

Monthly Config Watch #1

September 4, 2026 · MarketMania Research · Monthly series

MONTHLY · CONFIG WATCH · Issue #1 · one calendar month, 30 replays, one survival curve · Issue date September 4, 2026 · Language English · Methodology v1.1 + Config Watch amendments (approved 12.08, hash e66c7e8c864a2233)
Search window: **MONTH** Aug 1-31, 2026 -- the frozen calendar month (Sat Aug 1 00:00 -> Mon Aug 31 23:59 UTC, 31 days), cut-off edge Mon Sep 1 00:00 UTC, so slots dated Sep 1 fall outside it; report cut-off Mon Sep 1, 2026, 16:00 UTC. Slots fall inside the window; trades settle past its right edge at their horizon -- the same 2-3 day luft the engine has used since Config Watch #1 and that the live presets run on.

The series starts here: this is issue #1, so there is no prior monthly window and no month-over-month number in this report. The weekly Config Watch #4, published from the same night collect, carries the calendar week Aug 24-31, 2026 and the four-week window Aug 3-31, 2026; both lie inside this month and neither board is reprinted here.

RESEARCH SNAPSHOT

10,889 simulations, 0 errors, in one frozen collect	MONTH Aug 1-31, 2026	cut-off edge Mon Sep 1 00:00 UTC ; report cut-off Mon Sep 1, 2026, 16:00 UTC	7-stage search: scan -> grid -> refine -> stake sweep -> per-ticker -> OOS replays
engine 1.1; \$1,000 deposit, 10x leverage	stake is universe-aware from sim #1 (table below)	maturity gate: >=20 closed (MONTH)	criteria frozen in writing before launch (w200 plan)
walk-forward: 26 past cards + 4 live presets replayed verbatim on the month	20 of 26 past cards net-positive on the month	stage-2 grid: 4,112 cells on this window	models: 7 (grok axis = grok-4.6 (incl. 4.5-era))
best MONTH council +\$1,618.4 (+161.8%)	best MONTH solo +\$1,721.9 (+172.2%)	zero-trade cells: 1,188 (whole collect)	margin-skip alerts on leaders: 10 rows (MONTH)

KEY FINDING · OBSERVATION (one calendar month)

The best mature council on calendar August makes +\$1,618.4 (+161.8% of a \$1,000 deposit, 77.4% win rate, 40.5% maxDD, 31 closed), 20 of the 26 cards this series has published are net-positive on the same month, and all 7 cards that survived both fresh windows of Config Watch #4 are among them.

WEEKLY EVIDENCE -> MONTHLY VERDICT (rule below)


Metric	CW#1 wk	CW#2 wk	CW#3 wk	CW#4 wk	Month	Verdict
Published cards net-positive (majority)	n/a	6/12	19/20	7/26	20/26	Rejected (1/3)
Best council universe	SOL+XRP	SOL+XRP	SOL+XRP	BTC+ETH+ SOL+BNB+ XRP	SOL+XRP	Confirmed (3/4)
Calibration helps (matched pairs)	n/a	75.0%	0.0%	61.0%	37.9%	Rejected (1/3)

Verdict rule: each row states one reading of the month; an issue agrees when its own fresh week shows the same reading (same side of the line, or the same name). Confirmed = 3 or 4 of the 4 issues agree with the month; Inconclusive = 2; Rejected = 0 or 1. A cell that reads n/a has no comparable number in that issue: it never agrees and leaves the count, so a row with one n/a is scored out of three and Confirmed there reads 3 of 3. Issue cells come from the four published issues' own open data, each on the fresh week it published; the month column is this collect's calendar-August pull.

Sources, cell by cell. Row 1 counts the replays each issue ran on its own fresh week -- Config Watch #2 6 of 12, #3 19 of 20, #4 7 of 26 on the calendar week Aug 24-31, 2026 -- against 20 of 26 on the month in the walk-forward table below; Config Watch #1 was the in-sample first issue and ran no replays, so its cell is n/a. Row 2 reads the coin universe of the top council configuration on each issue's fresh-week board, and here the top council card of the month. Row 3 is the share of matched calibrate-ON / calibrate-OFF pairs the axis improved on that fresh week: 4 pairs in issue #2, 279 exported top-of-board pairs in issue #3 (99.2% of their calibrate-OFF legs are net-positive -- a tail sample, not the middle of the grid), 1,468 whole-grid pairs in issue #4 and 1,419 whole-grid pairs on the month; issue #1 carries no calibration axis, so that cell is n/a too.

TL;DR

- OBSERVATION -- 20 of the 26 published cards are net-positive on the calendar month, 7 of 26 were on the fresh week inside it (Aug 24-31, 2026).** On the four-week window of Config Watch #4 (Aug 3-31, 2026) the same 26 cards read 19 of 26. By generation: 4 of 6 issue-3 cards, 8 of 8 issue-2 cards, 8 of 12 issue-1 cards, and 4 of 4 live presets. This month contains both of those windows, so the three readings are not three independent tests.
- The month's ceiling is +\$1,618.4 on the council board and +\$1,721.9 on the solo board.** The council card runs at +161.8% of a \$1,000 deposit, 77.4% win rate and 40.5% maxDD on 31 closed trades; the solo card at +172.2%, 73.8% win rate and 54.9% maxDD on 42 closed with 3 unfunded entries. The grid behind them is flat at the centre: median \$0.0 and 41.8% positive across 4,112 stage-2 cells, 427 of them above +\$500, range -\$805.1 to +\$1,721.9. On the four-week window of the same collect that grid ran 47.7% positive, and the two windows price the same cards differently: the month leader gives back \$36.0 on the four weeks and Config Watch #4's four-week leader gives back \$243.1 on the month (section below).
- The survival curve, first published here, reaches +3 fresh weeks for the issue-1 generation.** Share of a generation's cards net-positive on the k-th fresh week after its own window edge: issue-1 cards 50.0%, then 91.7%, then 8.3%; issue-2 cards 100.0%, then 50.0%; issue-3 cards 33.3%. Every cell is read out of the issue that ran that replay, each is a different calendar week of August, and the same cards read 8 of 12 (66.7%), 8 of 8 (100.0%) and 4 of 6 (66.7%) on this calendar month.
- The 7 survivors and the 19 faded cards differ on the recorded traits.** The cards that cleared both of Config Watch #4's fresh windows against those that did not: median fresh-week maxDD 13.5% vs 45.8%, median win rate 82.4% vs 47.1%, 6 of 7 against 3 of 19 on the 1d horizon, and 1 of 7 (14.3%) vs 0 of 19 (0.0%) calibrated. 7 rows against 19 rows on one week is a description of two groups, not a test, and none of the traits is claimed to cause survival.



Share of each generation's cards net-positive on the k-th fresh week after its own window edge (Config Watch #1 12 cards, Config Watch #2 8 cards, Config Watch #3 6 cards). Each point is one published walk-forward table: for issue-1 the three points are Config Watch #2, #3 and #4; for issue-2 they are #3 and #4; for issue-3 it is #4. The x axis counts fresh weeks, not calendar weeks -- every point sits on a different week of August -- and neither the live presets (no publication date) nor the calendar-month column (a 31-day window) is plotted.

