LIQUIDITY GUIDE · RESEARCH PAPER · AUGUST 2026

Concentrated Liquidity on Orca SOL/USDC: A Meta-Analysis of 93 Days of Simulation

Generated: 2026-08-09 **Asset:** SOL/USDC · Orca CLMM · pool `Czfq3xZZ…u44zE` · 0.04% fee tier **Corpus:** 7 live simulation runs (2026-05-09 → 2026-08-09), 5 counterfactual studies, 1 real on-chain fee dataset **Total logged rebalance events:** 72 (parallel architecture) + 5 closed cycles (state-machine architecture) + 79 logged 1D out-of-range closes (telemetry) **Capital at risk in production:** $0. Every result below is paper. `EXECUTION_ENABLED=false` throughout.

Sources synthesized: `SIMULATION_ANALYSIS_FIRST_RUN.md` · `SIMULATION_ANALYSIS_2026-05-24.md` · `SIMULATION_ANALYSIS_2026-06-01.md` · `SIMULATION_ANALYSIS_2026-06-06.md` · `SIMULATION_ANALYSIS_2026-06-24.md` · `SIMULATION_ANALYSIS_2026-08-09.md` · `BAND_WIDTH_ANALYSIS_2026-08-09.md` · `CEILING_EXIT_ANALYSIS_2026-08-09.md` · `PIVOT_1YR_PROPOSAL_2026-08-09.md` · `COURSE_CORRECTION.md` · `Final.md` · `STATE_MACHINE.md` · source code in `app/`

Published by liquidityguide.com · Companion to the write-up *I Ran the Narrow-Range Orca Strategy for 93 Days*

> **Not investment advice.** This is a technical post-mortem of a simulated liquidity-provision system. No capital was ever deployed — every result reported here is paper, produced against live price and pool data with execution disabled throughout. Results cover a single asset, venue, and fee tier and are not shown to generalize. See §17, Limitations.

Abstract

Over 93 days the project ran seven distinct simulated liquidity-provision strategies against live SOL/USDC price and pool data, then validated a counterfactual engine against the production ledger to the penny and swept the policy space it could not reach live. The headline result is negative and unambiguous: **no configuration tested — across band widths from ±1.25% to ±90%, hold policies from 1-hour to never-close, and two complete architectural rewrites — reliably beat the trivial benchmark of holding a 50/50 SOL/USDC basket.** Two of two formal paper-run gates failed.

The reasons are now understood mechanically rather than statistically. Three findings carry the analysis:

1. **Rebalancing cost is ~97% realized impermanent loss and ~3% gas and swap fees.** Swap friction is 0.022% of principal per close; IL at a band edge is 0.31%–3.60%. The project spent months optimizing the 3%.
2. **The forced 50/50 reconstitution swap at each reopen accounts for 99.7% of all realized IL** in the most recent run — proven by FIFO-matching the five close-swaps (−$44.49) against the ledgered lifetime IL (−$44.35), a 0.3% residual. What the system logged as "impermanent loss" was the P&L of a buy-high/sell-low ladder the rebalance rule executed on itself, five times, mechanically.
3. **Concentration is a risk multiplier, not an edge.** Over 233 days spanning a bull-and-crash cycle, fee income varied only 3× across band widths (±5% earned $75.17, ±90% earned $24.45) while IL varied 4× and dominated the outcome. Only ±60% and ±90% bands produced positive excess return over 50/50 HODL in all four market

regimes tested. Nothing tighter did.

A fourth finding emerged from reconciling this corpus rather than from any single report, and is documented here for the first time (§5): **the parallel-architecture era's portfolio mark-to-market was a ratcheting figure that never debited realized losses.** `effective_principal_usd` in `app/rebalance_sim.py` is incremented by every fee compound and never decremented by IL, slippage, or gas, and each reopen resets position value to it. Correcting this turns the June 24 headline of *+2.1% while SOL fell 22.9%* into approximately *−13.7%*, and turns "the strategy beat HODL by $541" into "it trailed HODL by roughly $90." Every favorable verdict in the first three months rests on this figure.

The constructive result is that the profitable configuration is nearly the opposite of what was built: **open wide, centre on live price, never close at a boundary, and treat the fee stream as a ~7–9% annual carry rather than the 300%+ APR the dashboard implied.** The gap between those two numbers is most of what went wrong.

Table of Contents

---

1. Scope, provenance, and evidence grading

1.1 What this analysis covers

Seven live simulation runs against real SOL price feeds (LiveCoinWatch) and real Orca pool state (`/api/pool/stats` via the sidecar), spanning 2026-05-09 to 2026-08-09. Two architectures: a **parallel four-bucket simulator** (`app/rebalance_sim.py`, 1D/1W/1M/1Y positions held simultaneously) that ran from May through early July, and a **single-position state machine** (`app/strategy_sim.py`, `STATE_MACHINE.md`) that replaced it on 2026-07-10 and is the only architecture now considered live.

Plus five counterfactual studies executed with `analysis/2026-08-09/backtest.py`, an engine that re-implements the production CLMM math (`_calc_liquidity`, `_token_amounts`, `_position_value_usd`, `_swap_price_impact_frac`) and the exact production ledger convention, and which **reproduces the production realized-IL ledger to the penny on all five closed segments**.

1.2 Evidence grading

Not all results in this corpus deserve equal weight. Grading them is the first job of a meta-analysis.

GRADE	MEANING	APPLIES TO
A	Arithmetic identity or mechanism proof; not a statistical inference	Reconstitution-swap FIFO match (0.3% residual); break-even day economics; IL path-independence; ledger reconciliations
B	Adequate sample, single asset, real data, correct benchmark	Hold-duration stratification (66 events); band-width sweep across 4 regimes; reopen-policy walk-forward (17 windows); crash stress test (37 windows)
C	Small sample or contaminated by known defects	Individual run P&L before 2026-08-09; the 5-cycle state-machine ledger; cascade-savings estimates
D	Effectively 2–3 independent observations; hypothesis, not finding	Ceiling-exit rule sweep (18 heavily-overlapping windows, only 3 with the ceiling binding); the manual `rewards.db` baseline

Every conclusion in §8, §9, and §14 is grade A or B. Grade C and D material is reported for completeness and explicitly labelled where it appears.

1.3 A note on the benchmark

The single most consequential methodological failure of this project is that **it ran for three months without a benchmark.** Reports from 2026-05-13 through 2026-06-24 assess the strategy against zero, against its own compounded principal, or against gross fee APR. None of those is a decision-relevant comparison. The only correct question is *versus what you would have had if you had done nothing*, and for a two-token liquidity position that is a **50/50 SOL/USDC basket held over the same window.**

This benchmark first appears in the 2026-08-09 analysis. When it is applied retrospectively (§5), several favorable historical verdicts invert. Anyone repeating this work should compute excess return over 50/50 HODL from the first hour of the first run.

2. The experimental register

ID	WINDOW	DAYS	ARCHITECTURE	CAPITAL	KEY VARIABLE UNDER TEST	OUTCOME	GRADE
E1	2026-05-09 → 05-13	4.0	Parallel (1D, 1W live)	$3,000/slot	Baseline viability	+$26.83 realized — **but the cost model contained no IL term** (§4.1)	C
E2	2026-05-18 → 05-24	6.0	Parallel, 4 slots	$1,000/slot	Same, with IL accounting added	1D closed −$25.38; **5 logic bugs found**	C (contaminated)
E3	2026-05-21 → 06-01	10.6	Parallel, post-bugfix	$4,000	Do the bug fixes hold?	Realized +$12.86; 4 clean events averaging 50.1h hold, +$38.25	C
E4	2026-06-01 → 06-06	5.0	Parallel	$4,125	Behaviour in a −25% crash	Realized −$15.52; reported MTM −$662 (−16.1%)	B (mechanism)
E5	2026-05-21 → 06-24	34.0	Parallel	$4,000	Largest sample; 72 events	Realized **−$46.35**; reported MTM +$83 (+2.1%) — **invalid, see §5**	B (events) / F (equity)
E6	2026-07-18 → 07-26	~8.0	**State machine** (rewrite)	$1,000	Single position, 120s grace	**−3.56%** (−$35.57), daily net −$1.20. **Gate FAILED**	C
E7	2026-07-26 → 08-09	13.39	State machine + learned grace	$1,000	Learned grace (p75 = 933→1,048s)	**$992.41 (−0.76%)** vs $1,004.36 HODL = **−1.19%**. **Gate FAILED**	B

2.1 Counterfactual studies

ID	DESIGN	WINDOWS	QUESTION	GRADE
S1	Walk-forward, 17 overlapping 5-day windows, chop series	17	Passive one-sided reopen vs 50/50 recenter, at 5 widths	B
S2	43 days spanning the May–June crash, 37 overlapping 7-day windows, volume-scaled fees	37	Does *any* LP policy beat HODL in a trend?	B
S3	Hold-through (never close) across 4 real market regimes, 5 widths	4	What band width survives every regime?	B
S4	18 overlapping 120-day windows over the 2025 rally-and-crash, 18 exit rules	18 (≈2–3 independent)	Can the upside tail be timed?	D
S5	`rewards.db` — 565 real on-chain fee collections, 206 days, manual operator	1	Did the human approach work, and why?	D

3. Market regimes covered

The corpus is unusually well-covered on regime diversity, which is what makes the negative result credible. SOL traded across a 4× range in the price series used.

