
|---|---|---|---|
| $\mathbb{P}(b = 0)$ | 0.4998 | 0.5000 | 1.000 |
| $\mathbb{E}[b \mid b > 0]$ | $5.71 \times 10^{-5}$ | $5.64 \times 10^{-5}$ | 1.012 |
| $\mathrm{Var}(b \mid b > 0)$ | $2.01 \times 10^{-9}$ | $1.82 \times 10^{-9}$ | 1.108 |
| CDF $q = 0.10$ | $9.23 \times 10^{-6}$ | $8.89 \times 10^{-6}$ | 1.04 |
| CDF $q = 0.25$ | $2.24 \times 10^{-5}$ | $2.25 \times 10^{-5}$ | 0.99 |
| CDF $q = 0.50$ | $4.66 \times 10^{-5}$ | $4.77 \times 10^{-5}$ | 0.98 |
| CDF $q = 0.75$ | $8.21 \times 10^{-5}$ | $8.13 \times 10^{-5}$ | 1.01 |
| CDF $q = 0.90$ | $1.17 \times 10^{-4}$ | $1.16 \times 10^{-4}$ | 1.01 |
| $\mathrm{corr}(D_{\text{pool},T_1}, D_{\text{pool},T_2})$ | −0.027 | 0 | n/a |

**Rate-of-approach sweep across $N$.** To check that the asymptotic decoupling sets in at the rate suggested by the leading-order argument, the homogeneous hub-spoke chain is re-run at

$$N \in \{10, 50, 100, 200, 500\},$$

with parameters fixed at $L = 1$, $w = 0.01$, and $\sigma_{\text{slip}} = 0.01$, and with $10^5$ post-burn-in events per spoke. The simulator implements the full finite-$N$ coupling, drawing the slack-side leftover at the rebalance of pool $k$ as

$$|Lw + b_k - h| \cdot |\varepsilon_1 + \varepsilon_2|/2$$

rather than from the leading-order limit kernel. Any visible $N$-dependence therefore reflects the correction terms omitted by the leading-order argument.

Table 4. Rate-of-approach sweep against Corollary 1 (the leading-order half-normal mixture); $|\rho|$ medians taken across all $\binom{N}{2}$ spoke pairs.

| $N$ | $\mathbb{P}(b=0)$ | error vs $1/2$ | $\mathbb{E}[b \mid b>0]$ | rel. error | median $\lvert\rho\rvert$ |
|---|---|---|---|---|---|
| 10 | 0.5003 | $3.3\times10^{-4}$ | $5.641\times10^{-5}$ | 0.01% | 0.0021 |
| 50 | 0.5001 | $5.5\times10^{-5}$ | $5.644\times10^{-5}$ | 0.03% | 0.0022 |
| 100 | 0.5000 | $4.2\times10^{-6}$ | $5.643\times10^{-5}$ | 0.02% | 0.0021 |
| 200 | 0.5001 | $8.9\times10^{-5}$ | $5.644\times10^{-5}$ | 0.03% | 0.0021 |
| 500 | 0.4999 | $8.1\times10^{-5}$ | $5.643\times10^{-5}$ | 0.01% | 0.0028 |

The atom $\mathbb{P}(b=0) = 1/2$ and the slack-side scale

$$\mathbb{E}[b \mid b>0] = \beta\sqrt{\tfrac{2}{\pi}}$$

are recovered to within Monte Carlo noise across two orders of magnitude in $N$, with median pairwise spoke correlation remaining at the noise floor throughout. At the homogeneous parameters used here, the leading-order ansatz $|Lw + b_k - h| \approx Lw$ is already accurate to relative $O(\sigma_{\text{slip}})$, so the 1–4% gap in Table 3 between simulator and closed form is best understood as a fixed higher-order kernel correction (range half-width $w$ and swap-locus discreteness) rather than a finite-$N$ artefact. A rigorous rate-of-convergence statement à la[10] for this kernel class remains open, as does identifying a regime in which finite-$N$ departures from the leading-order limit become numerically visible, for example under heterogeneous spokes or larger $\sigma_{\text{slip}}$.

### 6.4. *Heterogeneous spokes and finite* $N$

**Conjecture 1** (Heterogeneous spokes)**.** *Drop the homogeneity assumption, and let spoke $k$ have its own parameters $(L_k, w_k, \sigma_{slip,k}, \lambda_k)$ and slippage law $\nu_{\varepsilon,k}$. If conditions analogous to (A1′)–(A5) hold on a spoke-by-spoke basis, then the factorisation structure of Proposition 2 and the leading-order atom-mixture structure of Proposition 3 are conjectured to persist in the following weak form: distinct spokes become asymptotically independent, and each spoke admits a stationary law of the form*

$$\mu_k = q_k\,\delta_0 + (1-q_k)\,\rho_{+,k}$$

*for some $q_k \in [0,1]$, with slack-side spoke leftover scaling to leading order as*

$$\frac{L_k w_k}{2}\,|\varepsilon_k|\,(1+s_h^2).$$

*No closed-form joint stationary law is claimed under arbitrary heterogeneity.*

Empirical V9 evidence (Section 7.3) is consistent with the conjecture's pairwise-decoupling claim, but the observed per-spoke variation in $\mathbb{P}(D_{\text{pool},T_k} = 0)$ over the range $[0, 0.59]$ across 32 V9 token-portfolio pairs indicates a heterogeneous production regime lying beyond this clean per-spoke form. That regime is treated as an open extension problem in Section 9.4.

**Numerical verification under heterogeneity.** The simulator-level claims of Conjecture 1 (atom mass, leading-order scaling, and pairwise decoupling under symmetric-Gaussian slippage) are tested by re-running the chain with per-spoke parameters drawn log-uniformly:

$$L_k \in [0.5, 2], \qquad w_k \in [0.005, 0.05], \qquad \sigma_{\text{slip},k} \in [0.005, 0.05], \qquad \lambda_k \in [0.5, 2].$$

This produces a 4−10× spread on each parameter and a ∼ 100× spread in $\beta_k$ across spokes. Pool selection at each event is weighted by $\lambda_k / \sum_j \lambda_j$, matching the Poisson superposition of independent rebalance streams. Table 5 shows that, for every $N$ from 10 up to 500, the half-mass atom holds per spoke (median $\mathbb{P}(b_k = 0)$ within $10^{-3}$ of 0.5), the leading-order leftover scaling holds per spoke (median relative error on $\mathbb{E}[b_k \mid b_k > 0]$ in the range 0.5−0.8%, with p90 below 2%), and pairwise spoke-spoke correlations remain at the noise floor (median $|\rho|$ below 0.005, p95 below 0.012).

Table 5. Heterogeneous-spoke verification of Conjecture 1; each row uses fresh log-uniform parameter draws.

| $N$ | med $\mathbb{P}(b_k = 0)$ | max $\lvert\mathbb{P} - 1/2\rvert$ | med rel err | p90 | max | med $\lvert\rho\rvert$ | p95 $\lvert\rho\rvert$ |
|---|---|---|---|---|---|---|---|
| 10 | 0.4994 | 0.0049 | 0.55% | 1.06% | 1.81% | 0.0027 | 0.0068 |
| 50 | 0.4999 | 0.0038 | 0.76% | 1.43% | 2.62% | 0.0028 | 0.0087 |
| 100 | 0.4998 | 0.0050 | 0.50% | 1.55% | 4.31% | 0.0033 | 0.0094 |
| 200 | 0.4996 | 0.0057 | 0.55% | 1.41% | 3.21% | 0.0032 | 0.0092 |
| 500 | 0.4999 | 0.0073 | 0.67% | 1.90% | 4.47% | 0.0041 | 0.0118 |

These simulator-level results support the conjecture's limited claims across two orders of magnitude in $N$ and across a heterogeneity range comparable to that observed in production V9 portfolios. They do not, however, yield a closed-form expression for the joint stationary law under arbitrary heterogeneity. Production V9 portfolios have $N \in \{2, \ldots, 8\}$, and the per-spoke atom-mass variation observed there indicates a heterogeneous production regime outside the conjecture's clean per-spoke form. Section 7.3 therefore tests the pairwise-decoupling claim directly on operator data, while Section 9.4 treats the residual per-spoke distributional gap as an open extension problem.

## 7. Empirical validation

The framework is validated across the four PM contracts in the observed architectural class on Aerodrome Slipstream (Base): PM V6 (active 1–5 March 2026, where the residual was originally observed), PM V7 (6–13 March, following post-exploit security hardening), PM V8 (13–21 March, introducing multi-portfolio support), and PM V9 (21 March onward, the current deployment). V9 is indexed natively from chain into a DuckDB lake, whereas V6, V7, and V8 are reconstructed from per-version diffusion logs. In conservation form, the Master Equation combines the rebalance transaction's Pool.Burn, Pool.Swap, and Pool.Mint amounts with the standing-balance state obtained by chronological `DustCredited` cumulation, and compares the result with the next `DustCredited` emitted at the rebalance. Contract addresses, indexer plumbing, and the supporting audits are given in Appendix G; the reconstructed credit history agrees with the on-chain `dustBalance` mapping at 99.83% across 16,388 comparison rows.

### *7.1. Cross-version replication: V6 to V9*

The same equation holds on every contract version back to where the residual was first observed. Table 6 reports the conservation-form result on each of the four contracts in the architectural class.