Why it matters

MarketMania publishes default LLM-council trading configs. The weekly Config Watch asks what one week of large-scale in-sample search extracts and how much of it survives the next week. This series asks the same question with the calendar month as the unit: one frozen 31-day window, the same engine, the same criteria and the same frozen search plan, and the same walk-forward table read on that window. Two measurements only a monthly cadence can carry are published here for the first time: the survival curve by generation, which places each published card's replay on the k-th fresh week after its own issue, and the structural comparison of the 7 cards that cleared both fresh windows of Config Watch #4 against the 19 that did not. The month is not an independent test of the weeks: it contains all four published Config Watch weeks and the whole 28-day four-week window, and every section that reads the two together says so.

How we searched

One frozen collect, run to the plan with the criteria frozen in writing *before* launch (w200 plan) and printed into the meta of the JSON deliverable: the night run, which sweeps the four-week window and the calendar month in one pass. Every stage count below is that run whole. Its four-week boards ship with Config Watch #4, so no count is split between the two issues; the stage-2 grid is 4,112 cells on each of the run's two windows.

- **1. Scan (s1)** -- 952 sims across council compositions and coarse settings.
- **2. Systematic grid (s2)** -- 8,224 sims sweeping the declared axes (archetype, FH/TF, entry filters, universe, council size, membership, TP/SL source, ladder, break-even, TP/SL shift, calibrate and stake) -- the denominator for every distribution claim below.
- **3. Refinement: ladder + stake sweep (s2b+s3b)** -- 243 sims around the grid winners.
- **4. Preset neighbourhood grid (s2p)** -- 106 sims: one-knob neighbours of the live presets that have a local grid this issue.
- **5. Per-ticker probes (s5t)** -- 130 sims: best finalists split onto single coins.
- **6. Dedicated per-ticker (s6t+s6tc)** -- 1,160 sims: a compacted independent per-coin grid plus calibrate twins (the standing branch from Config Watch #2).
- **7. OOS replays (s4cw3+s4cw3o+s4cw2+s4cw2o+s4cw1+s4pre)** -- 74 sims: every published card and all four live presets, verbatim, on both windows of this collect, plus origin re-runs of the issue-2 and issue-3 cards for the drift baseline.

Total 10,889 simulations, 0 errors, 1,188 zero-trade cells; stage sums reconcile exactly inside the collect. The axis grid was NOT widened relative to Config Watch #4 (anti-overfit rule): this issue changes the window, not the search space. Equity series, the grid histogram, the baselines and the calibrate pack are produced inside the same collect -- there is no separate extras pass.

Stake: universe-aware from simulation #1 (\$1,000 deposit)

Instruments in universe	Stake cap	Sweep values also simulated
1 coin	\$450	\$300 / \$375
2 coins	\$450	\$300 / \$375
3 coins	\$300	\$225 / \$375
4 coins	\$250	\$200 / \$300
5 coins	\$200	\$150 / \$250

A \$1,000 deposit cannot fund five concurrent \$450 positions. Stake is a function of how many instruments the config trades, recomputed on every universe change and swept as its own axis; any simulated stake above the cap for its universe size is marked **concentrated** and excluded from the headline boards (it stays in the open data). Every card also prints **skipped_nofunds** -- entries the engine could not fund. Margin-skip alert rows on this window's leaders: 10, of the 18 the collect raised over its two windows; the other 8 are on the four-week board and belong to Config Watch #4.

Models under test (exact engine IDs): claude-fable-5, claude-opus-5, gpt-5.6-sol, deepseek-v4-pro, gemini-3.1-pro, qwen-3.8-max, grok-4.6. The grok axis runs as **grok-4.6 (incl. 4.5-era)** -- one lineage; published grok-4.5 configs replayed out-of-sample get the same 4.5 -> 4.6 lineage remap ([remap] rows).

TOP councils -- calendar month (Aug 1-31, 2026)

#	Council	Coins	FH/TF	Net	Net %	maxDD	Trades	WR
#1	fable, gemini	SOL+XRP	1d/1d	+\$1,618.4	+161.8%	40.5%	31	77.4%
#2	fable, gpt, deepseek, gemini, qwen	ETH+SOL+BNB+XRP	1d/1d	+\$756.6	+75.7%	6.4%	41	92.7%
#	Ladder	BE	Shifts	Conf filter	Source	Entry	Positional	Stake
#1	50:40,100:60	on	TP -10 / SL +75	>=60 (each)	farthest	1/0	reenter	\$450
#2	50:40,100:60	on	TP -10 / SL +75	>=60 (each)	farthest	3/1	reenter	\$250

Against the 4,112-cell stage-2 grid on this window: card #1 is the same simulation as the raw net-board leader (i=6910), so its exported percentile is the card's -- above **99.9%** of the grid; card #2 lands in the +\$750 to +\$1,000 bucket (106 cells) with 95.2% of the grid finished below it. Grid median \$0.0, 427 of the 4,112 cells above +\$500. Net % = net PnL as a percentage of the fixed \$1,000 starting deposit; every simulation runs at **10x leverage** (liquidation modeled by the engine). Entry = max_sideways / max_diff_side.


Two councils are published on this window. All ten slots of the mature net board are the same 2-coin universe (SOL+XRP), so the pairwise-different-universes rule admits exactly one; #2 is the best mature multi-coin council outside that universe on the quality boards -- it comes off the **Calmar** board (+\$756.6 on ETH+SOL+BNB+XRP, 6.4% maxDD, R2 0.894, 0 unfunded entries). Read the cards as two universes deep, not three.

The two cards answer the calibrate axis differently: card #1 is an uncalibrated cell and card #2 a **calibrate=ON** one (Calibrate axis below). Card #1 runs the \$450 two-coin table stake and card #2 the \$250 four-coin stake, so its +\$756.6 is earned on the smaller of the two; neither card carries an unfunded entry, and card #1 pays for its +\$1,618.4 with a 40.5% drawdown on 31 closed trades against card #2's 6.4% on 41. The mature pool on this window is 3,265 configs against 3,184 on the four-week window of the same collect: three more days of trading put 81 more configs over the ≥ 20 -trade gate.

OPEN IN SANDBOX marketmania.ai/s/cwm1-m1 marketmania.ai/s/cwm1-m2


Each link opens this exact frozen window; results visible without sign-in (embedded share signature). The codes are minted at publish time and are NOT live in this draft; every card's full frozen query signature is in the open-data JSON next to its simulation index.

Published card equity -- MONTH (Aug 1-31, 2026) -- \$1,000 at 10x -- 31 settled trades



Daily settled equity for card #1 above (\$1,000 start, 10x leverage, frozen window; #1 red -- fixed rank colors across all issues). A curve running past the window end is an open-at-cutoff trade settling at its horizon (the 2-3 day left); the last point reconciles to the card's net to the cent. There is no muted contrast line on this window: the mature leader is also the net-board maximum (the same simulation, i=6910), so a second curve would be the same line. The line peaks at +\$2,023.1 on 2026-08-27 and ends the window at +\$1,618.4. Card #2, the solo card and the board leaders keep their daily series in the open data -- the PDF carries one equity chart.

Stage-2 systematic grid -- net PnL distribution -- MONTH window (Aug 1-31, 2026) -- n=4,112



Net PnL across all 4,112 stage-2 grid sims, MONTH window (Aug 1-31, 2026), universe-aware stake throughout. Grid median **\$0.0**, 41.8% positive; the same grid on the four-week window of this collect ran 47.7% positive and on Config Watch #4's fresh week 31.0%. The spike at \$0 is mostly cells whose entry filters produced no trades (562 zero-trade sims on this window). Range -\$805.1 to +\$1,721.9; 427 cells finished above +\$500 and the maximum of the grid is the mature solo board's leader. Red marker = best solo default at table stake; green = the published cards.