REGIME	WINDOW	SOL MOVE	CHARACTER	USED IN
Compressed / trending	May 9–13	$93.70 → $94.50 (range $93.70–$96.30)	ATR ratio 0.26 — extreme compression; BB squeeze 22.8th pctile	E1
Oscillating chop	May 18–24	$81–$87, no net direction	ATR ratio 0.39–0.53; repeated boundary punches	E2, E3
Sustained crash	Jun 1–6	$80.89 → $60.60 (**−25.1%**)	regime_score = −1 every event after Jun 2	E4, S2
Crash + partial recovery	May 21 – Jun 24	$87.71 → $67.64 (**−22.9%**), low $60.61	53 of 66 rebalances under bearish regime	E5
Tight chop	Jul 26 – Aug 9	$70.99–$77.24 (8.8% band), net **+0.87%**	The friendliest regime for LP in the corpus	E7, S1
Full bull-and-crash cycle	Jun 2025 – Feb 2026	$141 → **$247.91** → **$78.15**	52.4% peak-to-trough drawdown, round trip	S3, S4, S5

Interpretation. The strategy was tested in one regime it should love (tight chop, E7/S1), one it should hate (sustained trend, E4/E5/S2), and one full cycle containing both (S3/S4). It failed its gate in the regime it should love. That is the strongest single piece of evidence in the corpus, because a chop-only carry trade that cannot make money in chop has no remaining habitat.

---

4. Results: the unified ledger

4.1 Realized P&L decomposition, all runs

Every run reduced to the same four terms. This is the table the project never had.

RUN	DAYS	FEES	REALIZED IL	FRICTION (GAS/SLIP/SWAP)	REALIZED NET	COVERAGE (FEES ÷ IL)	BREAK-EVEN COVERAGE[1]
E1	4.0	+$37.16	**not modelled**	−$10.32	+$26.83	—	—
E2	6.0	+$41.19	−$55.69	−$9.55	−$25.38	**0.74**	1.17
E3	10.6	+$119.12	−$88.22	−$18.62	+$12.86[2]	**1.35**	1.21
E4	5.0	+$239.37	−$212.86	≈−$42.03	−$15.52	**1.12**	1.20
E5	34.0	+$766.16	−$676.04	−$136.47	**−$46.35**	**1.13**	1.20
E6	~8.0	+$28.77	−$31.17	−$7.23	−$9.63[3]	**0.92**	1.23
E7	13.39	+$22.95	−$44.35	−$1.05	−$22.46	**0.52**	1.02

[1] *Break-even coverage* = 1 + friction ÷ IL. This is the fee/IL coverage ratio at which realized net crosses zero. It sits at **1.20 ± 0.03** for the whole parallel era — a remarkably stable constant, and a directly usable operating gate. [2] The June 1 report's own components sum to +$12.28 against a stated net of +$12.86; a $0.58 internal residual, carried forward unresolved. [3] E6's realized ledger nets −$9.63; the full −$35.57 equity change includes ≈−$24 of directional basket exposure attributed to over-trading.

Reading the table.

- **Coverage never sustainably exceeded its own break-even.** The only run above 1.20 was E3 (1.35), a 10.6-day window inside a larger run that ended at 1.13. Two of seven runs are below 1.0.
- **E1's +$26.83 is not a strategy result.** Git confirms `realized_il_usd` first entered `app/rebalance_sim.py` on 2026-05-18 (`0ef6e39`) — five days *after* the first run's report was written. The first analysis's cost table lists slippage $4.53, swap fee $0.60, gas $0.02, downtime $0.02, and no IL line at all. Its conclusion ("the 1D slot is the clear winner", "$428/month net") was drawn from a ledger missing the dominant cost term. SOL also moved 2.8% end-to-end over the window, so there was little IL to miss.
- **Coverage collapsed after the architecture rewrite**, from 1.13 (E5) to 0.52 (E7), even though every metric the rewrite targeted improved. This is the central paradox of the project and §4.3 resolves it.

4.2 The hold-duration law (E5, 66 normal rebalances) — Grade B

The cleanest empirical result in the corpus. Events stratified by hold duration, with two derived columns added here:

HOLD	EVENTS	WIN RATE	NET P&L	TOTAL IL	TOTAL FEES	IL / EVENT	FEES / EVENT	COVERAGE
< 2h	6	**0%**	−$82.18	−$92.83	+$23.51	$15.47	$3.92	**0.25**
2–6h	25	12%	−$179.20	−$249.65	+$121.73	$9.99	$4.87	**0.49**
6–12h	14	21%	+$23.32	−$126.80	+$178.89	$9.06	$12.78	**1.41**
12–24h	10	70%	+$84.47	−$93.09	+$197.94	$9.31	$19.79	**2.13**
> 24h	11	**82%**	+$109.40	−$113.14	+$243.59	$10.29	$22.14	**2.15**

The derived columns are the finding. IL per close is essentially **constant at $9–10** across every duration bucket, while fees per close rise **5.6×** from $3.92 to $22.14. Coverage is monotone in hold duration and crosses 1.0 between the 2–6h and 6–12h buckets — precisely where the win rate flips from 12% to 21% and then to 70%.

This is the mechanism stated in numbers: **IL is a per-close constant set by band geometry; fees are a rate integrated over time.** Profitability is therefore not a function of skill in choosing when to close. It is a function of how long you stay open relative to the fixed toll charged for closing.

The one exception is instructive. The `< 2h` bucket has *higher* IL per close ($15.47) than every other bucket. Those are whipsaws where price gapped well past the boundary before the stop fired. So the sharper statement is: **IL per close is bounded below by the band half-width and exceeds it when price moves violently past the edge before the close executes.**

4.3 Why fixing hold duration did not fix profitability

The June analyses correctly identified hold duration as the dominant variable, and the state-machine rewrite (2026-07-10) fixed it decisively:

METRIC	PARALLEL ERA (E5)	STATE MACHINE (E7)	CHANGE
Average hold	~6 h	**30.1 h**	5× longer
In-range rate	~60%	**86.8%**	+27pp
Cycles per day	2.1	**0.37**	6× fewer
Fee/IL coverage	1.13	**0.52**	**halved**

Every metric the rewrite was built to improve, it improved. Profitability got worse. The resolution: **fewer, longer holds mean fewer but larger closes, at wider boundaries — more dollars crossing the spread per close.** The rewrite removed the whipsaw layer and exposed the layer underneath, which was the reopen swap itself. It was not a regression; it was a successful experiment whose success revealed the real problem.

4.4 The root cause, proven arithmetically (E7) — Grade A

At every close the machine liquidates to one token (an out-of-range CLMM position is 100% one side by construction), then **swaps back to ~50/50 to open a new centred band.** That swap is a discrete trade executed at a range boundary, and a range boundary is by definition a local extreme.

The five realized swaps:

#	ACTION	SIZE	PRICE
1	BUY 6.5600 SOL	$501.56	**$76.46**
2	SELL 6.6588 SOL	$484.28	**$72.73**
3	BUY 6.5618 SOL	$486.28	**$74.11**
4	SELL 6.6295 SOL	$471.24	**$71.08**
5	BUY 6.4999 SOL	$475.15	**$73.10**

Average buy $74.56, average sell $71.91. **The machine bought 3.69% above where it sold, five times, mechanically.** FIFO-matching the round trips:

```
6.5600 SOL  bought $76.46 → sold $72.73  =  −$24.46
0.0987 SOL  bought $74.11 → sold $72.73  =   −$0.14
6.4630 SOL  bought $74.11 → sold $71.08  =  −$19.55
0.1664 SOL  bought $73.10 → sold $71.08  =   −$0.34
                          matched total  =  −$44.49
              ledgered lifetime IL       =  −$44.35
                          residual       =   +$0.14  (0.3%)
```

This is an identity, not a correlation. The number the system reported as "impermanent loss" and the P&L of the buy-high/sell-low ladder are the same dollars. Price ended the window **+0.87%**; the strategy finished 4.2% behind holding, entirely inside these five swaps.

It self-reinforces because the direction of the swap is determined by the direction of the exit: exit *above* the band → 100% USDC → the reopen buys SOL at the high. Exit *below* → 100% SOL → the reopen sells SOL at the low. In an oscillating market — the regime LP is supposed to want — **every oscillation costs one buy-high plus one sell-low.**

4.5 Why waiting longer does not help — Grade A

The obvious response is "close later." It was tested directly. Shadow watchers (§12.1) recorded the true recovery time after every force close, solving the censoring problem:

#	BAND	EXIT SIDE	IL @ CLOSE	RECOVERED	IL @ RECOVERY	FORGONE
1	74.29–76.16	above	−$5.12	never	—	—
2	72.76–79.01	below	−$14.81	+2 m	−$16.63	−$1.82
3	71.45–74.04	above	−$4.64	+16 m	−$4.25	+$0.39
4	71.44–76.35	below	−$12.36	+94 m	−$10.05	+$2.31
5	69.33–73.08	above	−$7.42	+1 m	−$6.28	+$1.14
					Total	**+$2.03**

Perfect grace timing — a machine with an oracle for recovery times — is worth $2.03 across five cycles. Recovery happens *at the boundary you just left*, so the price you would swap at barely changes. The learned-grace machinery is correct and functioning; it is simply a second-order parameter. This closed off an entire branch of the search space and is one of the most valuable negative results in the corpus.

---

5. Reconciliation: the ratcheting-principal defect

This section documents a defect found while reconciling the corpus for this meta-analysis. It is not reported in any prior document, and it re-grades the project's most favorable historical result.

5.1 The defect

In `app/rebalance_sim.py` (the parallel architecture, live May–early July), `effective_principal_usd` is:

- **incremented** on every fee compound (line 555) and by `add_to_principal` (line 1018);
- **reset** to base principal on initialization (lines 462, 482, 987);
- **never decremented** — by realized IL, by slippage, by swap fees, or by gas.