Table 6. Master Equation conservation form, V6→V9.

| Version | Active window | Groups | Events | Median ratio | Within-2× |
|---|---|---|---|---|---|
| PM V6 | 2026-03-01 to 03-05 | 5 | ~738 | 1.0000 uniformly | 91.2–100.0% |
| PM V7 | 2026-03-06 to 03-13 | 5 | ~3,600 | 1.0000 uniformly | 98.2–100.0% |
| PM V8 | 2026-03-13 to 03-21 | 13 | ~22,000 | 1.0000 uniformly | 97.8–100.0% |
| PM V9 | 2026-03-21 onward | 8 | ~8,200 | 1.0000 uniformly | 99.8–100.0% |
| Total | 23-day span | 31 | ~34,500 | 1.0000 across 31/31 | 91.2–100.0% |

The conservation-form match is exact in median terms across the entire version history. The median ratio is 1.0000 in all 31 groups, and the within-2× rate is ≥ 97.8% on V7–V9 and ≥ 91.2% on V6, whose lower floor reflects the retroactive diffusion-log reconstruction, since contemporaneous engine logging begins only at V7. The V6 to V7 security hardening, V7 to V8 schema extension, and V8 to V9 AERO-leak fix each touched substantial portions of the contract code, yet the dust-ledger dynamics obey the same equation throughout. The cross-version evidence is therefore consistent with the Master Equation governing every observed contract version in this architectural class, though the four contracts still share one operator (Section 9.3), so operator-independent invariance remains open. Fig. 4 anchors the kernel structure of Theorem 1 on the swap-corrected branch independently of the conservation form, by comparing the per-event leading-order prediction to the on-chain dust credit on each of 804 substantial-swap V9 events.

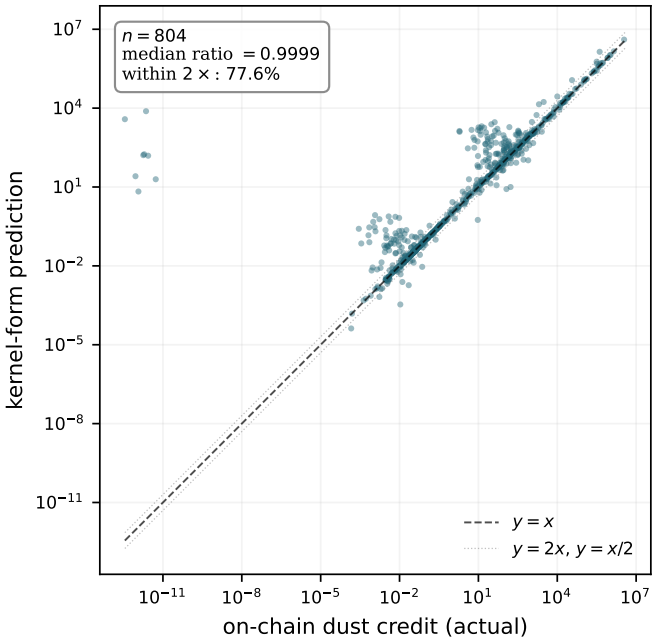


Fig. 4. Per-event kernel prediction of slack-side dust credit versus on-chain emission, on 804 V9 substantial-swap events. Geomedian predicted/actual = 0.9999; 77.6% within a factor of 2 of the diagonal.

### *7.2. V9 per-portfolio detail and predictive-form lift*

The V9 row of Table 6 is broken down by portfolio in Table 7. The depositor's on-chain universe contains 13 active portfolios, of which eight have a sufficient sample of events together with the diffusion-log inputs and pool coverage required for the prediction.

Table 7. Master Equation conservation form on V9, per portfolio.

| Portfolio | Events | Binding | Within 2× | Within 10% | Median | Geomean |
|---|---|---|---|---|---|---|
| A | 2,332 | 100.0% | 100.0% | 100.0% | 1.0000 | 1.00 |
| B | 2,556 | 99.9% | 99.9% | 99.8% | 1.0000 | 1.00 |
| C | 1,060 | 99.5% | 100.0% | 100.0% | 1.0000 | 1.00 |
| D | 559 | 100.0% | 100.0% | 100.0% | 1.0000 | 1.00 |
| E | 504 | 99.2% | 99.8% | 99.8% | 1.0000 | 1.00 |
| F | 443 | 99.8% | 100.0% | 100.0% | 1.0000 | 1.00 |
| G | 381 | 100.0% | 100.0% | 100.0% | 1.0000 | 1.00 |
| H | 345 | 99.7% | 100.0% | 100.0% | 1.0000 | 1.00 |

Across 8,180 V9 events, the within-10% rate (defined as $|\log_{10}(\text{predicted}/\text{actual})| < \log_{10}(1.1)$) is 99.8–100% on every portfolio. The four conservation inputs are read from distinct on-chain events, and their algebraic combination matches the next `DustCredited` exactly as a closed-form identity would predict. The same equation reproduces swap-free and swap-present events alike.

**Predictive form.** The conservation form above reads its mint inputs directly from the on-chain Pool.Mint event. The predictive form of Theorem 1 instead derives the mint inputs from the LP solve at the engine-detected post-swap currentTick, and derives $(x_w, y_w)$ from the engine's `preValue` inverted at the engine-detected pre-swap currentTick. Table 8 compares the two forms.

Table 8. Theorem 1 as forward predictor (LP-solve column) versus mass-conservation check (conservation column).

| | | Within 2× | | Geomean | |
|---|---|---|---|---|---|
| Portfolio | Events | LP-solve | Conservation | LP-solve | Conservation |
| A | 2,332 | 80.2% | 100.0% | 1.26 | 1.00 |
| B | 2,556 | 65.4% | 99.9% | 1.39 | 1.00 |
| C | 1,060 | 62.2% | 100.0% | 1.55 | 1.00 |
| D | 559 | 73.8% | 100.0% | 2.52 | 1.00 |
| E | 504 | 57.3% | 99.8% | 1.37 | 1.00 |
| F | 443 | 80.2% | 100.0% | 1.78 | 1.00 |
| G | 381 | 71.8% | 100.0% | 1.43 | 1.00 |
| H | 345 | 96.7% | 100.0% | 1.04 | 1.00 |

The accuracy of the predictive form is capped by engine-tick observability. On narrow-range stable pools, the on-chain tick can drift between firing and burn, shifting $(x_w, y_w)$ in ways the engine snapshot does not capture. In that regime, the predictive form recovers the residual within a factor of two on 57–97% of events, whereas the conservation form (which reads Pool.Burn directly from chain) is unaffected ($\geq 99.8\%$).

### *7.3. Framework predictions verified on V9*

Beyond the per-event conservation check, Proposition 1, Theorem 2, and Propositions 2–3 make distinct predictions about V9 dust dynamics; each is verified on V9 chain data below.

**Connector Rule (Proposition 1).** Per-pool Spearman correlations between signed normalised displacement and signed connector-side dust quantity are computed across the eight portfolios containing a multi-pool connector, giving 37 pool-portfolio groups in Table 9, a 3.7× increase over the 10 groups reported in §5.2 of the prior paper. Sixteen pools (43%) attain $|\rho| \geq 0.5$, and the sign separates cleanly by connector position: the $T_0$-connector subset has median +0.42, while the $T_1$-connector subset has median −0.44. The original empirical sample, which ranged from −0.56 to −0.98 across 10 groups, lay entirely in the $T_1$-connector subset; the V9 evidence now covers both connector positions and recovers the predicted sign in each.

Table 9. Connector Rule Spearman $\rho$ summary, V9.

| Connector position | Pools | Median $\rho$ | Range |
|---|---|---|---|
| $T_0$ (token0) | 20 | +0.42 | [+0.32, +0.67] for $\|\rho\| \geq 0.3$ |
| $T_1$ (token1) | 17 | −0.44 | [−0.83, −0.50] for $\|\rho\| \geq 0.3$ |
| All pools | 37 | n/a | [−0.83, +0.67] overall |

**Donor/absorber decomposition (Theorem 2).** Every multi-pool V9 portfolio observed under $S = 0$ partitions into net donors and net absorbers over the observation window, while preserving the event-level identity of Theorem 2. Table 10 reports the partition: all 8 multi-pool portfolios contain both donors and absorbers, while the two single-pool portfolios (D and J) are degenerate. Portfolio B provides the canonical hub-spoke example at portfolio scale, with a perfect 4:4 donor-absorber split and near-zero net total (+$1,285); dust moves between positions sharing the same depositor-keyed dust mapping while the aggregate remains approximately balanced over the window. By contrast, Portfolio H records a net total of −$7,697, indicating that donor and absorber flows need not offset exactly over a finite observation horizon even when the event-level conservation law holds.