The 30-day cards against the calendar month

Config Watch #4 published two councils and one solo on the four-week M30 window (Aug 3-31, 2026, four full weeks, edge Mon Aug 31 00:00 UTC); this issue publishes two councils and one solo on the calendar month (Aug 1-31, 2026, edge Mon Sep 1 00:00 UTC). Both sets come out of the same night collect, and the month contains the four weeks plus 3 days -- Aug 1, Aug 2 and Aug 31. The table measures all six cards on both windows: M30 is Config Watch #4's window and Month is this issue's.

Card	Council / solo	Coins	FH/TF	Net M30	Net month	WR M30/Mo	maxDD M30/Mo	ClS
cw4-m301	gpt, deepseek, gemini	SOL+XRP	1d/1d	+\$1,676.0	+\$1,433.0	80.6% / 74.3%	53.6% / 69.5%	31 / 35
cw4-m302	fable, gpt, deepseek, gemini	BNB+XRP	1d/1d	+\$1,155.2	+\$912.5	75.0% / 68.8%	11.3% / 36.1%	28 / n/a
CW#4 solo	gemini	SOL+XRP	1d/1d,4h	+\$1,741.7	+\$1,721.9	75.0% / 73.8%	54.9% / 54.9%	40 / 42
cwm1-m1	fable, gemini	SOL+XRP	1d/1d	+\$1,582.4	+\$1,618.4	76.7% / 77.4%	40.5% / 40.5%	30 / 31
cwm1-m2	fable, gpt, deepseek, gemini, qwen	ETH+ SOL+ BNB+XRP	1d/1d	n/a	+\$756.6	n/a / 92.7%	n/a / 6.4%	n/a / 41
cwm1-s1	gemini	SOL+XRP	1d/1d,4h	+\$1,741.7	+\$1,721.9	75.0% / 73.8%	54.9% / 54.9%	40 / 42

Every paired column reads **M30 first, calendar month second**; Cls = closed trades. **CW#4 solo and cwm1-s1 are one configuration** -- gemini on SOL+XRP at 1d/1d,4h and the \$450 two-coin stake -- published by Config Watch #4 on its four-week window and by this issue on the calendar month, so those two rows carry the same pair of numbers; the other four rows are four distinct configurations. Place on each window's mature net board (councils and solos in one list, as the collect ranked them): cw4-m301 3 on M30 and 14 on the month; CW#4 solo 1 on M30 and 1 on the month; cwm1-m1 8 on M30 and 3 on the month. Two cells are not in this collect: cw4-m302 reaches no calendar-month board, so its month figures come from the matched calibration pair the collect exported for that window (which carries opened trades, 32 here, and not the closed count), and cwm1-m2 -- the \$250 four-coin calibrated cell from the stake sweep -- reaches no four-week board and has no four-week pair in the export, so its M30 column reads n/a. Boards and a pair sample are what the collect exports, not every cell on both windows.


Every configuration both windows can measure is worth between \$19.8 and \$243.1 more on one window than on the other. The three cards Config Watch #4 published are all larger on its four-week window: cw4-m301 makes +\$1,676.0 there and +\$1,433.0 on the month (-\$243.1, on 35 closed trades against 31 and a 69.5% drawdown against 53.6%), cw4-m302 +\$1,155.2 against +\$912.5, and the solo +\$1,741.7 against +\$1,721.9. This issue's council card runs the other way: +\$1,618.4 on the month against +\$1,582.4 on the four weeks, a \$36.0 difference on one closed trade. The solo is one configuration in both issues, so its larger number sits on Config Watch #4's window and its smaller on this one. The two windows do not keep the order of the same rows either: on M30 the mature net board puts the solo 1st, cw4-m301 3rd and cwm1-m1 8th; on the calendar month it puts the solo 1st, cwm1-m1 3rd and cw4-m301 14th.

Walk-forward on the calendar month

Every config this series has ever published -- all 26 cards from Config Watch #1, #2 and #3 -- replayed VERBATIM (same knobs, same stake as printed, no re-tuning) on this issue's calendar month, plus the four hand-tuned presets running live on the platform, listed first. The replays and their origin re-runs cost 74 simulations inside this collect, which replays every row on both of its windows in one pass. Config Watch #4's own five cards are not in the table: the collect was frozen on 2026-09-03T12:07:22Z, before that issue published, so their first out-of-sample read belongs to the next issues.

Verdict rule, on the one window this issue has. Survived = net-positive on the calendar month; **Faded** = net-negative. It is the standing rule (net-positive on every fresh window) with one window in it, and it is therefore easier to pass than the two-window version Config Watch #2, #3 and #4 applied. The column is not pure out-of-sample either: this month contains all four published Config Watch weeks (Config Watch #1 Aug 5-12, 2026; Config Watch #2 Aug 10-17, 2026; Config Watch #3 Aug 17-24, 2026; Config Watch #4 Aug 24-31, 2026) and the whole 28-day four-week window, and the issue-1, issue-2 and issue-3 search windows all end inside it -- so the number a card shows here is partly scored on the days it was selected on. Read the column as decay context over a 31-day window.

Config Watch #3's six published cards: home number vs calendar-month replay



Config Watch #3's six published cards: the number its own issue printed (grey, in-sample on its own window) against the same configuration replayed verbatim on this calendar month (green). 4 of the six are net-positive on the month; wk1 and m301 are the same configuration published on both of issue #3's windows, so their replays are identical by construction.