At every reopen (line 734) the new position's capital is taken from it, and position value is assigned directly from it:

```
new_capital = state.effective_principal_usd or self._base_principal
new_L = _calc_liquidity(new_capital, price, actual_lower, actual_upper)
...
state.current_position_value_usd = new_capital
state.unrealized_il_usd = 0.0
```

There is no bankroll concept anywhere in the file. The per-event ledger (`fees_earned_usd`, `realized_il_usd`, costs, `net_usd`) was maintained correctly and separately — **but it fed nothing.** Each rebalance reset the position to *initial capital plus every fee ever compounded*, with every realized loss discarded.

This is not the same as the already-fixed "Bug 6". Bug 6 (fixed 2026-06-25, `rebalance_sim.py:936`) was a *display* double-count: the portfolio total added `total_compounded` on top of `current_position_value_usd`, showing +18.7% when the file's own figure was +2.1%. That fix corrected the addition. It did not touch the ratchet described here, which is upstream of it: the base figure that Bug 6 was double-counting was itself already unable to fall from trading losses. The June 24 report's +2.1% is the *post-Bug-6* number, and it is still overstated by ≈$631.

5.2 The magnitude

For the 1D bucket at 2026-06-24, the June report states: MTM $1,585.08, effective principal $1,595, compounds 595 (= $595), realized net over 34 days −$46.35.

```
Reported "MTM":     $1,000 initial + $595 compounds              = $1,595   (drifted to $1,585.08)
True equity:        $1,000 + fees $766.16 − IL $676.04 − friction $136.47 = $953.65
Overstatement:                                                     ≈ $631
```

The two figures differ by exactly IL + friction − uncompounded fee remainder: $676.04 + $136.47 − ($766.16 − $595.00) = **$641**, consistent to rounding with the observed gap. The arithmetic closes.

The three passive buckets (1W, 1M, 1Y) never reopened, so their liquidity was never reset and **their reported values are honest.** Only the actively rebalancing bucket ratchets — which is to say, only the bucket whose performance was being evaluated.

5.3 The corrected result

BUCKET	REPORTED MTM (JUN 24)	CORRECTED	BASIS
1D	$1,585.08	**$953.65**	$1,000 + realized net (−$46.35)
1W	$801.65	$801.65	never reopened — honest
1M	$827.98	$827.98	never reopened — honest
1Y	$868.49	$868.49	never reopened — honest
Total on $4,000	**$4,083 (+2.1%)**	**≈$3,451.77 (−13.7%)**	

Against the correct benchmark — 50/50 HODL with SOL at −22.9% over the window — the basket would be worth **$3,542 (−11.45%)**.

> ***The 34-day parallel-sim run did not beat holding by $541. It trailed holding by roughly $90 (−2.3%).***

This is entirely consistent with everything measured afterwards, and it removes the last apparently-successful result from the record. **Every one of the seven runs either failed its benchmark or, in E1's case, was measured without the dominant cost term.**

5.4 Why this matters more than the dollars

The state machine fixed this architecturally: `app/strategy_sim.py` maintains a real `_bankroll_usd`, updates it at close (`line 748: bankroll = exit_pos_value + fees − costs`) and funds the next open from it (`line 529: capital = self._bankroll_usd`). Losses compound forward correctly. That is a genuine and underappreciated improvement from the rewrite.

But for two months, **the number on the dashboard was structurally incapable of going down from trading losses.** Decisions were made against it. This is the most important governance lesson in the corpus: a plausible-looking wrong number is more dangerous than a crash, because a crash stops you and a wrong number does not.

---

6. What made money

Honesty requires listing these, because several are real and reproducible.

#	RESULT	MAGNITUDE	WINDOW	GRADE	CAVEAT
1	**Hold-through ±90%, never close, May–Jun crash**	**+24.3% annualized excess vs HODL**	43 d, SOL −13%	B	Wide band, zero closes
2	**Hold-through ±10% in tight chop**	**+51.8% annualized excess**	13 d, SOL +1%	B	Highest number in the corpus — and the most regime-dependent
3	**Hold-through ±60%, 2025 bull-and-crash**	+1.5% annualized excess	233 d, SOL −37.5%	B	The only tight-ish width positive through a full cycle
4	**Passive one-sided reopen, ±2.0%, on the E7 window**	$1,055.63 vs $992.41 production (**+$63.22**)	13.4 d	B	Counterfactual, in-sample on fee calibration
5	**Passive one-sided reopen, ±2.5%**	$1,051.50 (**+$59.09** vs production, +$47.13 vs HODL)	13.4 d	B	Wins 17/17 walk-forward chop windows
6	**Never close, ±6% band**	$1,035.28 (+$30.92 vs HODL)	13.4 d	B	Zero closes, zero IL
7	**> 24h holds in production**	+$109.40 across 11 events, **82% win rate**	E5	B	The strategy working as designed
8	**The 53 non-catastrophic events of E5**	+$128.93	34 d	B	i.e. profitable once whipsaw is excluded
9	**Four clean post-bugfix events**	+$38.25, avg hold 50.1h	E3	C	Best clean production sample ever recorded
10	**Correct band centring (counterfactual)**	+$42.37 over production on the same policy	E7	B	Free — production opened off-centre on 5 of 5 opens

The pattern across all ten: every profitable result is either *wide*, *never-closing*, *long-held*, or *swap-free*. Not one of them comes from better prediction, better timing, or a tighter range. **There is no timing edge anywhere in 93 days of data.**

6.1 The single best-performing configuration tested

Over the manual-operator window (206 days, SOL −5.1%, 50/50 HODL = $974.45), hold-through with no rebalancing at all:

BAND	FINAL	VS HODL	FEES	IL	IN-RANGE
±10%	$1,037.80	**+$63.36**	$68.96	−$6.85	30.8%
±25%	$1,020.87	+$46.43	$46.20	−$2.80	51.4%
±40%	$1,018.00	+$43.56	$40.58	−$1.76	73.7%
±60%	$1,013.07	+$38.62	$32.78	−$1.16	92.1%
±90%	$1,005.91	+$31.46	$21.25	−$0.70	100.0%

Note that **in-range percentage is inversely correlated with return here** — ±10% was in range only 30.8% of the time and returned the most. This single table should end the use of in-range% as a success metric permanently.

But note also the fragility: hold the *same ±10% position four weeks longer*, to 2026-02-15 with SOL at −37.1%, and +$63.36 becomes **−$92.19**. Fees moved $6; IL moved **26×**. Identical management, opposite outcome, decided entirely by where price happened to stop.

---

7. What lost money, ranked

RANK	LOSS SOURCE	DOLLARS	RUN	NATURE	FIXED?
1	**Short-hold whipsaw** (13 events < 3h)	**−$173.11**	E5	Rule design	Yes — state machine (holds 5× longer)
2	**Passive stranded buckets** (1W + 1M unrealized)	**−$370**	E5	Architecture	Yes — parallel buckets abolished
3	**Cascade clusters** (4 identified periods)	−$82.76	E5	Rule design	Partially (cascade/wide-hold)
4	**Bug 5 — spike-guard override on out-of-range rebalance**	≈−$46.80	E2	Logic bug	Yes
5	**Reconstitution swap ladder**	**−$44.35 of −$44.49** (99.7% of all IL)	E7	**Rule design — the core defect**	**No — proposed, not shipped**
6	**Compounded-fee ledger leak**	−$29.89	E7	Accounting	Yes (code); persisted state still $21.00 short
7	**Voluntary contractions**	−$2.17 over 6 events, **0 wins**	E2/E3	Rule design	Yes — removed

7.1 The two worst individual events on record

EVENT	WHEN	HOLD	FEES	IL	NET	WHAT HAPPENED
E5 #33	Jun 5 19:09	1.9 h	+$11.62	**−$35.37**	**−$25.96**	5.1% price swing in under 2 hours; exit token0 was 18.63 SOL, all in a falling asset
E5 #71	Jun 24 17:42	2.0 h	+$2.39	−$21.44	**−$21.34**	$2.65 (3.9%) drop in 2 hours, immediately after a +$20.79 33-hour hold

Event #71 is the cleanest illustration of the whole problem: **a 33-hour profitable hold (+$20.79) was erased and then some by two rebalances over the following 7 hours (−$21.34, −$15.46).** Net over the cluster: −$16.01. The system's response to a large profitable close was to immediately make three bad trades.

7.2 The most misleading loss

The June 1–6 crash (E4) produced the highest fee APR the system ever reported: **+424% gross annualized**, with a single 9.8-hour position earning $55.81 (a 1,205% annualized rate). Over the same five days the reported portfolio fell **$662 (−16.1%)**.

This is not a bug and not specific to this system. It is structural: **price volatility drives pool trading volume, volume drives fee income, and the fee APR formula does not subtract the value of the asset you are holding while it falls.** Fee APR is at its most impressive exactly when an LP position is at its most destructive. Peak APR is a warning indicator.

The stress test makes this unavoidable. The worst single week in the corpus (Jun 1, SOL −19.7%) under a production-like policy:

```
final $779.72   vs 50/50 HODL $901.69   vs 100% SOL $803.38
fees +$26.71    IL −$141.73    costs $3.14    17 closes    91% in-range
```

91% in range, and it finished below simply buying and holding the crashing asset. Seventeen recentering swaps in seven days, each selling SOL lower than the last.

---

8. The mechanism: what is actually true about CLMM liquidity

Everything above reduces to seven statements. These are grade A — they follow from the payoff structure, not from this sample.