Table 10. Per-portfolio donor/absorber decomposition over the V9 window.

| Portfolio | Pools | Donors | Absorbers | Neutrals | Net total (USD) |
|---|---|---|---|---|---|
| A | 6 | 0 | 5 | 1 | +17,783 |
| B | 8 | 4 | 4 | 0 | +1,285 |
| C | 8 | 2 | 4 | 2 | +13,003 |
| D | 1 | 1 | 0 | 0 | −146 |
| E | 5 | 3 | 2 | 0 | +2,497 |
| F | 6 | 2 | 4 | 0 | +24,004 |
| G | 6 | 4 | 1 | 1 | +847 |
| H | 6 | 3 | 2 | 1 | −7,697 |
| I | 2 | 1 | 1 | 0 | +14,724 |
| J | 1 | 0 | 1 | 0 | +63,443 |

**Asymptotic spoke independence (Proposition 2).** The propagation-of-chaos prediction that distinct spokes become asymptotically independent remains empirically testable on V9, even at finite $N$. Across the eight portfolios there are 12 degree-1 token pairs, that is, pairs in which each token appears in exactly one pool. Table 11 reports the pairwise correlations of their standing-balance trajectories: median $|\rho| = 0.020$ and maximum 0.14. Portfolio C alone contributes six pairs with $|\rho| \leq 0.025$, with all four of its spokes effectively independent. The two largest correlations occur at the smallest $N$, in Portfolios H and I where $N \leq 4$, consistent with the $O(1/N)$ finite-$N$ correction discussed in Section 6.4.

The factorisation diagnostic is also stable across slippage-law families. Bootstrap simulators driven by V9-realistic asymmetric heavy-tailed slippage reproduce the same noise-floor scaling

$$|\rho| \cdot \sqrt{\text{events}/N} \approx \text{const}$$

across $N \in \{100, \dots, 10\,000\}$ under three independent conditional-bootstrap kernels (i.i.d. marginal, sign-conditional, magnitude-conditional), with max/min of the implied noise constant at most 1.34 in each case.

**Atom-mixture closed form (Proposition 3).** The kernel-level prediction that the spoke marginal takes the binary atom-mixture form is supported at the simulator level under bootstrap from V9 events: the simulator's $\mathbb{P}(b = 0)$ matches the conditional prediction $\mathbb{P}_{\nu_\varepsilon}(\varepsilon < 0)$ to within 0.001 across $N \in \{100, \dots, 10\,000\}$ under both sign-conditional and magnitude-conditional bootstraps. The half-normal special case (Corollary 1) has, by contrast, only narrow empirical scope. Across 32 V9 token-portfolio pairs, three exhibit the half-mass atom

$$\mathbb{P}(D_{\text{pool},T_k}(\tau^+) = 0) \in [0.40, 0.60]$$

predicted by the symmetric-Gaussian Corollary, distributed across three distinct V9 portfolios; the remaining 29 spokes behave as persistent slack-side accumulators with $\mathbb{P}(b = 0)$ near zero. This per-spoke variation lies outside the homogeneous symmetric closure and is therefore treated, in Section 9.4, as an open extension problem rather than as a within-scope refinement of Proposition 3.

Table 11. Pairwise correlation between distinct spoke-token standing trajectories on V9 portfolios. Median $|\rho| = 0.020$, max 0.14 across 12 pairs.

| Portfolio | Spoke pair | $\lvert\rho\rvert$ | Sample size |
|---|---|---|---|
| C | ZORA / BNKR | 0.0007 | 13,216 |
| C | ZORA / DIEM | 0.0055 | 13,216 |
| G | UP / MORPHO | 0.010 | 8,757 |
| C | BNKR / DIEM | 0.018 | 13,216 |
| E | cbBTC / deSPXA | 0.018 | 6,824 |
| C | DIEM / VIRTUAL | 0.020 | 13,216 |
| C | ZORA / VIRTUAL | 0.020 | 13,216 |
| C | BNKR / VIRTUAL | 0.023 | 13,216 |
| E | WETH / deSPXA | 0.0067 | 6,824 |
| E | cbBTC / WETH | 0.052 | 6,824 |
| H | CADC / COOKIE | 0.138 | 3,669 |
| I | deSPXA / USDC | 0.139 | 1,314 |

## 8. Contract-level verification

The cross-version replication of Section 7.1 is anchored in operational data. A complementary track verifies the closed-form predictions of Theorem 1, Theorem 2, and Proposition 1 directly against unmodified V3 `NonfungiblePositionManager` contracts under controlled mint, rebalance, and withdraw flows. The full Foundry suite, including contract addresses and per-test residual tables, is given in Appendix H.

On both DEXes exercised by the live fork tests, Aerodrome Slipstream and Uniswap V3, the predicted dust credit matches the emitted credit to within 1 wei (token0) and 1 raw unit (token1) on every event, including the swap-corrected and cross-pool cases. Theorem 2's mass conservation under $S = 0$ is reproduced bit-exactly across 200 randomised rebalances on the mock pool; Proposition 1's sign prediction holds in both connector-position configurations on a controlled hub-spoke graph.

The two NFPMs are independently developed and differ in pool-key encoding (`tickSpacing` on Slipstream, `fee` on Uniswap V3). Bit-exact agreement across both is therefore consistent with the per-event jump being a property of the architectural class rather than of any particular DEX implementation.

## 9. Discussion

### *9.1. What the framework subsumes*

The Master Equation (5) gathers several earlier results in the programme into a single dynamical law. In the swap-free single-pool case, its per-event jump recovers the geometric residual of §3.2 of the prior paper; the iff vanishing condition of Theorem 1 there reappears here as the measure-zero ratio-preserving locus inside Theorem 1. The single-pool corner $m = 1$ then separates by swap regime: under non-zero $S$, it is the convergence-in-expectation regime of Theorem 2 there; under $S = 0$, it is the geometric extinction of Theorem 3.

The Connector Rule conjecture of §3.5 is promoted here to Proposition 1 under explicit V3-mechanical assumptions, while the hub-spoke study of §5.4 may now be read, in retrospect, as the empirical instance of the canonical hub-spoke case of Section 6.1. The rate-ratio identity $\tau_{\text{hub}}/\tau_{\text{spoke}} = 1/N$ is likewise consistent with Proposition 2 of[2] on homogeneous-rate hub-spoke graphs.

### *9.2. Architectural precondition*

Stock NFPMs do not maintain a depositor-keyed dust mapping. The architectural-precondition fork test of §7.1 of the prior paper confirms the absence of cross-position dust absorption on a stock NFPM, while Proposition 1 of[2] gives the formal iff condition (token degree $\geq 2$ in the portfolio graph) for cross-pool circulation paths to exist. The four PM contracts examined here remain the full observed population of this architectural class.

The mechanism studied here is therefore not a generic property of concentrated liquidity as such. It arises when the V3 mint geometry is coupled to a contract architecture that retains and recycles residual balances across positions. In that sense, the dust ledger is not merely an implementation detail but the hidden microstructural state through which the portfolio becomes dynamically coupled.

### *9.3. Limitations*

The empirical anchor is single-operator and single-chain. The cross-version replication of Section 7.1 shows that the framework is not a peculiarity of any one contract version, but the four contracts still share an operator on Aerodrome Slipstream. Cross-operator replication on V9 is straightforward and remains to be carried out.

The atom-mixture closed form of Proposition 3 is exact at the kernel level under arbitrary $\nu_\varepsilon$ but the rate of convergence in $N$ is conditional on (A4); finite-$N$ corrections are visible at chain scale $N \leq 8$ in the V9 data. The half-normal special case (Corollary 1) has correspondingly narrow empirical scope: 3 of 32 V9 spokes match its half-mass atom, while the remaining 29 act as persistent slack-side accumulators with $\mathbb{P}(b = 0)$ close to zero. This per-spoke variation lies outside the homogeneous symmetric closure (Section 9.4). The single-pool swap-corrected stationary distribution likewise remains characterised only as a contraction-mapping fixed point.

The predictive form of Theorem 1 (Section 7.2) recovers per-event dust credits to within a factor of two on 57–97% of events given engine-observable inputs. The gap to the conservation form's ≥ 99.8% is engine-tick observability at the burn moment, not framework misspecification. A protocol that records the actual on-chain pre-burn currentTick alongside its rebalance event would close that gap directly.

### *9.4. Future work*

The empirical V9 spoke-marginal distribution exhibits per-spoke variation in the atom mass $\mathbb{P}(D_{\text{pool},T_k}(\tau^+) = 0)$ ranging over $[0, 0.59]$ across 32 token-portfolio pairs. Most spokes act as persistent slack-side accumulators; the three half-mass-atom matches identified in Section 7.3 are the exception. This per-spoke heterogeneity lies outside the homogeneous symmetric leading-order closure of Proposition 2 and points instead to a richer extension problem: characterising the dust ledger's stationary distribution under per-pool LP-asymmetry, per-pool $\nu_\varepsilon$, and joint $(\sigma_x^*, \varepsilon, \text{state})$ dynamics on heterogeneous portfolio graphs.

The natural analytical candidates are quantitative McKean–Vlasov rates for jump systems with empirical-measure dependence,[9,10] adapted to the discrete-event multi-token kernel of Theorem 1. The full production law is therefore best treated as an open extension problem rather than a within-scope refinement: Propositions 2–3 are exact under their respective assumption blocks, whereas the heterogeneous production regime occupied by V9 in practice remains open.

Three additional theoretical questions remain open: the single-pool swap-corrected fixed point, the finite-$N$ rate of approach to the chaos limit (a quantitative propagation-of-chaos question along the lines of[10]), and a non-parametric closed form for $\rho_+$ in Proposition 3 under empirical V9 slippage (no standard 3–4 parameter family appears to capture the observed heavy tails and asymmetry simultaneously). A separate empirical direction concerns the directional-asymmetry result of the prior paper, which the per-pool swap-event tick stream of Section 4.2 now makes it possible to decompose into numéraire and drift components. The V9 data suggest that the asymmetry persists on volatile–volatile pairs (median $D^+/D^- = 1.11$, against 1.05 on volatile–stable pairs), making that decomposition a natural next step. Even so, the most informative immediate extension would be cross-operator replication of the operational anchor of Section 7.1 on an independent deployment of the same architecture.