Config	Family	Coins	Home	New MONTH (n)	Verdict
preset 4h4	fable, qwen, deepseek, opus	ETH+SOL	-- live	+\$371.6 (47)	Survived
preset bc2	gemini, gpt, deepseek	SOL+XRP	-- live	+\$1,200.9 (22)	Survived
preset sol	deepseek, fable, opus	SOL [conc]	-- live	+\$2,615.8 (40)	Survived
preset xrp	gemini, gpt, deepseek	XRP [conc]	-- live	+\$1,333.4 (11)	Survived
cw3-m301	opus, deepseek, gemini	SOL+XRP	+\$2,137.6	+\$1,290.8 (86) +	Survived
cw3-m302	fable, opus, deepseek, gemini, qwen	BTC+ETH	+\$412.3	+\$288.7 (25)	Survived
cw3-s1	gemini	SOL+XRP	+\$1,604.3	-\$349.4 (130) +	Faded
cw3-s2	gemini	SOL+XRP	+\$1,449.0	+\$1,721.9 (42) +	Survived
cw3-wk1	opus, deepseek, gemini	SOL+XRP	+\$2,175.1	+\$1,290.8 (86) +	Survived
cw3-wk2	opus, deepseek, gemini, grok	BTC+ETH	+\$837.6	-\$562.5 (27) +	Faded
cw2-30d1	opus, gpt, deepseek, gemini, qwen	BTC+ETH+ SOL+BNB+XRP [conc]	+\$744.7	+\$1,594.5 (39) +	Survived
cw2-30d2	opus, deepseek, gemini	ETH+SOL+ BNB+XRP [conc]	+\$655.1	+\$1,269.0 (51) +	Survived
cw2-30d3	opus, deepseek, gemini	BNB+XRP	+\$489.1	+\$927.8 (33)	Survived
cw2-30ds1	gemini	BNB+XRP	+\$383.8	+\$793.3 (39)	Survived
cw2-7d1	fable, opus, deepseek, gemini	SOL+XRP	+\$299.1	+\$881.8 (84) +	Survived
cw2-7d2	gpt, deepseek, gemini, qwen	BNB+XRP	+\$223.7	+\$813.3 (29)	Survived
cw2-pt1	fable, deepseek, grok [remap]	SOL	+\$698.3	+\$1,285.9 (56)	Survived
cw2-pt2	opus	XRP	+\$543.7	+\$179.4 (14)	Survived
cw1-30d1	fable, deepseek, opus, qwen	ETH+SOL	+\$868.4	+\$14.9 (49) +	Survived
cw1-30d2	gemini, gpt, deepseek	SOL+XRP+ETH	+\$590.6	+\$949.5 (33)	Survived
cw1-30d3	fable, deepseek, opus, qwen	BTC+ETH+SOL	+\$580.5	-\$185.9 (70) +	Faded
cw1-30ds1	opus	SOL+XRP	+\$442.8	+\$687.5 (90) +	Survived
cw1-30ds2	grok [remap]	SOL+XRP	+\$262.9	+\$237.5 (72) +	Survived
cw1-30ds3	deepseek	XRP+BNB	+\$119.0	+\$86.7 (90) +	Survived
cw1-7d1	fable, deepseek, opus, gpt	SOL+XRP	+\$439.2	+\$1,145.1 (118) +	Survived

Config	Family	Coins	Home	New MONTH (n)	Verdict
cw1-7d2	fable, qwen, gemini, opus	XRP+BNB	+\$429.5	-\$130.5 (93) +	Faded
cw1-7d3	gpt, opus	BNB+XRP+SOL	+\$398.6	+\$271.4 (139) +	Survived
cw1-7ds1	gemini	XRP+BNB	+\$417.8	-\$294.5 (190) +	Faded
cw1-7ds2	gpt	XRP+BNB	+\$366.3	+\$292.6 (123) +	Survived
cw1-7ds3	qwen	XRP+BNB	+\$267.0	-\$301.5 (103) +	Faded

n in parentheses; * = fewer than 10 closed trades (insufficient cell); + = non-zero skipped_nofunds (entries the engine could not fund -- the printed economics are then not the economics that ran). Home = the number the card's origin issue printed; the live presets have no home issue. [conc] = stake above this issue's cap for its universe size; [remap] = grok-4.5 replayed on the grok-4.6 lineage. Verdict rule: Survived = net-positive on the calendar month. Survivors are tinted green.

Counted: 20 of 26 past cards are net-positive on the calendar month -- by cohort 4 of 6 issue-3 cards, 8 of 8 issue-2 cards, 8 of 12 issue-1 cards. Config Watch #4 read the same 26 cards on two shorter windows inside this one: 7 of 26 on its fresh week, 19 of 26 on its four weeks, 7 surviving both. All 7 of those survivors are net-positive here, and so are 13 of its 19 faded cards. The best replay is cw3-s2 at +\$1,721.9 on 42 closed trades and the worst cw3-wk2 at -\$562.5 on 27. Live presets: 4 of 4 net-positive on the month, 2 of 4 survived both of Config Watch #4's windows.

The four live presets lead the table by standing rule. All four are net-positive on the month (sol +\$2,615.8 on 40 closed, xrp +\$1,333.4 on 11, bc2 +\$1,200.9 on 22, 4h4 +\$371.6 on 47), where Config Watch #4 recorded 2 of 4 on its fresh week. sol and xrp are marked [conc]: both run a \$900 stake on a single coin, twice the \$450 cap the stake table sets for a 1-coin universe, so their printed economics assume funding this issue's own rule would not grant. The bc2 row is also this issue's hand-tuned anchor (section below) and its +\$1,200.9 is the same simulation on both sides of the report.

Origin re-runs (drift baseline). Re-simulating the issue-3 cards on their OWN origin windows today lands within -\$251.1..\$0.0 of the printed issue-3 numbers: 5 of the six reproduce to the cent, and the exception is cw3-m302 (+\$412.3 printed, +\$161.2 today), the calibrated month card -- calibration is the one axis whose multipliers are learned from a forecast history that keeps growing. The issue-2 cards drift -\$177.7..+\$88.6. Snapshot principle: no past issue is restated; the full drift table is in the open data.

What this column is. 20 of 26 past cards are net-positive over 31 days that contain every week this series has searched. Three measured facts bound the reading. The window overlaps the cards' own: issue-3's search week (Aug 17-24, 2026) and issue-2's (Aug 10-17, 2026) lie inside it, and so does the whole 28-day four-week window of Config Watch #4. The rows that lost the most on the fresh week did not lose here: the 19 cards Config Watch #4 dropped on its two-window rule are 13 net-positive on the month. And the fresh search moved with them: the month ceiling is +\$1,618.4 against +\$1,676.0 on the four-week window and +\$399.4 on the fresh week, all three from the same engine and the same criteria. The month is a longer window, not a cleaner one.

Survival curve by generation

Each published generation of cards, read on the fresh weeks that followed it. A cell is the share (and count) of that generation's cards that finished net-positive on the k-th fresh week after the issue's own window edge, taken from the issue that ran that replay; the last column is this issue's calendar-month replay of the same cards. Nothing is re-simulated for this table: every number in it is already in a published open-data file, and this issue only lines them up. The chart under TL;DR draws the three card rows of this table; it is printed once, there.

Generation	Cards	+1 week	+2 weeks	+3 weeks	Calendar Aug
Config Watch #1	12	6 of 12 (50.0%)	11 of 12 (91.7%)	1 of 12 (8.3%)	8 of 12 (66.7%)
Config Watch #2	8	8 of 8 (100.0%)	4 of 8 (50.0%)	--	8 of 8 (100.0%)
Config Watch #3	6	2 of 6 (33.3%)	--	--	4 of 6 (66.7%)
Live presets	4	2 of 2 (100.0%)	2 of 4 (50.0%)	--	4 of 4 (100.0%)

Cell by cell: issue-1 cards +1 week = Config Watch #2's walk-forward (Aug 10-17, 2026), +2 weeks = Config Watch #3's (Aug 17-24, 2026), +3 weeks = Config Watch #4's (Aug 24-31, 2026); issue-2 cards +1 week = Config Watch #3's, +2 weeks = Config Watch #4's; issue-3 cards +1 week = Config Watch #4's. The live presets are not a generation and have no publication date: their two cells are the two issues that replayed them -- Config Watch #3 replayed 2 of the 4 and Config Watch #4 all 4 -- and that row is not drawn on the chart. Calendar Aug is this issue's window (Aug 1-31, 2026) for every row. Two properties of the table to read with it: a column is not one calendar week across rows (issue-1's +1 week is Aug 10-17, 2026, issue-3's is Aug 24-31, 2026), and issue-1's +1 week overlaps its own search window (Aug 5-12, 2026) by 2 days.