8.1 IL is path-independent for a position you never close

For a held position, divergence loss depends **only on entry price, exit price, and band geometry** — not on the route between them. The 2025 excursion from $141 to $247.91 and back to $133.86 cost essentially nothing (−1.58% at ±5%, −0.04% at ±90%) because it came back.

Closing converts path-independent IL into path-dependent IL. Every close crystallizes a boundary excursion permanently. This single sentence explains why the manual operator in `rewards.db` outperformed the bot's architecture: they collected fees on a schedule (565 collections, mean gap 8.8h) and never realized IL at a bad price.

8.2 Every close happens at a local extreme, by construction

You close because price left the band. Price left the band by moving to an edge. **The edge is the worst price in the band for the token you now hold 100% of.** There is no version of boundary-triggered rebalancing that does not transact at the local extreme; it is definitional.

8.3 IL at the edge is roughly linear in width; fee rate is roughly inverse; break-even is roughly quadratic

Calibrated to the live pool ($14,972/24h fees, $26.36M TVL, pool L 1.265e15):

BAND ±W	IL AT EDGE (% PRINCIPAL)	FEES/DAY IN RANGE (%)	BREAK-EVEN DAYS	SWAP COST (%)	TOTAL BE DAYS
1.25%	−0.312%	0.876%	0.36	0.022%	**0.38**
2.00%	−0.497%	0.568%	0.88	0.022%	**0.91**
2.50%	−0.621%	0.467%	1.33	0.022%	**1.38**
3.00%	−0.744%	0.399%	1.86	0.022%	**1.92**
4.00%	−0.990%	0.316%	3.13	0.022%	**3.20**
6.00%	−1.477%	0.234%	6.31	0.022%	**6.41**
10.00%	−2.435%	0.141%	17.21	0.022%	**17.37**
15.00%	−3.600%	0.095%	37.84	0.023%	**38.08**

Break-even days scale approximately as **w^2**: doubling width from 1.25% to 2.5% multiplies required in-range time by 3.7×; quadrupling from 2.5% to 10% multiplies it by 12.9×.

This table is the core economics of the venue and should be the first thing any operator computes. Treat it as economics, not as a trading trigger — tested against the five E7 cycles it predicted the sign correctly only 3 of 5 times, because realized fee rates varied **33×** between cycles (0.030%/day to 0.987%/day), which swamps the signal at n=5.

8.4 Rebalancing cost is ~97% IL and ~3% fees and gas

Swap friction is **0.022% of principal**. IL at the edge is **0.31%–3.60%**. Every hour spent on slippage modelling, priority-fee caps, Jupiter routing, and retry policy was spent on 3% of the cost of the thing it was optimizing. This is the most actionable finding in the document.

8.5 Fee income is far less width-sensitive than IL

Over 233 days spanning the 2025 bull-and-crash:

BAND	FINAL	VS HODL	FEES	IL	IN-RANGE
±5%	$708.69	−$103.91	$75.17	−$183.71	15.8%
±10%	$713.58	−$99.02	$70.65	−$178.82	27.9%
±15%	$728.89	−$83.71	$75.31	−$172.61	43.7%
±25%	$732.20	−$80.40	$52.91	−$155.53	51.6%
±40%	$780.07	−$32.53	$47.91	−$115.21	75.2%
±60%	$826.63	**+$14.03**	$38.61	−$75.93	93.4%
±90%	$871.37	**+$58.77**	$24.45	−$45.72	100.0%

Fees vary 3.1× across an 18× range of widths. IL varies 4.0× and is 2–7× larger in absolute terms. Narrowing the band buys a little more fee income and a lot more impermanent loss. Concentration is a risk multiplier, not an edge.

8.6 The payoff is short gamma and the tails are symmetric in USD but not in kind

A log-symmetric band has symmetric tails. For ±90% at spot $77.24 on $1,000:

EXIT PRICE	MOVE	SOL HELD	USDC	USD VALUE	VS HODL
$40.65 (floor)	−47.4%	**15.396**	0	$625.90	**−17.99%**
$60.00	−22.3%	9.647	283.92	$862.76	−2.89%
$77.24 (entry)	0.0%	6.473	500.00	$1,000.00	0.00%
$100.00	+29.5%	3.617	751.04	$1,112.73	−3.02%
$146.76 (ceiling)	+90.0%	**0.000**	1,189.20	$1,189.20	**−17.99%**
$200.00	+158.9%	0.000	1,189.20	$1,189.20	−33.74%

Identical −17.99% at both bounds — **but the floor pays you in SOL (15.396 vs 6.473 at entry, +138%) and the ceiling pays you in USDC.** Only one of those is a consolation, and only if SOL is what you wanted. Above the ceiling the position stops participating entirely and the gap to HODL grows without limit.

Most people designing these fixate on the downside. The upside tail is worse and has no consolation prize.

8.7 Fee APR is anti-correlated with your P&L

Volume spikes when price trends. Crash-period SOL volume ran a **$2,241M/day median versus $1,202M/day** in quiet periods — a 1.9× uplift arriving exactly when the position is being destroyed. Two consequences:

- Any backtest using a flat fee rate **understates LP income by ~1.9× in precisely the regime being stress-tested.** Always scale fees by volume. (The counterfactual engine does; this is why its crash results are credible.)
- The dashboard metric that looks best is the one that should alarm you.

---

9. The lever hierarchy: five orders of magnitude

Every tunable in the system, measured, on a common basis. **This is the single most useful table in the document.**

RANK	LEVER	MEASURED WORTH	SOURCE	EFFORT SPENT
1	**Whether to be in LP at all during a trend**	−1 to −2%/week typical, **−8.7%/week worst**	S2 (37 windows)	~none
2	**Band width, across regimes**	Annualized excess spread of **67.6pp** at ±10% (−15.8% to +51.8%) vs **24.1pp** at ±90%	S3 (4 regimes)	~none until Aug
3	**Whether you close at all**	Never-close ±6% beat production by **+$42.87** on 13.4 days	S1	none
4	**Reopen policy (passive vs 50/50 recenter)**	**+$59.09** per $1,000 over 13.4 days (+5.9%)	E7/S1	none — proposed, unshipped
5	**Band centring at open**	**+$42.37** per $1,000 over 13.4 days	E7	small (freshness guard, partial)
6	**Fixed ±2.5% vs the ATR-derived 1D/1W selector**	Selector added no measurable value: **0 of 79** logged 1D closes had atr_ratio ≥ 1.2 (max 0.89)	E7 telemetry	large
7	**Grace / stop-out timing (perfect foresight)**	**+$2.03** total across 5 cycles	E7 shadow watchers	**very large**
8	**Slippage & priority-fee optimization**	0.022% of principal per close ≈ **3%** of close cost	S3	**very large**
9	**Compounding cadence (daily vs monthly)**	**$0.52/year** per $1,000	PIVOT §0	moderate

Levers 1–3 were never systematically explored until the final week of the project. Levers 7–9 absorbed most of the engineering effort. The ratio between the top lever (−8.7%/week ≈ −$87 per $1,000 per week) and the bottom (compounding cadence, $0.52/year) is roughly **9,000×**.

This inversion is the project's defining failure, and it is not a failure of intelligence — every one of levers 7–9 is *visible*, *tunable*, and *satisfying to work on*, while levers 1–3 require admitting the strategy might be wrong. That is a general hazard, not a local one.

9.1 The negative results that closed branches

Equally valuable, because each eliminated a large search space:

FINDING	CONSEQUENCE
Perfect grace timing worth $2.03	Stop-out timing is exhausted. Do not tune further.
ATR gate never escalated in 79 closes	The volatility-based range selector is inert. Replace with a fixed width.
Kaufman Efficiency Ratio gate made **every** upside-exit rule worse, monotonically as the threshold tightened	By the time a trend is statistically confirmed, price is near the top of it. **Do not build the ER gate.** (This retracted a recommendation from the same day's earlier analysis.)
No exit rule beats hold-through on average; the only rule that helps when the ceiling binds (`convert to SOL at +20%`) needs a **74% hit rate** to break even and is the worst rule tested when it doesn't	**Do not build an upside exit rule.** The band absorbs the tail, or a delta hedge does. Timing does not.
Voluntary contractions: 6 events, 0 wins, −$2.17	Removed. "Take profits by tightening" is value-destroying by construction — you pay slippage for a fee-density gain you never hold long enough to collect.
Every LP policy loses to HODL in a trend (win rate 30–40% across all widths, all grace values, both reopen modes)	LP is a **chop-only** trade. Regime selection is the whole game, and no tested signal selects regime.

---

10. Defect catalogue and what it teaches

Thirteen defects found across 93 days. Classifying them is more instructive than listing them.