## 10. Conclusion

The dust ledger of a multi-pool concentrated liquidity manager is a stochastic flow network on the portfolio bipartite graph, governed by a closed-form Master Equation extending the binding-token residual of the prior paper.[1] That equation subsumes the single-pool results developed there, promotes the Connector Rule conjecture to a sign-keyed correlation, generalises zero-swap extinction to a per-token mass-conservation identity across multi-pool rebalances, and yields a hierarchy of asymptotic results on the homogeneous hub-spoke topology: a conditional propagation-of-chaos limit (Proposition 2) and a closed-form binary atom-mixture spoke marginal under arbitrary slippage law (Proposition 3), with the half-normal mixture as a narrow corollary in the symmetric-Gaussian special case.

The mechanism follows directly from the V3 amount equations once depositor-keyed shared balances are present. The V6 through V9 evolution of a PM contract family on Aerodrome Slipstream (Base) provides an operational anchor for that law, while contract-level verification exhibits the same mechanism on unmodified Aerodrome Slipstream and Uniswap V3 infrastructure. Within the architectural class observed here, the dust ledger is not merely a residual store. It is a governed dynamical state through which concentrated liquidity positions become coupled across the portfolio.

## References

[1] Ryan, K.R. "The Geometric Siphon." *SSRN preprint 6686798* (2026) https://ssrn.com/abstract=6686798.

[2] 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.

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

[5] 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.

[6] 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.

[7] Sznitman, A.-S. "Topics in Propagation of Chaos." In *École d'Été de Probabilités de Saint-Flour XIX, 1989*, Lecture Notes in Mathematics 1464, pp. 165–251, Springer (1991) https://doi.org/10.1007/BFb0085169.

[8] Graham, C., Méléard, S. "Propagation of chaos for a fully connected loss network with alternate routing." *Stochastic Processes and their Applications* **44**(1), 159–180 (1993) https://doi.org/10.1016/0304-4149(93)90043-4.

[9] Chaintron, L.-P., Diez, A. "Propagation of chaos: a review of models, methods and applications. I. Models and methods." *Kinetic and Related Models* **15**(6), 895–1015 (2022). arXiv:2203.00446 https://arxiv.org/abs/2203.00446.

[10] Andreis, L., Dai Pra, P., Fischer, M. "McKean–Vlasov limit for interacting systems with simultaneous jumps." *Stochastic Analysis and Applications* **36**(6), 960–995 (2018). arXiv:1704.01052 https://arxiv.org/abs/1704.01052.

[11] Kelly, F. P. *Reversibility and Stochastic Networks*. Wiley, Chichester (1979). Reprinted by Cambridge University Press, 2011.

[12] Carmona, R., Delarue, F. *Probabilistic Theory of Mean Field Games with Applications II: Mean Field Games with Common Noise and Master Equations*. Probability Theory and Stochastic Modelling, vol. 84, Springer (2018) https://doi.org/10.1007/978-3-319-56436-4.

[13] Cardaliaguet, P., Delarue, F., Lasry, J.-M., Lions, P.-L. *The Master Equation and the Convergence Problem in Mean Field Games*. Annals of Mathematics Studies 201, Princeton University Press (2019) https://press.princeton.edu/books/paperback/9780691193137/.

## Appendix A: Notation

The notation matches the prior paper[1] throughout, extending scalar quantities specific to one pool to vector-valued, time-indexed analogues without changing the names.

Table 12. Notation, with cross-reference to the prior paper.

| Symbol | Meaning |
|---|---|
| ***Pool geometry and position state.*** | |
| $s$, $s_a$, $s_b$, $s'_a$, $s'_b$ | Sqrt-prices: current; old range bounds; new range bounds. |
| $L$, $L_{\text{new}}$ | Position liquidity before and after the rebalance. |
| $x_w$, $y_w$ | Withdrawn amounts of token0 and token1. |

*Continued on next page.*

*Table 12 (continued).*

| Symbol | Meaning |
|---|---|
| $\delta$, $\bar{s}$, $w$ | Signed normalised displacement; mid-sqrt-price of the new range; new-range width $s'_b - s'_a$ (see Appendix D). |
| $V$, $V_{\text{new}}$, $V_{\text{old}}$ | Position USD value: instantaneous ($V$); pre- and post-rebalance ($V_{\text{old}}$, $V_{\text{new}}$). |
| $\Delta R$ | Signed residual value (scalar). |
| $\phi(s, s_a, s_b)$ | Per-unit-$L$ position USD value, in-range, under the stable-numéraire convention $p_{T_0} = s^2$, $p_{T_1} = 1$: $\phi = 2s - s^2/s_b - s_a$. For out-of-range cases, see Appendix B. |
| $g_x = 1/s - 1/s'_b$, $g_y = s - s'_a$ | Per-unit-$L$ token amount at the *new* range. |
| ***Dust ledger.*** | |
| $x_{\text{dust}}, y_{\text{dust}}$ | Pre-rebalance standing dust available at mint, per pool token; snapshot of $D_{\text{pool}}$ on the pool's tokens. |
| $x'_{\text{dust}}, y'_{\text{dust}}$ | Per-token mint leftover (Theorem 1). Swap-free: one is zero, the other equals $\lvert\Delta R\rvert$. |
| $D_{\text{pool}}(t) \in \mathbb{R}^n_{\geq 0}$ | Vector of standing-dust balances at time $t$, indexed by portfolio token $T_i$. Pool-specific snapshot: $x_{\text{dust}} = D_{\text{pool},T_a}(t)$, $y_{\text{dust}} = D_{\text{pool},T_b}(t)$. |
| ***Swap correction.*** | |
| $S$, $\sigma_x$, $\sigma_y$ | $S$ (scalar): fraction of $\Delta R$ absorbed by the contract-level swap (prior paper). $(\sigma_x, \sigma_y)$: signed per-token decomposition at the mint inputs (pool-frame). Swap-free: $\sigma_x = \sigma_y = 0 \iff S = 0$. |

The $g$-notation is local shorthand for the new-range amount-per-$L$ functions, used here only where the V3 amount expressions would otherwise repeat within an equation;[1] uses inline forms.

## Appendix B: V3 amount formulas

The V3 amount functions for a position over $[s_a, s_b]$ at sqrt-price $s$ are

$$a_x(s, s_a, s_b) = \begin{cases} \dfrac{1}{s_a} - \dfrac{1}{s_b} & \text{if } s \leq s_a \\ \dfrac{1}{s} - \dfrac{1}{s_b} & \text{if } s \in [s_a, s_b] \\ 0 & \text{if } s \geq s_b, \end{cases} \tag{9}$$

$$a_y(s, s_a, s_b) = \begin{cases} 0 & \text{if } s \leq s_a \\ s - s_a & \text{if } s \in [s_a, s_b] \\ s_b - s_a & \text{if } s \geq s_b. \end{cases} \tag{10}$$

These are the $L = 1$ normalisations; for liquidity $L$, the corresponding token amounts are $L\,a_x(s, s_a, s_b)$ and $L\,a_y(s, s_a, s_b)$. The functions are monotone in the standard way: $a_y$ is non-decreasing in $s$, while $a_x$ is non-increasing, reflecting the standard V3 single-sided regimes at the range boundaries.

## Appendix C: Proof of Theorem 1

The mint chooses $L_{\text{new}}$ to maximise deployed liquidity subject to the available token balances. This appendix solves the resulting linear programme explicitly in each of the three positional sub-cases.

***The mint as a linear programme***

Let

$$g_x = a_x(s, s'_a, s'_b), \qquad g_y = a_y(s, s'_a, s'_b).$$

By Appendix B, both satisfy $g_x \geq 0$ and $g_y \geq 0$. The mint problem is

$$L_{\text{new}} = \max\{L \geq 0 : L\, g_x \leq \hat{x} \text{ and } L\, g_y \leq \hat{y}\}. \tag{11}$$

This is a one-dimensional linear programme in $L$. Since the objective is increasing in $L$, the optimum is attained when at least one constraint binds. The two candidate binding values are

$$L^*_x = \hat{x}/g_x \text{ (when } g_x > 0), \qquad L^*_y = \hat{y}/g_y \text{ (when } g_y > 0). \tag{12}$$

Hence

$$L_{\text{new}} = \min(L^*_x, L^*_y),$$

with the binding side determined by whichever candidate is smaller.

***Sub-case (i): in range***

Suppose $s \in (s'_a, s'_b)$. Then

$$g_x = 1/s - 1/s'_b > 0, \qquad g_y = s - s'_a > 0,$$

so both candidate liquidities are finite. If $L^*_x \leq L^*_y$, the $x$-constraint binds, so

$$L_{\text{new}} = \hat{x}/g_x, \quad x'_{\text{dust}} = 0, \quad y'_{\text{dust}} = \hat{y} - L_{\text{new}}\, g_y = \hat{y} - \hat{x}\,(g_y/g_x). \tag{13}$$

If $L^*_y < L^*_x$, the $y$-constraint binds, so

$$L_{\text{new}} = \hat{y}/g_y, \quad x'_{\text{dust}} = \hat{x} - L_{\text{new}}\, g_x, \quad y'_{\text{dust}} = 0. \tag{14}$$

In the swap-free regime $\hat{x} = x_w + x_{\text{dust}}$ and $\hat{y} = y_w + y_{\text{dust}}$, so the binding side is determined by the supply ratio $\hat{x}/\hat{y}$ versus $g_x/g_y$, with $x'_{\text{dust}} \cdot y'_{\text{dust}} = 0$ in either case.