Survivors vs faded: structure

Config Watch #4 replayed the same 26 cards on two fresh windows and its standing rule kept 7 of them (net-positive on both) and dropped 19. This table puts the two groups side by side on the traits recorded for every card: the frozen query (horizon/timeframe, council size, universe size, calibration, minimum confidence, table stake, membership) and the fresh-week replay statistics that decided the verdict. 7 rows against 19 rows on one week: this is a description of two groups, not a test, and no trait here is claimed to cause survival.

Trait	Survivors (7)	Faded (19)
FH/TF	1d/1d (4), 1d/1d,4h (1), 1d/4h (1), 4h/1h (1)	4h/1h (12), 1d/1d (3), 4h/4h (3), 1h/1h (1)
Council size	3 models (3), 5 models (2), 4 models (1), solo (1)	solo (9), 4 models (6), 3 models (3), 2 models (1)
Universe size	2 coins (3), 1 coin (1), 3 coins (1), 4 coins (1), 5 coins (1)	2 coins (16), 3 coins (2), 1 coin (1)
Calibrate ON	1 of 7 (14.3%)	0 of 19 (0.0%)
Min confidence	60 (7)	60 (12), 0 (3), 50 (2), 55 (2)
Table stake	\$450 (6), \$300 (1)	\$450 (17), \$300 (2)
Median maxDD	13.5%	45.8%
Median win rate	82.4%	47.1%
Median smoothness R2	0.431	0.589
Model membership	deepseek (6), gemini (6), gpt (3), opus (3), qwen (3), fable (2), grok (1)	opus (12), deepseek (9), gemini (9), fable (5), qwen (4), gpt (3), grok (2)

Counts are cards, not simulations, and each card contributes once per row; the membership row counts a model once for every card it appears in, so it sums above the group size. Medians are over each group's Config Watch #4 fresh-week (Aug 24-31, 2026) replay -- the window that decided every verdict in that issue -- and the R2 median is taken over the rows that have one (7 of 7 survivors, 18 of 19 faded: a replay with too few days carries no fit). On this issue's calendar month the same two groups read 7 of 7 and 13 of 19 net-positive.

Calibrate axis: does TP/SL calibration help a config?

Matched pairs -- identical knobs, calibration OFF vs ON -- measured on net PnL. This issue reports **two different measurements** and does not pool them: the whole stage-2 grid of the calendar month (every matched pair the search produced) and the board-leader slice (pairs whose OFF or ON leg reached the top of a board). They answer different questions.

Measurement	Pairs	Calibrate wins	Mean delta	Best delta	Worst delta
Whole stage-2 grid (this window)	1,419	37.9%	-\$46.7	+\$1,108.1	-\$1,225.9
-- mature rows only	858	42.1%	-\$51.6	+\$895.6	-\$995.2
Board-leader slice (top of board)	14	0.0%	-\$574.7	-\$25.0	-\$1,225.9

On this window the axis is negative in the population and negative at the top. Calibration improves 37.9% of the 1,419 matched pairs on the whole grid (mean -\$46.7, median -\$14.2) and 42.1% of the 858 mature pairs (mean -\$51.6). On the board-leader slice it improves 0.0% of 14 pairs (mean -\$574.7, median -\$584.3): the OFF legs there are at a median of +\$1,538.9 and the ON legs at +\$996.8, and all 14 ON legs are still net-positive, so the slice measures size given up, not losses created. The same axis on the other window of this collect improved 29.1% of its pairs, and on Config Watch #4's fresh week 61.0% -- three windows of one engine, three answers. Where the axis reached the cards is visible above: card #2 is a calibrate=ON cell, card #1 is not, and 37 of the 62 finalists run it.

The open-data pack also carries a 196-pair export for this window. Its delta distribution is not the population's (median -\$472.1 against the grid's -\$14.2), so no claim in this section is computed from it; it is published for inspection only.

Secondary analysis -- solo-model top

Sidebar to the council narrative. Solo configs run inside the same pipeline and face the same gates: ≥ 2 coins, the maturity gate for the window (≥ 20 closed), one config per model.

Solo top -- MONTH (Aug 1-31, 2026)

#	Model	Coins	FH/TF	Net	Net %	maxDD	Trades	WR
#1	gemini	SOL+XRP	1d/1d,4h	+\$1,721.9	+172.2%	54.9%	42 *	73.8%
#	Ladder	BE	Shifts	Conf filter	Source	Entry	Positional	Stake
#1	50:40,100:60	on	TP -10 / SL +75	≥ 60 (each)	farthest	1/0	reenter	\$450

One solo card, not two. Every one of the ten rows on the mature solo board is the same model (gemini) on the same universe (SOL+XRP), so the one-config-per-model rule admits exactly one; no single-coin row outranks it on that board this window. The card carries 3 unfunded entries (flagged * on Trades) and its clean twin, 0 skips, makes +\$1,415.5 on 45 closed. The configuration is also **cw3-s2**, a card Config Watch #3 published: query for query, window aside, it is the same configuration, so the +\$1,721.9 in this table and the +\$1,721.9 in the walk-forward row above are one simulation -- in-sample as this month's search maximum, out-of-sample as issue #3's replay.

Council vs solo. On this window the best solo sits **above** the best council (+\$1,721.9 against +\$1,618.4 on the mature boards), on 42 closed trades against 31, with a 54.9% drawdown against 40.5% and 3 unfunded entries against 0. Config Watch #4 measured the same comparison on its two windows: the council above the solo on the week, the solo above the council on the four weeks. Same in-sample caveats apply to both sides.

OPEN IN SANDBOX marketmania.ai/s/cwm1-s1

Each link opens this exact frozen window; results visible without sign-in (embedded share signature). The codes are minted at publish time and are NOT live in this draft; every card's full frozen query signature is in the open-data JSON next to its simulation index.

Quality boards

Net PnL stays the primary board (continuity with the weekly series); four selection views sit on top of it at zero additional simulations. **Calmar-like** = net / max(maxDD, 1.0), positive net only; **WR board** = win rate among mature configs (>=20 closed); **Smoothness** = R2 of a linear fit through the daily closed-equity series, mature rows with net > 0 and at least 5 days, ties broken by ulcer index; **Quality composite** = mean rank over net, Calmar-like, win rate and R2, lower is better. Top-3 of each board below; all 15 rows of all four boards are in the open-data JSON.

Board	Config	Coins	FH/TF	Net	WR	maxDD	CIs	R2
MONTH Calmar	fable, opus, deepseek [cal]	SOL	4h/1h	+\$892.4	94.3%	0.0%	53	0.976
MONTH Calmar	fable, opus, deepseek [cal]	SOL	4h/1h	+\$905.5	93.1%	1.3%	58	0.970
MONTH Calmar	fable, opus, deepseek	SOL	4h/1h	+\$1,344.2	81.0%	9.1%	58	0.946
MONTH WR	gemini [cal]	SOL	1d/1d,4h	+\$820.7	95.5%	13.5%	22	0.977
MONTH WR	gemini [cal]	SOL	1d/1d,4h	+\$820.7	95.5%	13.5%	22	0.977
MONTH WR	gpt, deepseek, gemini [cal]	BNB+XRP	1d/1d,4h	+\$558.7	95.5%	15.4%	22	0.914
MONTH Smoothness	gemini [cal]	SOL	1d/1d,4h	+\$820.7	95.5%	13.5%	22	0.977
MONTH Smoothness	gemini [cal]	SOL	1d/1d,4h	+\$820.7	95.5%	13.5%	22	0.977
MONTH Smoothness	fable, opus, deepseek [cal]	SOL	4h/1h	+\$892.4	94.3%	0.0%	53	0.976
MONTH Quality	fable, opus, deepseek [cal]	SOL	4h/1h	+\$892.4	94.3%	0.0%	53	0.976
MONTH Quality	fable, opus, deepseek [cal]	SOL	4h/1h	+\$905.5	93.1%	1.3%	58	0.970
MONTH Quality	gemini [cal]	SOL	1d/1d,4h	+\$820.7	95.5%	13.5%	22	0.977

[cal] = calibrate=ON cell; the board column names the window. The Calmar board is degenerate by construction -- its top rows sit at 0.0% and 1.3% maxDD, so the score collapses onto net PnL; read it as 'net among configs that barely drew down'. 11 of the 12 printed rows are single-coin SOL configs and 11 of them are calibrated: over 31 days the per-coin branch produces the straightest and most accurate mature lines in the pool, and the >=2-coin rule keeps them out of the cards, not off the boards. The pool is 3,265 mature configs against 3,184 on the four-week window of the same collect.