#	DEFECT	CLASS	COST / EFFECT	STATUS
1	Spike-guard override applied during out-of-range rebalances → 28–32% bands	Trading logic	≈−$46.80	Fixed 2026-05-24
2	Static velocity threshold (0.7) fired on normal noise at compressed ATR	Trading logic	Root trigger for #1	Fixed (dynamic, band-fraction scaled)
3	Contraction cooldown gated only on prior contractions, not on normal rebalances	Trading logic	Double contractions 1h apart	Fixed
4	Voluntary contraction fired without an in-range check	Trading logic	Bypassed displacement/streak gates	Fixed
5	Stale width-ratio denominator (`state.center` vs band midpoint)	Trading logic	Measured against a frozen reference	Fixed
6	`SLIPPAGE_BPS` set in env, code reads `SLIPPAGE_PCT` — **env var never read**	Config / silent no-op	Slippage modelled at a value nobody chose	Fixed
7	Container ran a **stale pre-rewrite build for 17 days**	Operational	17 days of data from dead code	Fixed 2026-07-18
8	Poll thread died on ENOSPC (2026-07-01), unguarded	Operational	Silent data loss	Fixed (try/except in `_poll_loop`)
9	**Compounded fees destroyed at close** — `_accrue_fees` raises `effective_principal_usd` without recomputing `liquidity_l` ; both consumers credited only the sub-threshold remainder	**Accounting**	**−$29.89** real ledger leak	Fixed 2026-08-09 + 5 regression tests; persisted state still $21.00 short
10	**`effective_principal_usd` never debited for realized loss; reopen resets position value to it**	**Accounting**	**≈$631 overstatement of 1D equity**; inverted the sign of the 34-day verdict	**Documented here (§5).** Architecturally resolved by the state machine's bankroll
11	`add_to_principal` grows `liquidity_l` without updating `entry_token0/entry_token1`	**Accounting**	Distorts **every subsequent IL calculation**	Open — blocks real compounding
12	Live-mode double-count hazard: `_accrue_fees` and `add_to_principal` would count the same dollars	**Accounting**	Would corrupt live ledger	Open — **gates `EXECUTION_ENABLED=true`**
13	`rewards.db` `sol_price_at_redemption` is a backfilled constant ($139.42 on 1,014 of 1,132 rows)	Data integrity	Understated real fee value by $33.45 (12%)	Corrected in analysis by repricing at real daily SOL

10.1 The pattern

- **5 of 13 are trading-logic bugs.** All were found within days, because they produced visibly absurd behaviour (30% bands, back-to-back contractions). Loud failures get fixed.
- **4 of 13 are accounting defects (#9–12).** These were the expensive ones. Not because their direct dollar cost was highest — #1 cost more than #9 — but because **they made the system unmeasurable.** Defect #10 sustained a false positive verdict for two months. Defect #9 made a −0.76% run look like −3.75%, a 5× error in the headline number, and it was *visible all along* as an internal contradiction: the event ledger summed to −$22.46 while the bankroll fell −$49.26. Any reconciliation of those two numbers would have caught it.

- **2 of 13 are operational (#7, #8)** and both are the same failure: **the system was not verifying that it was running the code it thought it was running.** Seventeen days of a stale build and a silently dead poll thread.
- **1 is upstream data integrity (#13)**, found only because someone checked a column's cardinality.

The lesson: build the reconciliation before the feature. Two independent computations of equity that must agree — event ledger vs bankroll — would have caught #9 and #10 immediately, and #10 would never have been written, because it cannot survive a reconciliation.

---

11. Metrics that misled, and metrics that worked

11.1 Metrics that actively inverted the signal

METRIC	WHAT IT SHOWED	WHAT WAS TRUE
Gross fee APR	+424% annualized	Portfolio −16.1% over the same 5 days
In-range %	91% during the worst week on record	$779.72 final vs $901.69 for 50/50 HODL — below even 100% SOL
Hold duration	Improved 5× (6h → 30.1h)	Fee/IL coverage *halved* (1.13 → 0.52)
`effective_principal_usd` as P&L baseline	1D grew $1,000 → $1,595	True equity $953.65 (§5)
Compound count	661 compounds = "$661 reinvested"	That is fee income re-labelled, not profit
Snapshot in-range rate	1W bucket: 92.53%	Historical probability of full-period containment: **0%** — SOL had never once stayed inside a 6.88% band for 7 consecutive days
Unrealized IL alone	1W: −$141.31	Total hole in that position: **−$315.68**. Price decline was invisible in the reporting

The 1W snapshot-vs-containment discrepancy deserves emphasis. **Snapshot in-range rate treats every poll as independent and therefore massively overstates position quality.** A position that is in range 92.5% of polls but breaks out at least once per period earns nothing like 92.5% of the fees, because the breakout forces a close. Use full-period containment probability.

11.2 Metrics that worked

METRIC	WHY IT WORKS	OPERATING GATE
Fee / IL coverage ratio	Directly compares the income to the dominant cost	**> 1.20** (empirically the break-even including friction, stable ±0.03 across the parallel era)
Excess return vs 50/50 HODL	The only benchmark that answers "should I have done this at all"	> 0 over a rolling 7 days
Realized IL per close	Isolates the toll from the income	< $2 per $1,000
Shadow-watcher recovery time	Solves right-censoring: measures what waiting *would* have earned, at zero risk	Feeds the learned grace p75
FIFO swap-ladder reconciliation	Turned a statistical suspicion into an identity (0.3% residual)	Run it on any strategy that swaps
Event ledger vs bankroll reconciliation	Two independent paths to equity that must agree	Must reconcile to $0.00

---

12. Where we went right

12.1 Shadow watchers — solving the censoring problem

After every force close, the machine kept watching the closed range with no capital at risk, recording whether and when price re-entered. This is a genuinely good piece of experimental design: **the naive approach (measure recovery only when you didn't close) is right-censored and systematically biased**, because you only observe recoveries in the cases where you happened to wait. Shadow watchers observe the counterfactual directly.

The payoff was a decisive negative result — perfect grace timing is worth $2.03 — which is exactly what good instrumentation is for. It cost a branch of the search space to build and saved an unbounded amount of further tuning.

12.2 The validated counterfactual engine

`analysis/2026-08-09/backtest.py` re-implements production CLMM math and **reproduces the production realized-IL ledger to the penny on all five closed segments.** This is the single highest-leverage artifact in the repository. It converted the project from "run it and see for two weeks" to "sweep the entire policy space in seconds against real price data."

It should have existed in week one. Every policy question answered in the final week — reopen policy, band width, exit rules, regime gates — was answerable from day one with this engine and the price history that was already being collected.

12.3 Volume-scaling the fee model

The engine scales fee income by observed SOL volume rather than assuming a flat rate. Without this, the crash stress test would have understated LP fee income by ~1.9× **in exactly the regime being stress-tested** — biasing the analysis toward the conclusion the author might have wanted. Getting this right is why "LP loses to HODL in trends" is credible rather than an artifact.

12.4 Withdrawing a recommendation on contradicting evidence

The `SIMULATION_ANALYSIS-2026-08-09.md` §5.4 recommendation to build a Kaufman Efficiency Ratio regime gate was **formally withdrawn the same day**, when the ceiling study found the same signal made every rule it touched worse, monotonically. The withdrawal is documented in-place with reasoning, rather than quietly dropped.

This is rare and it is the correct behaviour. The stated principle — "*if a regime filter is revisited it needs its own study, with the signal chosen before the rules rather than reused from a prior finding*" — is a direct guard against the multiple-comparisons failure that produces most bad quant results.

12.5 Refusing to go live

Through seven runs, thirteen defects, and two failed gates, `EXECUTION_ENABLED` stayed `false`. Two formal gates were defined in advance, both were failed, and both failures were accepted rather than rationalized. **Total capital lost across the entire project: $0.**

Given that defect #12 (the live-mode double-count) is still open and would have corrupted the live ledger from the first compound, this discipline was not merely prudent — it was load-bearing.

12.6 Correct architecture on the second attempt

The state machine introduced a real bankroll that carries losses forward (§5.4), a single position instead of four phantom parallel ones, learned rather than guessed parameters, and full ledger logging per close (no APR estimates). These are the right foundations. The strategy running on them was wrong, but the foundations are sound and reusable.

---

13. Where we went wrong

13.1 We optimized the 3%

Months of work on slippage modelling, `SLIPPAGE_BPS` / `SLIPPAGE_PCT` alignment, priority-fee hard caps, compute-unit limits, retry policy without fee escalation, gas cost constants, and a 0.75% slippage ceiling. **Total addressable cost: 0.022% of principal per close, roughly 3% of what a close actually costs.** The other 97% — realized IL at the boundary — was not examined until the final week.

The tell was available early: the first-run analysis noted "slippage dominates" (87.6% of the $5.17 rebalance cost) — but that cost model **had no IL term** (§4.1). A dominant-cost finding derived from a ledger missing the dominant cost set the optimization agenda for two months.

13.2 We never isolated a variable

Every run changed several things at once: bug fixes plus market regime, architecture rewrite plus grace learning plus escalation patches plus freshness guard. The first genuinely single-variable experiment in the project is the one *proposed* on 2026-08-09 and not yet run.

The consequence is stated explicitly in that analysis: "*this window's whole lesson is that the reopen rule was invisible for three months because it was never isolated.*"

13.3 We ran without a benchmark for three months

See §1.3 and §5.3. Without a HODL comparison, "profitable" and "unprofitable" are unanchored. The June 24 report's headline — +2.1% while SOL fell 22.9% — reads as a triumph and was in fact (once corrected) a −2.3% underperformance.

13.4 We accepted a target that was never sanity-checked

The founding goal was **$100/day on $3,000 principal** — an annual return of **1,217%**. At the observed net APR of 35–50% after IL drag, the arithmetic gives $3.29/day. This was first refuted on 2026-06-01, seven weeks and four analyses into the project.

A five-minute calculation at the outset — fee APR × principal ÷ 365, minus an IL estimate — would have reset the entire project's design constraints. Instead, an impossible target shaped the architecture toward tight, high-density, frequently-rebalanced ranges: **precisely the configuration that §8 shows is worst.** The target did not merely fail to be met; it actively caused the wrong strategy to be built.

13.5 Architecture drift

The specification was a single-position state machine. The code ran four parallel perpetual positions for roughly two months. This was not a subtle divergence — it changed what was being measured (four simultaneous bets vs one sequential one), and the parallel buckets generated the −$370 stranded-position loss that has no analogue in the specified design.