***Boundary sub-cases***

Sub-cases (ii) and (iii) cover the boundary regimes $s \leq s'_a$ and $s \geq s'_b$. These lie outside Theorem 1's hypothesis but are reproduced for completeness; the LP solution remains explicit in each.

If $s \leq s'_a$, then

$$g_x = 1/s'_a - 1/s'_b > 0, \qquad g_y = 0.$$

Only the $x$-constraint is active, so

$$L_{\text{new}} = \hat{x}/g_x, \quad x'_{\text{dust}} = 0, \quad y'_{\text{dust}} = \hat{y}.$$

The new position is therefore single-sided in $T_a$ at the lower boundary of the range.

If $s \geq s'_b$, the argument is symmetric. In that case

$$g_x = 0, \qquad g_y = s'_b - s'_a > 0,$$

so only the $y$-constraint is active and

$$L_{\text{new}} = \hat{y}/g_y, \quad x'_{\text{dust}} = \hat{x}, \quad y'_{\text{dust}} = 0.$$

***Swap-corrected inputs***

With a non-zero swap correction $(\sigma_x, \sigma_y)$ as in Section 4.2, the inputs become

$$\hat{x} = x_w + x_{\text{dust}} - \sigma_x, \qquad \hat{y} = y_w + y_{\text{dust}} - \sigma_y.$$

The LP structure in (11) is unchanged, so the same binding-side and leftover formulas apply verbatim with the corrected inputs. The ratio-preserving locus in Theorem 1 of the prior paper becomes

$$\hat{x}/\hat{y} = g_x/g_y,$$

which is the routing target of the swap step. Slippage and sqrt-price drift during the swap then generate a measure-non-zero noise floor on which $x'_{\text{dust}} \cdot y'_{\text{dust}} \neq 0$. ■

## Appendix D: Proof of Proposition 1 (Connector Rule)

The argument has two steps. First, Theorem 1 identifies the binding side at a rebalance and therefore the side on which dust is left over. Second, under the connector-token hypothesis, the signed displacement of a pool inherits a shared component from the connector token's marginal-price evolution, which induces the predicted cross-rebalance sign pattern in connector-side dust.

***Per-rebalance binding side***

Consider a pool $P_k$ with tokens $T_a$ at the $T_0$ position and $T_b$ at the $T_1$ position, current sqrt-price $s$, old range $[s_a, s_b]$, and new range $[s'_a, s'_b]$, with $s \in [s'_a, s'_b]$. For clarity, first take the swap-free regime $\sigma_x = \sigma_y = 0$; the swap-corrected case is identical with $\hat{x}, \hat{y}$ replacing $x_w + x_{\text{dust}}$ and $y_w + y_{\text{dust}}$ throughout.

By the in-range case of Appendix C, the binding side is determined by

$$T_b \text{ binds} \iff \frac{\hat{x}}{\hat{y}} > \frac{g_x}{g_y}, \qquad T_a \text{ binds} \iff \frac{\hat{x}}{\hat{y}} < \frac{g_x}{g_y}, \tag{15}$$

where

$$g_x = 1/s - 1/s'_b, \qquad g_y = s - s'_a.$$

Dust appears on the slack side: $T_a$ leaks when $T_b$ binds, and $T_b$ leaks when $T_a$ binds.

At the rebalance moment, the carried position composition is

$$x_w = L\, a_x(s, s_a, s_b), \qquad y_w = L\, a_y(s, s_a, s_b), \tag{16}$$

so the supply ratio $\hat{x}/\hat{y}$ depends on $s$ through the V3 amount functions at the old range and through the standing-balance ratio at this rebalance.

Define the signed normalised displacement

$$\delta = \frac{s - \bar{s}}{w/2},$$

where $\bar{s}$ is the mid-sqrt-price of the new range and $w = s'_b - s'_a$. Holding the new-range geometry fixed and treating the standing-balance ratio as locally fixed across nearby rebalances of $P_k$, condition (15) reduces to a sign condition on $\delta$: when $\delta > 0$ (sqrt-price above the new-range mid), the position composition tilts toward $T_b$, $\hat{x}/\hat{y}$ falls, and $T_a$ tends to bind, with $T_b$ leaking; when $\delta < 0$, the converse holds.

This establishes the per-rebalance sign rule. The sign of displacement determines which side is slack, and therefore which token receives the dust residual.

***Population-level correlation***

Now let $T^*$ be the connector token of Proposition 1, so that $\deg(T^*) \geq 2$ and $\deg(T^*) > \deg(T)$ for every non-connector token $T$ paired with it. Across all pools containing $T^*$, the connector contributes a shared component to the pool's marginal-price evolution, whereas the non-connector token contributes a pool-specific component.

Under the stationarity assumption that the connector-token marginal price has no systematic drift over the rebalance window (relative to the rest of the portfolio), and the homogeneity assumption that in-range rebalances dominate, the signed displacement of pool $P_k$ may be written in reduced form as

$$\delta_k \approx \delta_{T^*} + \epsilon_k,$$

where $\delta_{T^*}$ is the connector-shared component and $\epsilon_k$ is a pool-specific term.

The signed connector-side dust quantity, as defined in the statement of Proposition 1, is the dust credited to the non-connector side minus the dust credited to the connector's side: positive when the non-connector side leaks and negative when the connector's side leaks.

If $T^* = T_a$ (the $T_0$ position), then by the per-rebalance sign rule above, $\delta_k > 0$ implies $T_b$ leaks and signed connector-side dust is positive; $\delta_k < 0$ implies $T_a$ leaks and signed connector-side dust is negative. Hence displacement and signed connector-side dust move with the same sign across rebalances, implying a positive Spearman rank correlation. By symmetry, if $T^* = T_b$ (the $T_1$ position), the sign flips and the correlation is negative.

The magnitude of the correlation depends on the relative importance of the shared connector component $\delta_{T^*}$ versus the pool-specific noise $\epsilon_k$. As a heuristic, let $r := \mathrm{Var}(\delta_{T^*})/\mathrm{Var}(\delta_{T^*} + \epsilon_k)$ denote the variance share of the connector-shared component; one then expects the magnitude of the observed rank correlation to increase with $r$, approaching one as $r \to 1$ and falling toward zero as pool-specific noise dominates. The V9 sample (Section 7.3) recovers both the sign-by-position prediction and the magnitude-scaling pattern; pools where $|\rho|$ is weakest correspond to pool-specific volatility regimes where $\epsilon_k$ dominates, consistent with the heuristic rather than a framework failure.

■

## Appendix E: Proof of Theorem 2 and partition lemma

The proof of Theorem 2 is a per-token conservation argument at the contemporaneous sqrt-price. The donor/absorber partition is then discussed via Lemma 1, which gives a sufficient condition under which a non-trivial partition arises with positive probability, and asymptotically with high probability under additional mixing assumptions.

***Proof of Theorem 2***

Consider a rebalance of pool $P_k$ at time $\tau$, with contemporaneous sqrt-price $s = s(\tau)$, in the swap-free regime $\sigma_x = \sigma_y = 0$. If a token $T_i \notin P_k$, then the dust-ledger entry $D_{\text{pool},i}$ is unchanged by (5), and no position outside $P_k$ moves any amount of $T_i$. Hence the per-token total $M_i$ defined in (6) is unchanged.

Now suppose $T_i \in P_k$, and write $T_i = T_a$ without loss of generality. By the V3 amount equations and Theorem 1, position $k$'s amount of $T_a$ jumps from

$$L^{(k)} a_x(s, s_a, s_b) = x_w$$

to

$$L^{(k)}_{\text{new}} g_x = \hat{x} - x'_{\text{dust}},$$

while the corresponding dust-ledger entry jumps from $x_{\text{dust}}$ to $x'_{\text{dust}}$. In the swap-free regime $\hat{x} = x_w + x_{\text{dust}}$, so

$$L^{(k)}_{\text{new}} g_x + x'_{\text{dust}} = (\hat{x} - x'_{\text{dust}}) + x'_{\text{dust}} = \hat{x} = x_w + x_{\text{dust}} = L^{(k)} a_x(s, s_a, s_b) + x_{\text{dust}}.$$

Thus the total allocation of token $T_a$ across the moving position and the dust ledger is preserved across the rebalance at price $s$. All other positions containing $T_a$, evaluated at the same contemporaneous sqrt-price, are unchanged by this event. Summing over positions and adding the dust ledger therefore gives

$$M_{T_a}(\tau^+) = M_{T_a}(\tau^-).$$

The same argument applies to the other token in the rebalanced pool. Summing over all tokens yields per-token conservation at the event, and multiplying by the fixed reference prices $p_i$ gives invariance of $V_{\text{total}}$.