Stability, and what it is across this issue. The board is defined as presence in the mature net top-40 of BOTH windows of the run that computes it. That run is this issue's collect, whose two windows are the four-week Aug 3-31, 2026 and this calendar month, so the board reads **stable across the four-week window and the calendar month** -- one of the two windows is this issue's own, and the four-week one lies inside it (28 of 31 days), so this is a persistence check inside one month, not a cross-period one. 25 configs qualify, all of them SOL+XRP, 19 of them councils and 14 clean (zero unfunded entries on both windows). The top three by rank sum are solo gemini (+\$1,721.9); solo gemini (+\$1,721.9); solo gemini (+\$1,582.2). Config Watch #4 printed the same board from the same run and read it the same way. Full list in the open data.

Hand-tuned baseline vs the grid

Standing section, carried from the weekly series. One hand-tuned configuration is treated as the reference the search has to beat -- not on net PnL, where any large search wins by construction, but on the three properties a configuration is actually tuned for: **win rate, shallow drawdown and a straight equity line**. The reference (the **hand-tuned anchor**) is the live preset 1d_3_llm_BC2_v02, a 3-model 1d council on SOL+XRP at the \$450 two-coin stake; it is distinct from the pre-search hand-tuned set A/B/C carried in Baselines below. Its replay row is in the walk-forward table above; here it is the yardstick.

Beat rule (frozen with the criteria): a config beats the anchor only if it does so on all three at once -- win rate higher, maxDD lower and smoothness R2 higher -- among mature, non-concentrated search rows. Net PnL is not part of the rule; it is reported next to the winners so the cost of the improvement is visible.

The anchor on this window

Window	Config	Coins	FH/TF	Net	WR	maxDD	R2	Ulcer	Cls
MONTH	Hand-tuned anchor (1d_3_llm_BC2_v02)	SOL+XRP	1d/1d	+\$1,200.9	81.8%	31.9%	0.926	4.16	22

12 of the 3,265 mature configs (0.4%) clear all three bars at once. The anchor's own month is +\$1,200.9 net at 81.8% win rate with a 31.9% drawdown, an R2 of 0.926 and 22 closed trades. Of the ten best of those 12 by net, 1 beats the anchor on net PnL as well; the other 9 buy the three quality bars with a smaller number, and the anchor's +\$1,200.9 is above every one of them. 9 of the ten are calibrated cells, 8 are councils, all ten run the \$450 stake and their universes are 1 coin and 2 coins. On the four-week window of the same collect the same rule admitted 19 of 3,184 configs (0.6%), so the share moves with the window: read it as one field at a time, not as a property of the anchor.

Top-10 by net among the 12 configs that beat the anchor on win rate, maxDD and smoothness at once

#	Config	Coins	FH/TF	Stake	Net	WR	maxDD	R2
#1	fable, opus, deepseek	SOL	4h/1h	\$450	+\$1,335.0	83.0%	9.1%	0.946
#2	fable, opus, deepseek [cal]	SOL	4h/1h	\$450	+\$905.5	93.1%	1.3%	0.970
#3	fable, gpt, deepseek, gemini [cal]	SOL+XRP	1d/4h	\$450	+\$905.3	85.7%	16.6%	0.936
#4	gemini, gpt, deepseek [cal]	SOL+XRP	1d/1d	\$450	+\$902.0	92.9%	13.6%	0.945
#5	fable, opus, deepseek [cal]	SOL	4h/1h	\$450	+\$892.4	94.3%	0.0%	0.976
#6	gemini [cal]	SOL	1d/1d,4h	\$450	+\$820.7	95.5%	13.5%	0.977
#7	gemini [cal]	SOL	1d/1d,4h	\$450	+\$820.7	95.5%	13.5%	0.977
#8	opus, deepseek, gemini, qwen [cal]	SOL+XRP	4h/1h	\$450	+\$750.8	84.3%	13.4%	0.949
#9	opus, gpt, deepseek, gemini, qwen [cal]	SOL	4h/1h	\$450	+\$691.2	85.5%	13.6%	0.972
#10	fable, gpt, deepseek, gemini, qwen [cal]	SOL	4h/1h	\$450	+\$624.5	85.2%	13.6%	0.967

Anchor passport (verbatim, as stored): members=gemini-3.1-pro,gpt-5.6-sol,deepseek-v4-pro tokens=SOL,XRP fh=1d tf=1d conf_mode=each min_conf=60 max_sideways=2 max_diff_side=0 tp_sl_source=nearest trade_mode=position same_side=update opposite=reverse no_signal=hold time_stop=3x steps=40:80,100:20 steps_on=1 be=1 tp_shift=30 sl_shift=50 deposit=1000 lev=10 refill=0 stake=450 calibrate=- variant=v0

Per-ticker bests

Which coin was extractable over the month, and by what. Single-coin results are excluded from the council cards by rule (idiosyncratic-coin risk) and reported here, never ranked against the cards. The dedicated per-coin branch contributed 1,160 sims -- a compacted grid, one best per coin per window; the branch tag in brackets says which kind won.

Coin	Best config (branch)	FH/TF	Net (trades)
BTC	fable [solo]	1d/4h	+\$481.8 (15)
ETH	deepseek, gemini [council]	1d/1d	+\$393.0 (12)
SOL	fable, opus, deepseek [council]	4h/1h	+\$1,344.2 (58)
BNB	opus, qwen [council]	1d/1d	+\$682.2 (15)
XRP	opus, deepseek, gemini [council]	4h/1h	+\$799.9 (73)

n in parentheses; * = fewer than 10 closed trades (thin -- direction only). All five winners run calibrate OFF and 4 of the five are councils.

The reads. Every coin is net-positive over the month and no row is thin: 12 to 73 closed trades, SOL leading at +\$1,344.2 and ETH lowest at +\$393.0. 3 of the five sit below the ≥ 20 -trade maturity gate the council boards apply (BTC 15, ETH 12, BNB 15), so those three rows are direction with a small sample rather than a measured edge. The 1d horizon wins 3 of the five coins. **Winner's curse applies to this whole table:** each row is the maximum of a per-coin grid, biased high by selection alone, and none has an out-of-sample read until the next issue.