13.6 Complexity was the response to poor results

The chronological record of mechanisms added: spike guard → dynamic velocity threshold → band-ratio recovery signal → cascade detection → wide-hold mode → vol-ratio exit gate → 1D-stopout escalation patch → freshness guard → learned grace → efficiency-ratio regime gate.

Each was a patch on a rule whose core premise — close at the boundary and reopen centred — was wrong. The final analysis found that the ATR gate these mechanisms surrounded had never once fired on its own in 79 logged closes; the escalation patch was doing all the work. A large fraction of the codebase exists to decide *when to close*, and the best-performing configurations in the entire corpus never close at all.

13.7 We trusted the dashboard

Defects #9 and #10 both survived because the displayed number was never checked against an independent computation. In #9's case the contradiction was sitting in the data: the event ledger summed to −$22.46 and the bankroll had fallen −$49.26.

---

14. Practical guidance: searching for the optimal SOL/USDC position on Orca

This section is the deliverable. It is written for someone about to open a position, and every rule cites the evidence behind it.

Step 0 — Decide whether to provide liquidity at all

This is lever #1 and it dominates everything else. Grade B, S2 (37 windows).

An LP position is **short gamma and long delta**. It earns when price oscillates inside your band and loses when price trends through it. Across 37 overlapping 7-day windows spanning a real crash, **every** policy — every width, every grace value, both reopen modes — beat 50/50 HODL only 30–40% of the time in trending conditions.

Ask three questions and answer them honestly:

1. **Do I have a directional view?** If yes, express it directionally. Do not provide liquidity as a way of being long. At the ceiling of a ±90% band you hold zero SOL and are 18% behind HODL — and 34% behind if price runs 159%. "*If the conviction is that SOL goes to $200, do not provide liquidity — just hold SOL.*"
2. **Do I expect SOL to range more than it trends over my horizon?** LP is a bet on realized volatility being *choppy* rather than *directional*. Nobody in this corpus found a signal that predicts that. Efficiency ratio was tested and made things worse.
3. **What unit do I care about?** USD or SOL. The answer changes the correct band entirely (Step 3).

If you cannot answer 2 affirmatively, **hold the basket.** That is not a defeatist conclusion; it is the measured one.

Step 1 — Compute the break-even table for your pool, first

Before choosing anything else, build the table in §8.3 from live pool data: 24h fees, TVL, and liquidity at tick. It takes minutes and tells you the only thing that matters operationally: **how many days in range a given width must survive to pay for one close.**

At the pool state observed on 2026-08-09, a ±2.5% band needs **1.38 days in range** to break even on a single close. A ±10% band needs **17.37 days**. If your expected hold is shorter than the break-even, the position is a guaranteed loser regardless of anything else you do.

Recompute it when pool conditions change. Fee rates in the observed cycles varied **33×**.

Step 2 — Choose width by the worst regime, not the best

Grade B, S3. Annualized excess return over 50/50 HODL, hold-through (never close):

REGIME	±10%	±25%	±40%	±60%	±90%
2025 bull+crash (233d, SOL −37%)	−15.8%	−12.6%	−6.0%	**+1.5%**	**+8.8%**
Manual window (206d, SOL −5%)	+11.2%	+8.2%	+7.7%	**+6.8%**	**+5.6%**
May–Jun crash (43d, SOL −13%)	−14.2%	+12.5%	+14.1%	**+17.2%**	**+24.3%**
Recent chop (13d, SOL +1%)	**+51.8%**	+19.8%	+11.2%	+5.5%	+0.2%
Worst case (maximin)	**−15.8%**	**−12.6%**	**−6.0%**	**+1.5%**	**+0.2%**
Spread (max − min)	**67.6pp**	32.4pp	20.1pp	15.7pp	**24.1pp**

Only ±60% and ±90% are positive in all four regimes. Nothing tighter is.

The ±10% column is the trap. It produced the best single number in the corpus (+51.8%) and the worst (−15.8%), and the difference between them was **not skill — it was which direction SOL happened to move.** A 67.6pp spread on a decision you make once is not a strategy; it is a coin flip with extra steps.

- **Default: ±60%.** Best maximin, ~12% fee APR at observed pool rates.

- **If you cannot tolerate the tail: ±90%.** Better mean (−5.28% vs −6.81%) and better worst case (−20.11% vs −23.20%) across the 2025 cycle, at a cost of ~2.8pp of fee APR. ±120% is better still on risk but drops fee income to ~7%, where the carry stops justifying the operational complexity.
- **Only go tighter than ±25% if you have specifically verified that your expected hold exceeds the break-even from Step 1, with margin.**

Step 3 — Set the geometry: log-symmetric, centred on the live price

Two rules, both grade A/B.

Log-symmetric. A band of `[p/k, p·k]` has symmetric percentage tails. At ±90% the loss versus HODL is −17.99% at *both* bounds, which makes the risk trivially easy to reason about. Arithmetic-symmetric bands do not have this property.

Centred on the live price at the moment of opening. Production bands were anchored to lagging SMAs and opened off-centre on **5 of 5** opens (`recentered_on_price: false` every time). Correcting this alone was worth **+$42.37 per $1,000 over 13.4 days** — free money, no forecast, no extra transaction.

Do not use a trailing-window band. The live "1Y model band" at spot $77.24 was $57.33–$171.29 — i.e. **−25.8% / +121.8%** — because spot happened to sit near the bottom of the trailing year's range. Entering it deposits **70% SOL / 30% USDC**. That is not a liquidity position; it is a leveraged directional bet with a fee coupon, and its ceiling case (−27.60% vs HODL) is the worst of every band tested.

A note on asymmetry. A test that appeared to show upward skew hurts was later found to be mis-parameterized: the "skew up" bands also had *tighter floors*, and the floor distance was driving the results (the widest-floor variant placed second overall). **Upward skew is untested, not disproven.** A correct test holds the floor fixed and varies only the ceiling.

Step 4 — Never close at a boundary

Grade A. This is the rule that would have saved the most money in this project.

Every close happens at a local extreme by construction (§8.2). The reopen swap that follows is therefore a systematic buy-high or sell-low, and in the audited run it accounted for **99.7% of all realized IL**.

Three acceptable policies, in order of preference:

1. **Never close.** Best-performing family in the corpus. IL becomes path-independent — the 2025 round trip from $141 to $248 and back cost 0.04% at ±90%. Out-of-range simply means fees stop accruing; that is the entire consequence. Collect fees on a schedule; do not touch the range.
2. **Passive one-sided reopen** — if you must run a tight, active band. At close, **do not swap.** Open a one-sided band on the side matching what you already hold: exited above (all USDC) → band strictly *below* price; exited below (all SOL) → band strictly *above* price. Zero swap, zero slippage, and the position becomes a resting limit order that accumulates the cheap asset — the exact inverse of recentering. It beat 50/50 recenter at **every** width on mean, median, worst case, and win rate, and won **17 of 17** walk-forward chop windows.

POLICY	MEAN	MEDIAN	WORST	BEST	WIN%
recenter ±2.5%	+1.14	+0.84	+0.27	+2.39	100%
passive ±2.5%	**+1.54**	+1.44	+0.54	+2.48	100%
recenter ±1.25%	+1.14	+1.29	−1.27	+2.91	76%
passive ±1.25%	**+2.03**	+1.71	−0.32	+4.47	**94%**

3. **Recentre only when price is near your original entry**, where IL ≈ 0. Never at a boundary, which is by construction the worst available point.

Step 5 — Do not build an exit rule for the upside

Grade D on the sample, grade A on the mechanism. Eighteen exit rules were tested across the 2025 rally-and-crash. **Every active rule lost to holding through on average.** Recentering upward was the worst systematic choice (−11.0 to −11.5% mean, **0% win rate**, negative in every window) — it is the same buy-high mechanism, in a different costume.

The one rule that helps when the ceiling actually binds (convert to 100% SOL at +20%) requires a **74% hit rate** to break even: +3.21% when right, −9.21% when wrong. No tested signal comes close. The efficiency ratio made every rule worse, monotonically as the threshold tightened — **by the time a trend is statistically confirmed, price is near the top of it.**

Timing is not the constraint. A +60% move took a *minimum* of 62 days to arrive across the 2025 cycle; there were always weeks of warning. **Decision quality is the constraint, and there is no evidence any implementable rule has it.**

If the ceiling tail is genuinely unacceptable, the correct instrument is a **delta hedge** (a short SOL perp), not an exit rule. That converts a bet on price into a bet on volume — which is the thing you actually believe you are being paid for. It is a substantially larger project and should be scoped on its own.

Step 6 — Compound because it is free, not because it is a lever

Grade A. At the live ±90% band's 10.73% fee APR on $1,000:

COMPOUNDING	VALUE AFTER 1 YEAR	GAIN
Never (fees held as cash)	$1,107.34	+$107.34
Monthly	$1,112.78	+$112.78
Weekly	$1,113.19	+$113.19
Daily	$1,113.30	+$113.30

Monthly versus daily is worth $0.52 a year. Never-compound versus monthly is worth $5.44. Compound monthly or when accrued fees exceed ~$25, whichever is later — and **never let compounding frequency justify an extra on-chain transaction.**

Note also what compounding does *not* do: it does not reduce IL. IL is a function of entry price, exit price, and band geometry, none of which compounding touches. What it does is accumulate an offsetting annuity — and over a one-year horizon the fee stream is ~11% of principal while a −25% move costs ~5% in IL. **Time is the ally: the annuity**

accrues linearly while IL only crystallizes if you close. That, and not the compounding cadence, is the argument for a long hold.