The $m = 1$ corner cases reduce to the corresponding single-pool statements of the prior paper: under $S = 0$, the lone position's residuals accumulate in dust, recovering Theorem 3; under non-zero $S$, the swap-corrected single-pool case yields the in-expectation convergence regime of Theorem 2. ■

***Lemma: a sufficient condition for non-trivial partition***

**Lemma 1** (Positive-probability non-trivial partition under heterogeneity)**.** *Under the assumptions of Theorem 2, suppose further that*

*(P1) per-event slippage $(\sigma_x, \sigma_y)$ is conditionally mean-zero with strictly positive conditional variance;*

*(P2) rebalance arrival rates $\{\lambda_k\}_{k=1}^m$ are heterogeneous;*

*(P3) position ranges $\{[s_a^{(k)}, s_b^{(k)}]\}_{k=1}^m$ are not all coincident.*

*Then, for any window $[t_0, t_1]$ of positive length in which every position rebalances at least twice, the cumulative net impacts*

$$\Big\{ \sum_{\tau \in [t_0, t_1]} \Delta V_k(\tau) \Big\}_{k=1}^{m}$$

*have a non-degenerate joint law. In particular, a non-trivial donor/absorber partition occurs with strictly positive probability. Under additional mixing assumptions on the dust process, that probability tends to one as $T = t_1 - t_0 \to \infty$.*

*Proof sketch.* By Theorem 2, the sum of cumulative position impacts plus the cumulative dust-value change is zero in the window. Under (P1), $\Delta V_k$ at each rebalance has a slippage-driven random component with strictly positive variance, conditional on the pre-event state. Under (P2)–(P3), the per-event $\Delta V_k$ distributions are not exchangeable across positions; positions with higher arrival rates accumulate more variance, and positions with different ranges absorb dust at different rates given the same standing balance. The vector of cumulative $(\Delta V_1, \dots, \Delta V_m)$ over the window therefore has non-degenerate support rather than collapsing onto a single deterministic sign pattern. Combined with the conservation constraint $\sum_k \Delta V_k + \Delta V_{\text{dust}} = 0$, this gives a strictly positive probability of mixed-sign cumulative impacts at any non-trivial window length. Under additional ergodic or mixing assumptions on the dust process, repeated accumulation over long windows makes such mixed-sign partitions asymptotically typical. ∎

The lemma's failure mode (partition triviality) requires zero slippage variance or perfectly homogeneous arrivals and ranges across positions; the V9 dataset has neither, and the partition is observed in eight of eight multi-pool portfolios.

## Appendix F: Proof of Propositions 2 (Conditional propagation of chaos) and 3 (Closed-form one-spoke marginal)

**Status of the argument.** Proposition 2 is stated as a conditional theorem: assumption (A4), namely uniform-in-$N$ geometric ergodicity of the hub-marginal chain, is the only step in the proof not derived directly from the model definition. The remaining steps are rigorous under (A1′)–(A3) and (A5). Proposition 3 is then a leading-order consequence of Proposition 2 in the small-$w$ regime, valid under arbitrary slippage law $\nu_\varepsilon$; the half-normal mixture (Corollary 1) is the symmetric-Gaussian special case.

**Empirical anchoring of the assumption block.** On V9 chain data, the assumption block can be calibrated and stress-tested. (i) The Pool.Swap routing conservation $\sigma_y^{\text{raw}} = -s^2 \sigma_x^{\text{raw}}$ collapses the swap-correction noise to a single one-dimensional residual $\varepsilon$ ($\text{corr}(\varepsilon_1, \varepsilon_2) = 1.0000$ across 942 V9 substantial-swap events), whose distribution under V9 is asymmetric ($\mathbb{P}(\varepsilon < 0) = 0.26$ globally, 0.06–0.39 per pool) and heavy-tailed (excess kurtosis $\in [73, 451]$). (ii) Sign and magnitude conditioning of $\varepsilon$ on $\sigma_x^*$ together close approximately 71% of the gap between the i.i.d.-$\varepsilon$ baseline and the V9 spoke median; the corrected closed form (Proposition 3) matches the simulator's atom mass to within 0.001 at the kernel level. (iii) The leading-order approximation $|Lw + b - h| \approx Lw$ holds in the typical V9 event (median dust-to-position ratio 0.55%); state accumulation is not the dominant source of any production gap. (iv) The factorisation diagnostic $|\rho| \cdot \sqrt{\text{events}/N}$ is approximately constant in $N$ (max/min $\leq 1.34$) under three independent slippage-bootstrap kernels with materially different marginal shapes, consistent with asymptotic factorisation across the $\nu_\varepsilon$ family. (v) A topology-aware replay on V9 portfolio graphs reveals heterogeneous production behaviour outside the homogeneous symmetric leading-order closure: per-pool LP-asymmetry, per-pool $\nu_\varepsilon$, and joint $(\sigma_x^*, \varepsilon, \text{state})$ dynamics. The full V9 spoke-marginal law is therefore a graph- and geometry-

dependent extension question rather than a within-scope refinement of the present theorem (Section 9.4).

***The finite-$N$ Markov chain***

The state of the homogeneous hub-spoke system is $X^N = (H, B_1, \dots, B_N) \in \mathbb{R}_{\geq 0}^{N+1}$, where $H = D_{\text{pool},T_0}$ and $B_k = D_{\text{pool},T_k}$. Each pool $P_k$ fires at rate $\lambda$. When $P_k$ fires, draw $\delta \sim \text{Unif}(\{\pm w/2\})$ and $\varepsilon \sim \nu_\varepsilon$, then update $(H, B_k)$ via the kernel (7) and leave $B_j$ $(j \neq k)$ unchanged. The full chain is exchangeable in its spoke coordinates.

***Per-event jump on a homogeneous hub-spoke pool***

Consider a rebalance of pool $P_k = (T_0, T_k)$ in the homogeneous hub-spoke topology, with $h = D_{\text{pool},T_0}$ and $b_k = D_{\text{pool},T_k}$ pre-event. Without loss of generality consider the right-shift case $\delta = +w/2$ (the left-shift case is symmetric in the role of hub and spoke). At the boundary trigger, the position withdraws $x_w = 0$ and

$$y_w = L\,(s_b - s_a) = L\,s_b\,(1 - e^{-w}) \approx Lw + O(w^2).$$

The new range re-centres around $s = s_b$, giving per-unit-$L$ token amounts

$$g_x \approx \frac{w}{2s}, \qquad g_y \approx \frac{ws}{2}.$$

With swap routed at the spot price, so that $\sigma_y = -s^2\sigma_x$, the perfect-routing solution is

$$\sigma_x^* = -\frac{Lw + b_k - h}{2} + O(w^2),$$

which yields post-swap mint inputs at the LP-locus of the form

$$\hat{x} = \hat{y} = \frac{h + Lw + b_k}{2}$$

to leading order.

***Single-$\varepsilon$ slippage perturbation and slack-side leftover***

The Pool.Swap routing conservation

$$\sigma_y^{\text{raw}} = -s^2\,\sigma_x^{\text{raw}}$$

collapses the swap noise to a single residual $\varepsilon$: write

$$\sigma_x = \sigma_x^*(1 + \varepsilon), \qquad \sigma_y = \sigma_y^*(1 + \varepsilon), \qquad \varepsilon \sim \nu_\varepsilon.$$

The post-swap mint inputs are then

$$\begin{aligned} \hat{x} &= \tfrac{1}{2}(h + Lw + b_k) + \tfrac{1}{2}(Lw + b_k - h)\,\varepsilon, \\ \hat{y} &= \tfrac{1}{2}(h + Lw + b_k) - \tfrac{1}{2}(Lw + b_k - h)\,\varepsilon \cdot s_h^2, \end{aligned}$$

where $s_h^2 := s^2 \cdot 10^{d_0 - d_1}$ is the swap-routing slope in the human-scaled token frame, equal to 1 to leading order in the homogeneous symmetric regime ($s \to 1$, equal token decimals). To leading order,

$$\hat{x} - \hat{y} \;=\; \tfrac{1}{2}(Lw + b_k - h)\,(1 + s_h^2)\,\varepsilon, \tag{17}$$

so the binding side at the mint is determined by $\operatorname{sgn}\big((Lw + b_k - h)\,\varepsilon\big)$: spoke binds (i.e. $b_k' = 0$) iff $(Lw + b_k - h)\,\varepsilon < 0$. Conditional on the binding side, the slack-side leftover is

$$\text{leftover} \;=\; |\hat{x} - \hat{y}|\,(1 + O(w)) \;=\; \tfrac{1}{2}\,|Lw + b_k - h|\,|\varepsilon|\,(1 + s_h^2)\,(1 + O(w)). \tag{18}$$

***Foster-Lyapunov drift for the full chain***

Let $\Phi^N(h, b_1, \dots, b_N) = h^2 + \sum_k b_k^2$. Under (A1′) and the kernel (7), the per-event drift on the firing pool $k$ is

$$\mathbb{E}[\,\Phi^N(X') - \Phi^N(X) \mid X\,] \;\le\; -\tfrac{1}{2}h^2 - \tfrac{1}{2}b_k^2 + C_w, \tag{19}$$

with $C_w = O((Lw)^2\,\sigma_{\text{slip}}^2)$, using $\mathbb{E}[\varepsilon^2] = \sigma_{\text{slip}}^2$ and the elementary bound $(a+b)^2 \le 2a^2 + 2b^2$. The drift is negative outside the compact set $\{h^2 + b_k^2 \le 2C_w\}$. Combined with (A1′)'s density-bounded-below assumption (a small set / petite set construction on a neighbourhood of the origin), Meyn–Tweedie's Theorem 13.0.1 gives:

**Lemma 2.** *Under (A1′)–(A3), the chain $X^N$ admits a unique stationary law $\pi^N$ with finite second moment, and the family $\{\pi^N_{(m)}\}_{N \ge m}$ is tight on $\mathbb{R}^m_{\ge 0}$ for every fixed $m$.*

***Asymptotic decoupling under (A4)***

For the spoke chain on $b_k$: the spoke is updated only when pool $P_k$ rebalances (rate $\lambda$), but between two consecutive spoke-events of $P_k$ the hub is touched on average $N$ times by other pools' rebalances. Provisionally suppose (A4): the hub-marginal chain in stationary spoke environment mixes in TV at rate $C\rho^n$ after $n$ hub events, uniformly in $N$. Consider a Markov chain $\bar{B}$ on $\mathbb{R}_{\ge 0}$ with kernel: at each step, draw $\bar{H} \sim \pi_H$ independently of past state, draw $(\delta, \varepsilon)$ as before, and apply (7) with $(\bar{B}, \bar{H})$ in place of $(B_k, H)$. The pair $(\mu, \pi_H)$ is the unique fixed point of the McKean–Vlasov system

$$\mu \;=\; \operatorname{Stat}(\widetilde{\mathcal{K}}_w \mid \bar{H} \sim \pi_H), \qquad \pi_H \;=\; \operatorname{Stat}(\widetilde{\mathcal{K}}_w \mid \bar{B} \sim \mu), \tag{20}$$

where $\operatorname{Stat}(\cdot)$ takes a kernel parametrised by an environment law to its unique stationary law. Uniqueness is by a contraction argument anchored not by the TV-Lipschitz constant $\Lambda$ but by the binding-side reset: at each event, with probability $q := \mathbb{P}_{\nu_\varepsilon}(\varepsilon < 0)$ the spoke chain resets to 0 (annihilating any difference between coupled chains). The contraction rate on $\mu$ is $\rho_B := 1 - q < 1$, giving uniqueness even when $\Lambda > 1$.

*Proof of Proposition 2, sketch.* Tag spoke $k = 1$. Couple $(B_1^N, \bar{B})$ via shared boundary triggers, slippage variables, and binding-side outcomes; the only difference is the hub state seen at each spoke event ($H_N(\tau_n^-)$ versus $\bar{H}_n \sim \pi_H$). By (A4), the TV distance between these distributions decays as $C\rho^N$ uniformly in $n$. By (A5), the coupled chains satisfy

$$\mathbb{E}\big[|B_1^N(\tau_n^+) - \bar{B}(\tau_n^+)|\big] \le \Lambda \cdot d_{\text{TV}}\big(\text{Law}(H_N(\tau_n^-)), \pi_H\big) + \rho_B \cdot \mathbb{E}\big[|B_1^N(\tau_{n-1}^+) - \bar{B}(\tau_{n-1}^+)|\big].$$

The unconditional contraction $\rho_B = 1 - q$ from the binding-side reset gives an exponentially small bound in $N$ at any fixed event count. The stationary fixed-point identification is by passing to the limit; uniqueness from (20). The same argument applied $m$-wise gives $\pi^N_{(m)} \to \mu^{\otimes m}$. ∎

The proof shows in particular that pairwise correlations decay as $|\rho^N_{ij}| = O(\rho^N)$ exponentially in $N$, modulo (A4)'s constants. This is stronger than the unconditional $O(1/N)$ bound discussed in the surrounding text and consistent with the simulator results of Table 4 and the V9 evidence of Table 11.

***Closed-form spoke marginal in the small-$w$ limit***

*Proof of Proposition 3, sketch.* In the small-$w$ regime, $|Lw + B_k - H| = Lw \cdot (1 + O(\sigma_{\text{slip}}))$, so the leftover in (7) reduces at leading order to

$$\frac{Lw}{2}\,|\varepsilon|\,(1 + s_h^2).$$

The binding side is determined by $\operatorname{sgn}\big((Lw + B_k - H)\,\varepsilon\big) \approx \operatorname{sgn}(\varepsilon)$, since $(Lw + B_k - H) > 0$ in the typical-state regime. The spoke therefore binds (giving $B'_k = 0$) iff $\varepsilon < 0$, with probability $q = \mathbb{P}_{\nu_\varepsilon}(\varepsilon < 0)$; otherwise $B'_k = \frac{Lw}{2}\,|\varepsilon|\,(1 + s_h^2)$ with $\varepsilon \geq 0$, contributing the conditional law $\rho_+$. ∎

*Proof of Corollary 1, sketch.* If $\nu_\varepsilon$ is symmetric mean-zero Gaussian with standard deviation $\sigma_{\text{slip}}$, then $q = \mathbb{P}(\varepsilon < 0) = 1/2$ and $|\varepsilon|$ is half-normal with scale $\sigma_{\text{slip}}$. Multiplying by $\frac{Lw}{2}(1 + s_h^2)$ gives $\rho_+ = \mathrm{HN}(\beta)$ with $\beta = Lw\,\sigma_{\text{slip}}\,(1+s_h^2)/2$, which reduces to $Lw\,\sigma_{\text{slip}}$ in the homogeneous symmetric limit ($s \to 1$, $d_0 = d_1$). ∎

***Open: uniform-in-$N$ rate and full McKean–Vlasov route***

The conditional theorem above establishes the asymptotic factorisation modulo (A4). Two unconditional routes to (A4) are open and form the natural next technical step in this programme.

*Andreis–Dai Pra–Fischer route.* The kernel $\widetilde{\mathcal{K}}_w$ depends on the empirical spoke measure $\nu^N := (1/N)\sum_k \delta_{B_k}$ through the firing-pool's $B_k$. The Lipschitz dependence (A5) matches the hypotheses of[10] §5 for piecewise deterministic McKean–Vlasov systems. Adapting their $W_1$-Wasserstein rate argument to the discrete-event formulation here would give a quantitative version of (A4) and hence a quantitative rate in Proposition 2.

*Loss-network route.* Lift Graham–Méléard's variation-norm theorem[8] from the integer-state-space loss network to the continuous-state-space dust ledger by approximation arguments. This is the route discussed in Section 6.2 but is technically heavier because it requires controlling the convergence of the discrete approximation uniformly in $N$.

Either route would convert the conditional Proposition 2 into an unconditional theorem with explicit rate. The Foster–Lyapunov drift (19) is the input both routes need from the model definition. ∎

# Appendix G: Methodology of the empirical validation

### *Hypotheses summary*

Table 13. Empirical hypotheses and where each is checked.

| ID | Hypothesis | Source | Result |
|---|---|---|---|
| H1 | Per-event jump matches on-chain dust credits | Thm 1 | Median 1.0000 across 31/31 groups (§7.1, §7.2) |
| H2 | Same equation across V6 to V9 | Thm 1 | ~ 34,500 events on four contracts (§7.1) |
| H3 | Connector-side correlation has predicted sign | Prop 1 | $T_0$ median +0.42, $T_1$ median −0.44, 37 groups (§7.3) |
| H4 | Multi-pool portfolios partition into donors and absorbers | Thm 2 | 8/8 V9 multi-pool portfolios (§7.3) |
| H5 | Distinct spokes asymptotically independent | Prop 2 | $\lvert\rho\rvert \leq 0.14$ on 12 V9 spoke pairs (§7.3) |
| H5b | Atom-mixture closed form for spoke marginal | Prop 3 | sim $\mathbb{P}(b{=}0){=}\mathbb{P}(\varepsilon{<}0)$ within 0.001; half-normal special case in 3 of 32 V9 spokes (§7.3) |
| H6 | Theorem 1 predicts forward | Thm 1 | 57–97% within 2×, capped by engine-tick observability (§7.2) |

### *Contracts and schemas*

The four PM contracts in the framework's class are PM V6 (active 1–5 March 2026), PM V7 (6–13 March), PM V8 (13–21 March), and PM V9 (21 March onward). V6 and V7 share a 2-tuple $(\text{depositor}, \text{token})$ dust mapping; V8 introduced the portfolio dimension giving the 3-tuple $(\text{depositor}, \text{portfolioId}, \text{token})$ that V9 inherits. The V8 to V9 changes were internal AERO-leak-fix logic that did not modify event signatures. The V9 lake indexes 126,339 `Rebalanced` events and 143,991 `DustCredited` events over the 50-day window from V9 deployment to the current cursor; V6, V7, and V8 are reconstructed from per-version diffusion logs.

### *Conservation-form prediction*

For each rebalance event, the conservation-form prediction takes $(x_w, y_w)$ from the rebalance transaction's Pool.Burn event, $(\sigma_x, \sigma_y)$ from its Pool.Swap event on the rebalance pool, $(m_0, m_1)$ from its Pool.Mint event, and the standing-balance vector from the per-(depositor, portfolio, token) chronological credit history. The conservation expression $(x'_{\text{dust}}, y'_{\text{dust}}) = (x_w + x_{\text{dust}} - \sigma_x - m_0,\ y_w + y_{\text{dust}} - \sigma_y - m_1)$ is then compared to the on-chain `DustCredited` amounts emitted at the rebalance. The credit-history reconstruction is verified at 99.83% agreement with the on-chain `dustBalance` mapping on 16,388 comparison rows.