Baselines

Three independent reference points, none of them search output, re-simulated fresh on this issue's frozen window inside the same collect (engine 1.1): the engine solo default for each of the 7 models, the default council-of-7, and the pre-search hand-tuned set A/B/C carried since Config Watch #1. The table shows the universe-aware **table-stake** run -- this issue's canonical mode (\$200 for the 5-coin defaults; A and B already at their \$300 3-coin cap; C moves \$400 -> \$450 on 2 coins) -- with the best and worst of the seven solo defaults; all eleven rows and both stake modes are in the open data.

Baseline	Net PnL	Trades	WR	maxDD
Best solo default (gemini-3.1-pro)	-\$38.2 *	459	56%	7.8%
Worst solo default (deepseek-v4-pro)	-\$180.6 *	1028	54%	18.1%
Council-of-7 default	-\$67.6 *	437	57%	7.1%
Hand-tuned config A (pre-search)	+\$379.4 *	26	77%	50.1%
Hand-tuned config B (pre-search)	+\$315.6 *	153	67%	63.0%
Hand-tuned config C (pre-search)	+\$14.9 *	49	61%	45.8%

* = non-zero skipped_nofunds at table stake. On the 5-coin defaults those counts are the stake table doing its job: at \$200 a \$1,000 deposit funds at most five concurrent positions while the defaults fire hundreds of entries over the month (392 to 1,274 skipped each). The other five solo defaults run between the two printed rows, and both stake modes carry the same sign for every one of the eleven rows -- **the conclusions do not change between modes**. Hand-tuned A and B are already at their \$300 3-coin cap, so their two modes are identical; C moves \$400 -> \$450 on two coins and reads +\$82.2 at its own stake against +\$14.9 at the table stake.

Over the month the standing ordering holds in both halves. Every default lost money in both stake modes -- the seven solo defaults run -\$38.2 to -\$180.6 at table stake and the council-of-7 -\$67.6 -- while all three pre-search hand-tuned configs made money (+\$379.4 / +\$315.6 / +\$14.9), and the fresh search's +\$1,618.4 is above the best reference point in the table (+\$379.4). That is search > hand-tuned > defaults, the ordering Config Watch #1-#3 measured and the one Config Watch #4's fresh week inverted in its lower half. The live presets are a separate row of evidence: 4 of the four are net-positive on the month (walk-forward table above), and two of them run stakes above this issue's cap.

Live presets & preset neighbourhoods

Four hand-tuned configurations trade live on the platform (not search output, distinct from the pre-search A/B/C set): preset bc2 = '1d_3_llm_BC2_v02'; preset 4h4 = '4h_4_llm_v02'; preset sol = 'Sol'; preset xrp = 'XRP'. Their replay rows lead the walk-forward table; passports and window detail are in the open data. The neighbourhood grid asks a narrower question: does a one-knob neighbour beat the preset as configured? Two of the four have a local grid this issue (53 cells on this window); the two single-coin presets do not (0 cells -- the local grid is not built for them, so no neighbour claim is made about either).

Live preset	As configured	Best neighbour in local grid	Cells	Verdict
preset bc2	+\$1,200.9	+\$1,442.1 (1d/1d, SOL+XRP, 35 cls)	34	neighbour +\$241.2
preset 4h4	+\$371.6	+\$677.7 (1d/4h, ETH+SOL, 17 cls)	19	neighbour +\$306.1
preset sol	no local grid this issue	--	0	--
preset xrp	no local grid this issue	--	0	--

In **both** preset cells with a grid, a one-knob neighbour beat the configuration as it is actually running: bc2 by +\$241.2 (+\$1,200.9 -> +\$1,442.1 on the same universe) and 4h4 by +\$306.1 (a 4h/1h -> 1d/4h horizon move, +\$371.6 -> +\$677.7). The presets are not at a local optimum on this window -- a testable, low-risk edit, unlike adopting a search champion wholesale. Cell counts are small (34 and 19) and every neighbour is an in-sample maximum of its own little grid.

Patterns: what the finalists look like

Knob modes across this window's unique finalists (mature net top-10 per branch plus all four quality boards, deduplicated by full signature; 62 unique configs, 49 of them councils and 13 solo).

Knob	MONTH finalists mode
Forecast horizon	1d (46/62)
Ladder shape (steps)	50:40,100:60 (59/62)
Break-even stop (BE)	on (62/62)
Min confidence	60 (61/62)
SL shift	75 (59/62)
TP shift	-10 (59/62)
TP/SL source	farthest (54/62)
Calibrate ON	37/62
Universe	SOL+XRP (32/62)

The mechanical core is the one the weekly series has printed for four issues: 50:40,100:60 two-rung ladder (59/62), break-even on (62/62), conf 60 (61/62), SL +75 / TP -10 (59/62), farthest source (54/62). The horizon mode is **1d** (46/62). **Calibrate=ON** runs in 37 of the 62 finalists, where Config Watch #4 measured 35/62 on the four-week window of this same collect. The universe is SOL+XRP in 32 of the 62 finalists, and 12 of them are single-coin rows. Membership across the 49 council finalists: gemini 44, deepseek 41, fable 31, gpt 26, qwen 17, opus 7, grok 1. Standing caveats: refinement seeds from the same grid winners, so part of the convergence is search-design echo, and a knob core that produced 7-of-26 survival on Config Watch #4's two windows and 20-of-26 net-positive on this month is describing two windows, not a recommendation.

Liquidation note

Liquidation is modeled by the engine at the fixed 10x leverage used throughout. Across the published cards, stop-loss placement stays inside the distance that would approach the liquidation threshold; maxDD is printed on every config so realized risk is visible directly. On this window the published cards draw down 40.5% (council #1), 6.4% (council #2) and 54.9% (the solo card), and the replay table is deeper: 18 of the 30 replayed rows drew down more than 40% over the month and 13 past 55%. A high headline net and a survivable path are different claims, and so are a shallow in-sample drawdown and a shallow one next month.

Watch amendments (this issue)

Amendment 5 -- the monthly cadence. The unit of this series is the **calendar month**: one frozen window from the first day of the month to the first day of the next, cut at the month edge plus the maturity lag (Mon Sep 1, 2026, 16:00 UTC). The four-week M30 window and the calendar week stay in the weekly Config Watch and are not reprinted here; the two series share one night collect, and each prints the counts of that collect whole. The walk-forward verdict rule keeps its wording and is evaluated on the one fresh window this issue has (see Walk-forward above).

Amendment 6 -- two standing monthly sections. The **survival curve by generation** and the **survivors-vs-faded structure** table are standing sections of this series. Both cost zero additional simulations: every cell of the curve is read out of a published walk-forward table (this issue reads Config Watch #2, #3 and #4) and the structure table is computed from the frozen queries and the replay statistics of the cards Config Watch #4 already replayed. Each monthly issue appends the fresh weeks published since and the generation published since.

Both amendments are provisional and scoped to this series. **Formalization is slated for methodology v1.2**; until then this issue runs on v1.1 plus the Config Watch amendments approved 12.08 (hash e66c7e8c864a2233).