Step 7 — Instrument correctly before you deploy a dollar

Required. Refuse to trade without all five:

1. **Excess return versus 50/50 HODL**, computed from hour one. Everything else is decoration.
2. **Fee / IL coverage ratio**, gated at **> 1.20**.
3. **Realized IL per close**, target < $2 per $1,000.
4. **Two independent equity computations that must reconcile to $0.00** — an event ledger (Σ fees − IL − costs) and a bankroll that carries forward. Defects #9 and #10 both die instantly under this check.
5. **A reference `initial_principal` that never changes.** Never use a compounding principal as a P&L baseline.

Forbidden as success metrics. Each of these was near-perfect during the worst week on record:

- Gross fee APR — showed +424% during a −16.1% week; it is a **warning indicator**, because high fee APR means high volume means price is trending.
- In-range percentage — **91%** during the worst week in the corpus, which finished below buying and holding the crashing asset.
- Hold duration — improved 5× while coverage halved.
- Compound count — that is fee income re-labelled.
- Snapshot in-range rate — use full-period containment probability instead (1W: 92.5% snapshot, **0%** containment).

Step 8 — Pre-flight checklist

- ☐ Objective and unit of account stated in writing (USD preservation vs SOL accumulation vs fee carry)
- ☐ Break-even table computed from **live** pool data (§8.3)
- ☐ Expected hold ≫ break-even days, with margin
- ☐ Band width chosen by **worst-regime** performance, not best
- ☐ Band is log-symmetric and centred on the **live** price at the instant of opening
- ☐ Reopen policy is never-close or passive one-sided — **no 50/50 reconstitution swap at any boundary**
- ☐ No upside exit rule encoded (discretionary calls stay outside the machine and are owned by a human)
- ☐ Compounding is monthly or ≥$25, whichever is later
- ☐ All five required metrics wired; all five forbidden metrics removed from the headline
- ☐ Two equity computations reconcile to $0.00
- ☐ Position sized so the **−18% tail in either direction** is survivable
- ☐ Deployment verified to be running the code you think it is (defect #7 cost 17 days)

14.1 The one-paragraph version

Decide first whether you actually want to be short volatility on SOL; if you have a directional view, trade it directionally instead. If you do want the carry, open **one** log-symmetric band of **±60% to ±90%**, centred on the live price, sized so a −18% shortfall versus holding is acceptable in either direction. **Never close it at a boundary.** Collect fees on a

schedule and compound them monthly. Measure yourself against a 50/50 basket and against nothing else. Expect **7–9% a year**, not 300%. If you find yourself tuning grace windows, slippage caps, or compounding cadence, you are working on levers worth cents while the levers worth percent sit untouched.

---

15. Honest expectations

Per $1,000, one year, ±90% log-symmetric band, based on the regimes actually measured:

SCENARIO	OUTCOME	VS 50/50 HODL
Base case — SOL range-bound	~+9% fees, IL under 2%, **net +7–8%**	≈ +5%
SOL −47% (floor)	~$626 USD, but **15.396 SOL vs 6.473 at entry (+138%)**	−18%
SOL +90% (ceiling)	~$1,189, all USDC	−18%
SOL +159%	still $1,189 — position stops participating	−34%
Full 2025-style rally-and-crash cycle	mean **−5.28%** per 120-day window, worst **−20.11%**	negative

For a tight, actively-managed band with the passive-reopen fix shipped, the measured expectation is **+1.5 to 2.0% per week in chop and −1 to −2% per week in trends** — positive only if SOL spends more time ranging than trending, which nobody in this corpus can predict.

Three statements to internalize:

1. **This is not a HODL substitute.** Over a full cycle containing a 30% drawdown it will not beat holding the basket. That was tested directly and comprehensively.
2. **The floor case pays in SOL and the ceiling case pays in USDC.** Only one of those is a consolation, and only if SOL was what you wanted. Choose the band that matches your unit of account, or the tail will be denominated in the wrong thing.
3. **$100/day on $3,000 is 1,217% annualized.** It was never achievable and the pursuit of it caused the wrong strategy to be built. The realistic figure at that principal is **$3–5/day gross of price risk**, and price risk is larger than the fee stream.

---

16. Open questions and the falsifiable next experiment

16.1 The next experiment (already specified, not yet run)

Single change, measured in isolation: passive one-sided reopen. ±2.5% fixed half-width, grace unchanged at the learned p75. Rebaseline to $1,000 first (the persisted state is still $21.00 short from defect #9). Change nothing else for the duration.

Success criteria over 7 days:

- Fee/IL coverage ratio **> 1.0** (currently 0.52)
- Realized IL per close **< $2** (currently $8.87)

- Performance versus 50/50 HODL **positive** (currently −1.19%)
- ≤ 3 closes/day — requires raising `MAX_DAILY_TXNS` from 2, since passive closes are gas-only with no swap

This experiment can fail cleanly, which is its main virtue and something no prior run in this project possessed.

16.2 Blocking issues before any live capital

- **Defect #12 (live-mode double-count) must be resolved.** With `EXECUTION_ENABLED=true`, a confirmed on-chain compound calls `add_to_principal`, which grows `liquidity_l` *and* `effective_principal_usd`, while `_accrue_fees` has already counted the same dollars. Crediting the full fee ledger at close would double-count from the first compound.
- **Defect #11 (`add_to_principal` distorting IL) must be fixed** before real compounding is enabled — it grows `liquidity_l` without updating `entry_token0` / `entry_token1`, so every subsequent IL number is wrong, and over a year of monthly compounds the distortion accumulates.
- **The devnet gate** (2 clean end-to-end rebalance cycles) has still not been executed.

16.3 Genuinely open questions

QUESTION	WHY IT IS OPEN	HOW TO CLOSE IT
Does upward band skew help?	The test that appeared to refute it was mis-parameterized — the "skew up" bands also had tighter floors, and floor distance drove the results	Hold the floor fixed, vary only the ceiling
Does the ceiling-exit result survive a second asset or cycle?	Only 3 of 18 windows had the ceiling binding; the 18 windows overlap almost completely (≈2–3 independent observations)	Re-run on a second asset or a second cycle before acting
Was the manual `rewards.db` approach net profitable?	The database contains fee income only — no principal, deposits, withdrawals, range adjustments, or IL. It cannot demonstrate net profit	Reconstruct positions from on-chain history
Is a delta-hedged LP position viable?	Theoretically the correct way to isolate the fee carry, and the only untested structural answer to the tails	Separate project: perp venue, margin management, liquidation handling
Is a CLMM the right venue at all at ±90%?	A ±90% band is effectively full-range. If that is the answer, a constant-product pool may be simpler and cheaper	Compare against a Raydium/constant-product equivalent
Does the mSOL fallback for a stranded floor position work?	Staking the stranded SOL at ~7% yields +1.078 SOL/yr at the ±90% floor — **more than the entire LP fee income at that band**. But withdrawing realizes IL, and the re-entry trigger is undesigned	Backtest a re-entry rule; investigate the mSOL/SOL conversion rates in `rewards.db` (0.4997 to 3.76582 across 225 conversions — almost certainly mixed transaction types, but unverified)

---

17. Limitations of this meta-analysis

Stated plainly, because the corpus's own credibility rests on having done this consistently.

- **One asset, one venue, one fee tier.** Every result is SOL/USDC on Orca's 0.04% pool. Nothing here has been shown to generalize to other pairs, tiers, or AMMs.

- **The fee model is calibrated to the live pool on 2026-08-09** and transferred backward to 2025, when SOL traded at $141 with different volume, TVL, and liquidity distribution. **Rankings across widths are robust; absolute APRs are not.**
- **Sampling granularity varies and biases tight bands optimistically.** 15-minute sampling for the recent window (validated against 1-minute data over 19.2h: close counts differed by 0–2, final values by <0.5%); **7.1-hour** granularity for the 2025 series (808 points over 240 days), which is adequate for wide bands and endpoint-driven IL but unreliable below ±10%.
- **`range_avg_liquidity` is modelled**, as a log-interpolation between liquidity-at-tick and a single observed wide-band value. The real liquidity distribution across a ±90% band was never measured.
- **Overlapping windows inflate apparent sample size.** The ceiling study's 18 windows are 120 days long with a 7-day step over a 240-day series — an effective independent sample of roughly 2–3. Its reported means are not estimates of a population mean. The same caution applies in weaker form to S1 and S2.
- **Chop results carry in-sample risk**, because the fee model was calibrated on the chop window.
- **§5 is my own reconciliation**, derived from source code (`app/rebalance_sim.py` lines 555, 734, 1018) and the June 24 report's stated figures. It has not been verified against the raw `rebalance_sim.json` state file, and the corrected figures are reconstructions from published aggregates rather than a replay.
- **All production results are paper.** No real capital was ever at risk. On-chain slippage during thin liquidity will exceed the simulated model; the simulation is optimistic on execution.
- **Survivorship in the defect record.** Thirteen defects were found. There is no way to know how many remain, and defect #10 sat undetected for two months while being the headline number on a dashboard.