### *Data sources*

The empirical validation in Section 7 uses a PM V9 event lake (all contract events on Base from deployment at block 43,630,898 onwards), an Aerodrome swap dataset for the top-50 CL pools by historical vote weight (each Pool.Swap event's `sqrtPriceX96`, `liquidity`, and `tick`), and the off-chain rebalancing engine's diffusion log (per-rebalance pre-event position value, token prices, engine-detected pre-burn tick, and new range geometry).

The prediction pipeline operates per (depositor, portfolio): load all rebalance events and dust credits; build per-pool sorted-tick lookups for tick recovery and per-(pool, transaction) swap-delta lookups for routing detection; walk events in chronological order, applying the closed-form jump of Theorem 1 and comparing to the actual on-chain credit; update the standing balance with the emitted credit before the next event.

Two auxiliary indexers support the validation. An admin-op indexer scans all 138,080 PM V9 transactions via the Subsquid Network HTTP gateway and confirms zero silent dust mutations (no `transferDust`, `mergePortfolios`, `reassignPortfolio`, `withdrawDust`, `forceSetState`, or `emergencyWithdraw` calls); the only 27 decoded admin operations are `depositDust` calls already captured as `DustCredited` events. A swap indexer filtering on the indexed sender topic records 73,378 PM V9-issued Pool.Swap events across all 52 PM V9-managed pools, confirming the routing fires on 74–96% of rebalance events and providing the swap-correction inputs $(\sigma_x, \sigma_y)$.

### *Simulator*

The hub-spoke simulator underlying Tables 3–5 and the kernel bootstraps of Section 7.3 is implemented in PyTorch with GPU vectorisation across $M = 32$ parallel chains, with a per-chain burn-in of $\max(5 \times 10^4, 10N)$ events before any state is recorded. The i.i.d. bootstrap draws $\varepsilon$ from the marginal empirical distribution of $\varepsilon_1$ across the 942 V9 substantial-swap events; the sign-conditional bootstrap stratifies by $\mathrm{sgn}(\sigma_x^*)$; the magnitude-conditional bootstrap additionally bins by $|\sigma_x^*|/x_{\mathrm{pre}}$ into four equal-frequency strata. Per-chain summaries are aggregated to medians with 95% percentile intervals across the $M$ chains; runs are deterministic under seed `20260510`.

## Appendix H: Foundry verification suite

The verification suite of Section 8 comprises 19 tests across 5 contracts. Eleven mock-pool tests are network-free and verify Theorem 1, Theorem 2, and Proposition 1 against the V3-exact tick math directly. Eight live fork tests, all pinned to Base block `43,175,000`, verify Theorem 1 against two unmodified V3 `NonfungiblePositionManager` contracts on Base.

### *Test–theorem mapping*

Table 14. Foundry test mapping.

| Test contract | Tests | Network | Verifies |
|---|---|---|---|
| `MasterEquationT1Mock` | 7 | none | Theorem 1 (eq:newdust) |
| `MasterEquationT2` | 2 | none | Theorem 2 (mass conservation) |
| `MasterEquationConnectorRule` | 2 | none | Proposition 1 (sign-keyed correlation) |
| `MasterEquationT1Fork` | 4 | Base fork | Theorem 1 against Aerodrome Slipstream NFPM |
| `MasterEquationT1ForkUniV3` | 4 | Base fork | Theorem 1 against Uniswap V3 NFPM |

The two fork tracks differ only in the underlying NFPM ABI (`tickSpacing` on Slipstream, `fee` on Uniswap V3); the closed form is the same on both. The reference PM wrapper under test (`ReferencePM`, `ReferencePMUniV3`) is a ~270-line Solidity contract that adds the depositor-keyed shared dust mapping; V3 mechanics are delegated to the unmodified underlying NFPM.

### *Reproduction*

The mock-pool tests run locally without network access:

```
forge test -vv --no-match-contract Fork
```

The full suite, including the eight live fork tests, runs with:

```
RPC_BASE_ALCHEMY=https://mainnet.base.org forge test -vv
```

Recorded `forge test -vv` output, mapped to the test–theorem table above, is provided in `foundry/PROOF_OUTPUT.md` and serves as the deterministic regression target.

***Theorem 1: per-event jump (mock)***

The seven mock-pool tests (`MasterEquationT1Mock`) verify Equations (3)–(4) on `MockCLPoolV2`, a Solidity-internal pool with the V3-exact `TickMath` constants from `v3-core`. They cover the swap-free zero-standing-dust case (recovering the canonical residual of[1] §6 within 5%), linear additivity on both bound and slack sides, the swap-corrected case both below and past the locus threshold (binding-side flip), the three positional sub-cases of the V3 amount functions, and multi-position recycle where one position's residual is consumed by another position's mint.

***Theorem 1: per-event jump (Slipstream fork)***

`MasterEquationT1Fork` runs four tests against the unmodified Aerodrome Slipstream `NonfungiblePositionManager`[2]. The closed-form prediction is computed at the contemporaneous `slot0` sqrt-price and compared to the dust credit `ReferencePM` returns; predictions match emitted credits to within 1 wei (token0) and 1 raw unit (token1) on every event:

Table 15. Theorem 1 prediction error against Aerodrome Slipstream NFPM.

| Test | Token | Predicted | Actual | Match |
|---|---|---|---|---|
| swap-free, single rebalance | WETH | 707 | 706 | –1 wei |
| | USDC | 249,980,634 | 249,980,633 | –1 raw |
| swap-corrected, single rebalance | – | – | – | one side ≤ tol |
| cross-pool absorption | WETH | 706 → 151 | 706 → 151 | exact (event-level) |
| 3-event sequence (event 0) | WETH | 70 | 69 | –1 wei |
| | USDC | 142,456,348 | 142,456,347 | –1 raw |
| 3-event sequence (event 1) | WETH | 595 | 594 | –1 wei |
| | USDC | 58,204,099 | 58,204,098 | –1 raw |
| 3-event sequence (event 2) | WETH | 705 | 704 | –1 wei |
| | USDC | 249,980,632 | 249,980,631 | –1 raw |

The single-wei discrepancy on each row is the integer-arithmetic floor of the V3 amount-function inversion (`getAmountsForLiquidity` of `getLiquidityForAmounts`); the closed form is recovered to that floor on every event.

***Theorem 1: per-event jump (Uniswap V3 fork)***

`MasterEquationT1ForkUniV3` runs the same four tests against the canonical Uniswap V3 `NonfungiblePositionManager`[3] on the WETH/USDC fee-500 pool, using `ReferencePMUniV3`. Predictions match emitted credits to the same 1-wei tolerance:

[2] `0x827922686190790b37229fd06084350E74485b72` on Base.
[3] `0x03a520b32C04BF3bEEf7BEb72E919cf822Ed34f1` on Base.

Table 16. Theorem 1 prediction error against Uniswap V3 NFPM (Base, WETH/USDC fee 500).

| Test | Token | Predicted | Actual | Match |
|---|---|---|---|---|
| swap-free, single rebalance | WETH | 214 | 213 | –1 wei |
| | USDC | 418,199,470 | 418,199,469 | –1 raw |
| 3-event sequence (event 0) | WETH | 293 | 292 | –1 wei |
| | USDC | 405,158,379 | 405,158,378 | –1 raw |
| 3-event sequence (event 1) | WETH | 369 | 368 | –1 wei |
| | USDC | 395,322,047 | 395,322,046 | –1 raw |
| 3-event sequence (event 2) | WETH | 212 | 211 | –1 wei |
| | USDC | 418,199,468 | 418,199,467 | –1 raw |

The Uniswap V3 cross-pool test threads the depositor's standing dust ledger between Pool 1 (WETH/USDC fee 500) and Pool 2 (WETH/cbBTC fee 500); the connector token's standing dust is consumed by the second pool's mint per Equation (5).

***Theorem 2: mass conservation***

`MasterEquationT2` verifies the conservation form of Theorem 2 on a two-position portfolio under a shared depositor-keyed dust ledger. Driven by 200 deterministically randomised swap-free rebalances at the GS canonical anchor sqrt-price, the per-token totals

$$\sum_k a_{k,0}(t) + D_{\text{pool},T_0}(t), \qquad \sum_k a_{k,1}(t) + D_{\text{pool},T_1}(t)$$

are preserved bit-exactly: `totalWETH` matches initial to the wei across $K = 200$ events; `totalUSDC` matches to the raw unit. A second test verifies that Equation (5) threads the dust ledger correctly across rebalances of distinct positions, with the closed-form (4) predicting the post-rebalance dust state on every event.

The donor/absorber decomposition of Theorem 2 is a population-level corollary verified empirically on V9 in Section 7.3 (eight of eight multi-pool portfolios) and is not directly Foundry-tested.

***Proposition 1: Connector Rule sign***

`MasterEquationConnectorRule` runs $K = 80$ controlled rebalances with i.i.d. $\pm 300$-tick signed displacements and computes

$$\sum_{i=1}^{K} \text{sign}(\delta_i) \cdot \text{sign}(D_{\text{non-conn}}(\tau_i^+) - D_{\text{conn}}(\tau_i^+))$$

on two configurations. With the connector token at the $T_0$ position, the sum is $+18$; at the $T_1$ position, the sum is $-4$. Both signs match the prediction of Proposition 1. The magnitude is bounded below in the proposition by the variance share of the connector-shared component in the per-pool displacement decomposition; the single-pool mock has no shared component, so the suite verifies the sign claim only. The empirical magnitude on V9 (Section 7.3) is in the range $\pm 0.42$ to $\pm 0.44$ per-portfolio Spearman.