Limitations

- **One calendar month, one regime.** Every number here is a single 31-day window (Aug 1-31, 2026) on a single regime, and this is issue #1 of the series: there is no prior monthly window, so no month-over-month figure is printed and no past issue is restated.
- **This month contains the weekly windows, so it is not independent of them.** All four published Config Watch weeks lie inside it and so does the whole 28-day four-week window; the month, the week and the four weeks are three readings of overlapping data, and agreement between them is partly arithmetic.
- **The replay column is decay context, not clean out-of-sample.** The issue-1, issue-2 and issue-3 search windows all end inside this month, so 26 of the 26 replayed cards are partly scored on the days they were selected on. The four live presets are not search output and have no home issue, and 2 of them run stakes above this issue's cap for their universe size ([conc] rows) -- their printed economics assume funding the current rule would not grant.
- **The survival curve mixes window lengths and calendar weeks.** Each +kw cell is a 7-day window from a different issue and a different week of August, the Calendar Aug column is a 31-day window, and the live-preset row has no publication date at all. Read down a column, not across a row, and not as a decay rate.
- **Stability is defined inside the run that computes it.** This board comes from this issue's collect and therefore reads across the four-week window and the calendar month -- two windows of which one contains the other -- so it measures persistence inside one month and is not comparable with a board built across two disjoint periods.
- **Research-to-date counter is pinned, and it is read at this issue's cutoff.** The counter below is pinned from wave 3 onward and prints 39,073 resolved forecasts at the monthly cutoff (Mon Sep 1, 2026, 16:00 UTC), where Config Watch #4 printed 38,385 at its own Aug 31 cutoff: the same counter at two cutoffs, not two measurements. The model line counts 7 current models with earlier versions folded into their successors' lineage.
- **Multiple testing, winner's curse and thin cells.** 10,889 simulations in this collect; at this scale some winners are expected from chance alone, and every card, board row and per-ticker best is the maximum of a search, biased high by selection. On this window 0 replay rows are under 10 closed trades and 3 of the five per-ticker winners are under the ≥ 20 -trade maturity gate. Antidotes: axes frozen before launch, the walk-forward table, independent baselines, in-sample labeling; no formal correction yet. Research output, not financial advice.

Counters & lineage

Night collect (the four-week and calendar-month windows together; the four-week boards ship with Config Watch #4)

Stage	Sims	Stage	Sims
Scan (s1)	952	Dedicated per-ticker (s6t+s6tc)	1,160
Systematic grid (s2)	8,224	OOS replays (s4cw3+s4cw3o+s4cw2+s4cw2o+s4cw1+s4pre)	74
Refinement: ladder + stake sweep (s2b+s3b)	243	Total	10,889
Preset neighbourhood grid (s2p)	106	Errors / zero-trade	0 / 1,188
Per-ticker probes (s5t)	130	Engine	1.1

Scale note. This issue and Config Watch #4 share one collect: the night run is 10,889 simulations over its two windows, and the weekly issue additionally carries a 5,558-simulation weekly collect that this issue does not use (16,447 across the two collects there). The per-window stage-2 grid is 4,112 cells here and 4,208 on the weekly issue's fresh week. The replay branch is 30 rows -- 26 past cards and 4 live presets -- for 74 simulations, because every row is replayed on both windows of this collect. The two selection boards and both new sections of this issue cost zero simulations.

No extras pass. Equity series, the grid histogram (0 extra sims), the baselines and the calibrate pack are produced by the same collect that wrote the boards. The baseline runs are verification, outside the search total: 44 sims in this collect (11 configs x 2 windows x 2 stake modes), of which the 22 calendar-month rows belong to this issue and the other 22 to Config Watch #4. 0 equity replays were needed.

Source files: summary_night.json (the w200m collect) with grid_hist_night, equity_night, baselines_night, presets_recon_night and the three oos_cw*_night packs beside it; the published open data of Config Watch #2, #3 and #4 for every number this issue quotes from them. Generated at: 2026-09-03T12:07:22Z · Methodology v1.1 + Config Watch amendments (approved 12.08), hash e66c7e8c864a2233.

Snapshot principle: this report is a frozen snapshot of the search taken at the generated-at timestamp above; numbers are not updated retroactively. Each issue re-runs the pipeline fresh over that issue's window.

Research to date

RESEARCH SNAPSHOT

THIS REPORT -- search effort: 10,889 simulations (0 errors) in one frozen collect, of which 74 walk-forward replays and origin re-runs · 44 verification sims (baseline runs inside the same collect; 0 equity replays needed)

MARKETMANIA RESEARCH TO DATE (as of cutoff Sep 1, 16:00 UTC): 39,073 resolved since Jul 11 · 7 models tracked (current line-up; earlier versions folded into their successors' lineage) · 5 assets · 5 horizons · hourly · 22 published reports

Counter pinned from wave 3 onward and read at this issue's cutoff (Mon Sep 1, 2026, 16:00 UTC), where Config Watch #4 read 38,385 at the Aug 31 cutoff; the model line counts the current line-up, with earlier versions folded into their successors' lineage. Monthly Config Watch #1 is the 22th published report: Monthly Consensus Watch #1, Monthly Calibration #1 and Monthly Model Watch #1 are 19, 20 and 21, and Config Watch #4 is 18.

Issue #1. Monthly Config Watch is a living monthly comparison: each issue appends one more calendar month of cards, walk-forward verdicts and survival data, and nothing in a past issue is restated. This first issue sets the baseline -- a best mature council of +\$1,618.4 and a best mature solo of +\$1,721.9 on August, 20 of 26 published cards net-positive on the month, a survival curve that reaches +3 fresh weeks for the issue-1 generation, and a structural table of Config Watch #4's 7 survivors against its 19 faded cards. The calendar week and the four-week window stay in the weekly Config Watch.

What we're testing next

- **Monthly Config Watch #2** -- the September calendar month (2026-09-01 to 2026-10-01) from the same engine, the same criteria and the same collect code. It carries the first month-over-month comparison this series can make, and the first replay rows for Config Watch #4's and #5's own cards.
- **The survival curve extended.** Every weekly issue adds one fresh week to every generation already on the curve and puts a new generation on it; Monthly #2 prints those columns for the September weeks, which is the first point at which a generation has four fresh-week readings.
- **The quarterly** -- the first cross-issue aggregate of the walk-forward table: four weekly issues and this monthly issue of verdicts on one axis, with the survival curve as its spine.

Related research

Monthly Consensus Watch #1 -- <https://marketmania.ai/research/reports/consensus-watch-monthly-2026-08.pdf>

Monthly Calibration #1 -- <https://marketmania.ai/research/reports/calibration-monthly-2026-08.pdf>

Monthly Model Watch #1 -- <https://marketmania.ai/research/reports/model-watch-monthly-2026-08.pdf>

Config Watch #4 -- <https://marketmania.ai/research/reports/config-watch-2026-08-26.pdf>

The three monthly reports publish together as one issue each month; this issue joins them on the same cadence. Config Watch #4 is the weekly issue this one shares a collect with: it carries the calendar week Aug 24-31, 2026 and the four-week window Aug 3-31, 2026, and no board of either is reprinted here.

CITE THIS REPORT

```
@techreport{mm_configwatch_monthly_2026m08,
  title      = {Monthly Config Watch #1},
  author     = {{MarketMania Research}},
  institution = {MarketMania},
  year       = {2026},
  month      = sep,
  day        = {4},
  type       = {Monthly Research Report},
  series     = {Monthly Config Watch},
  number     = {1},
  note       = {Methodology v1.1 + Config Watch amendments, hash e66c7e8c864a2233},
  url        = {https://marketmania.ai/research/reports/config-watch-monthly-2026-08.pdf}
}
```