Appendix A: Full event ledger, parallel architecture

A.1 The bug-cascade window (2026-05-21 → 05-23) — Grade C, defect-contaminated

#	DATE	TYPE	RANGE	WIDTH	PRICE	HOLD	FEES	IL	SLIP	NET
1	May 21 1:53 PM	Normal	$83.86→$89.63	4.8%	$87.71	64.6h	$23.32	−$8.37	−$1.51	**+$13.22**
2	May 22 3:35 PM	Normal	$85.37→$100.03	**28.1%**	$84.49	25.7h	$7.25	−$13.27	−$1.48	**−$7.73**
3	May 22 5:36 PM	Contraction	$100.03→$87.29	4.8%	$85.14	2.0h	$0.00	−$0.11	−$0.21	−$0.36
4	May 23 4:23 AM	Normal	$87.29→$100.75	**30.4%**	$82.07	10.8h	$7.01	−$12.23	−$1.49	**−$6.94**
5	May 23 5:23 AM	Contraction	$100.75→$89.09	11.6%	$81.99	1.0h	$0.12	−$0.00	−$0.51	−$0.48
6	May 23 6:23 AM	Contraction	$89.09→$83.69	3.6%	$82.17	1.0h	$0.18	−$0.02	−$0.47	−$0.32
7	May 23 11:15 AM	Normal	$83.69→$101.18	**30.3%**	$84.35	4.9h	$1.96	−$8.84	−$1.54	**−$8.67**
8	May 23 12:15 PM	Contraction	$101.18→$88.99	10.0%	$84.02	1.0h	$0.06	−$0.03	−$0.32	−$0.35
9	May 23 1:36 PM	Contraction	$88.99→$85.67	3.6%	$84.12	1.3h	$0.00	−$0.01	−$0.10	−$0.14
10	May 23 5:11 PM	Normal	$85.67→$102.88	**31.8%**	$86.92	3.6h	$1.15	−$12.46	−$1.55	**−$13.09**
11	May 23 8:00 PM	Contraction	$102.88→$87.40	3.9%	$85.65	2.8h	$0.14	−$0.35	−$0.18	−$0.43

Events 2, 4, 7, 10 are Bug 5 manifestations (28–32% bands). Events 3, 5, 6, 8, 9, 11 are the six voluntary contractions — **zero winners**.

A.2 The clean post-bugfix window (2026-05-26 → 06-01) — the best production sample recorded

#	DATE	RANGE	WIDTH	PRICE	HOLD	FEES	IL	SLIP	NET
12	May 26 1:53 PM	$87.40→$85.54	4.6%	$83.64	65.9h	$23.22	−$7.02	−$1.52	**+$14.46**
13	May 27 11:51 PM	$85.54→$82.75	4.1%	$81.04	34.0h	$21.47	−$10.43	−$1.53	**+$9.27**
14	May 29 1:43 PM	$82.75→$84.56	3.7%	$82.95	37.9h	$10.99	−$6.80	−$1.59	**+$2.36**
15	Jun 1 3:57 AM	$84.56→$82.81	4.6%	$80.89	62.2h	$22.25	−$8.28	−$1.57	**+$12.16**

Four events, four winners, average hold **50.1 hours**, total **+$38.25**. This is what the tight-range strategy looks like when it works: long holds in a market that is not trending hard. It lasted six days.

A.3 The crash window (2026-06-01 → 06-06) — 21 events, net −$15.52

#	DATE	PRICE	RANGE	WIDTH	HOLD	FEES	IL	NET
1	Jun 1 07:57	$80.89	$79.01–$82.81	4.7%	62.2h	$22.25	−$8.28	**+$12.16**
2	Jun 2 11:51	$79.00	$77.57–$80.45	3.6%	27.9h	$12.16	−$6.39	+$3.93
3	Jun 2 15:12	$77.03	$75.57–$78.57	3.9%	3.4h	$2.68	−$8.70	−$7.89
4	Jun 2 20:46	$75.07	$73.62–$76.52	3.9%	5.6h	$5.53	−$8.67	−$5.02
5	Jun 2 23:36	$73.40	$71.74–$74.89	4.3%	2.8h	$4.65	−$7.01	−$4.26
6	Jun 3 06:36	$74.94	$73.23–$76.69	4.6%	7.0h	$5.25	−$5.59	−$2.25
7	Jun 3 16:19	$72.98	$72.13–$73.95	2.5%	9.7h	$9.96	−$8.23	−$0.17
8	Jun 3 20:04	$71.70	$70.75–$73.06	3.2%	3.8h	$7.17	−$6.18	−$0.96
9	Jun 4 01:20	$70.13	$70.25–$72.24	2.8%	5.3h	$12.62	−$7.10	+$3.55
10	Jun 4 02:20	$68.79	$67.50–$69.99	3.6%	1.0h	$0.00	$0.00	−$2.09
11	Jun 4 03:20	$70.88	$69.20–$72.69	4.9%	1.0h	$0.76	−$11.67	**−$12.92**
12	Jun 4 08:19	$68.60	$67.34–$70.14	4.1%	5.0h	$8.15	−$11.25	−$5.04
13	Jun 4 14:10	$70.31	$68.89–$71.74	4.1%	5.9h	$11.10	−$8.53	+$0.56
14	Jun 4 17:44	$68.53	$67.45–$69.80	3.4%	3.6h	$7.43	−$8.82	−$3.37
15	Jun 5 03:09	$67.09	$66.15–$68.23	3.1%	9.4h	$22.98	−$7.04	**+$13.93**
16	Jun 5 06:23	$65.00	$64.25–$66.49	3.4%	3.2h	$13.22	−$12.85	−$1.72
17	Jun 5 16:13	$63.69	$62.76–$65.15	3.8%	9.8h	**$55.81**	−$5.77	**+$47.87**
18	Jun 5 17:15	$65.57	$63.16–$65.92	4.2%	1.0h	$1.40	−$13.21	**−$13.98**
19	Jun 5 19:09	$62.19	$60.98–$64.56	5.8%	1.9h	$11.62	**−$35.37**	**−$25.96**
20	Jun 5 21:07	$64.92	$63.34–$66.57	5.0%	2.0h	$4.63	−$19.80	**−$17.35**
21	Jun 6 03:35	$62.81	$61.00–$64.67	5.8%	6.5h	$19.99	−$12.41	+$5.45
						$239.37	**−$212.86**	**−$15.52**

Note event 17 (+$47.87, a 1,205% annualized rate) immediately followed by events 18–20 (−$57.29 combined). **The single best event in the corpus was more than erased by the three trades that followed it within 28 hours.**

A.4 The state-machine cycle ledger (2026-07-26 → 08-09)

#	BAND	SRC	HOLD	IN-RNG	FEES	IL	NET	BANK (LOGGED)	BANK (CORRECTED)
1	74.29–76.16	1D	6.3 h	41.9%	$0.08	−$5.12	−$5.26	$1,002.99	$1,002.99
2	72.76–79.01	1W	66.1 h	**99.0%**	$9.05	−$14.81	−$5.97	$968.41	$977.40
3	71.45–74.04	1D	15.8 h	97.0%	$6.31	−$4.64	**+$1.46**	$972.66	$987.75
4	71.44–76.35	1W	54.7 h	**99.5%**	$6.87	−$12.36	−$5.70	$943.14	$963.86
5	69.33–73.08	1D	7.7 h	96.4%	$0.64	−$7.42	−$6.99	$950.74	$971.62

One winner in five. Cycles 2 and 4 are the important ones: 99% in-range, 55–66 hour holds, textbook execution by every metric the system was built to optimize — and both lost ~$6.

Appendix B: Data provenance

ARTIFACT	CONTENTS	USED FOR
`app/data/rebalance_sim.json`	Parallel-sim state, 72 events, 661 compounds	E1–E5
`app/data/strategy_sim.json`	State-machine state, 5 closed cycles + 1 open	E6, E7
`app/data/history.json`	1-minute SOL price series (1,000 points)	Sampling validation
LiveCoinWatch	15-min series (1,363 pts, recent); 7.1-hour series (808 pts, 2025 cycle)	S1–S4
`/api/pool/stats` (live)	Pool 24h fees $14,972.05, TVL $26.36M, L 1.265e15	Fee calibration, break-even table
`rewards.db`	565 on-chain fee collections, 2025-06-27 → 2026-01-19, one wallet	S5 (grade D)
`analysis/2026-08-09/backtest.py`	Counterfactual engine — validated to the penny vs the production IL ledger	S1–S3
`analysis/2026-08-09/ceiling.py`	Upside-exit rule sweep	S4
`analysis/2026-08-09/regime.py`	ER-gated variants (result: withdrawn)	S2, S4
`app/tests/test_fee_ledger.py`	5 regression tests for defect #9 (3 fail pre-fix, all pass post-fix)	Defect verification

---

Closing statement

Ninety-three days, seven strategies, two architectures, thirteen defects, five counterfactual studies, and zero dollars lost. The strategy does not work in the form it was built, and the corpus now explains why with arithmetic rather than with statistics: **the rule closed positions at local extremes and swapped there, which cost 97% IL and 3% fees, and the project spent its effort on the 3%.**

The constructive finding is unglamorous and durable. **Open wide. Centre on the live price. Do not close at a boundary. Compound monthly. Measure against holding the basket. Expect a single-digit annual carry and a −18% tail in either direction.** That is a real strategy with a real edge in the right regime, and it is roughly the opposite of the high-frequency, tightly-concentrated, aggressively-rebalanced system that three months of engineering produced.

The most valuable output of this project is not a configuration. It is a validated counterfactual engine, a set of correctly-specified metrics, and a documented record of a search that eliminated most of its own hypothesis space. That is what the next attempt should start from.

Meta-analysis compiled 2026-08-09 from the full document corpus, source code in `app/`, and the counterfactual engines in `analysis/2026-08-09/`. Every figure is traceable to a cited source; derived figures show their arithmetic. §5 is a new reconciliation produced for this document and is labelled accordingly.