

Monthly Config Watch #2

October 3, 2026 · MarketMania Research · Monthly series

MONTHLY · CONFIG WATCH · Issue #2 · one calendar month, 57 replays, one survival curve · Issue date October 3, 2026 · Language English · Methodology v1.1 + Config Watch amendments (approved 12.08, hash e66c7e8c864a2233)
Search window: **MONTH** Sep 1-30, 2026 -- the frozen calendar month (Tue Sep 1 00:00 -> Wed Sep 30 23:59 UTC, 30 days), cut-off edge Thu Oct 1 00:00 UTC, so slots dated Oct 1 fall outside it; report cut-off Thu Oct 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.

Prior issue: Monthly Config Watch #1 (September 4, 2026, calendar August, Aug 1-31, 2026) -- best mature council +\$1,618.4, best mature solo +\$1,721.9, 20 of 26 published cards net-positive on that month; every month-over-month number below is read from its open data. The weekly Config Watch #8 carries the calendar week Sep 21-28, 2026 and the four-week window Aug 31-Sep 28, 2026 from two collects of its own; this issue's collect recomputes that four-week window beside the month, and neither weekly board is reprinted here.

RESEARCH SNAPSHOT

11,075 simulations, 0 errors, in one frozen collect	MONTH Sep 1-30, 2026	cut-off edge Thu Oct 1 00:00 UTC ; report cut-off Thu Oct 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 (w298 plan)
walk-forward: 53 past cards + 4 live presets replayed verbatim on the month	17 of 53 past cards net-positive on the month	stage-2 grid: 4,176 cells on this window	models: 7 (grok axis = grok-4.6 (incl. 4.5-era))
best MONTH council +\$1,041.5 (+104.2%)	best MONTH solo +\$923.2 (+92.3%)	zero-trade cells: 1,686 (whole collect)	margin-skip alerts on leaders: 14 rows (MONTH)

KEY FINDING · OBSERVATION (one calendar month)

The best mature council on calendar September makes +\$1,041.5 (+104.2% of a \$1,000 deposit, 71.1% win rate, 37.7% maxDD, 97 closed), 17 of the 53 cards this series has published are net-positive on the same month, and 7 of the 10 cards that survived both fresh windows of Config Watch #8 are among them.

WEEKLY EVIDENCE -> MONTHLY VERDICT (rule below)

Metric	CW#5 wk	CW#6 wk	CW#7 wk	CW#8 wk	Month	Verdict
Published cards net-positive (majority)	8/34	8/38	33/43	24/48	17/53	Confirmed (3/4)
Best council universe	SOL+XRP	SOL+XRP	SOL+XRP	SOL+XRP	BNB+XRP	Rejected (0/4)
Calibration helps (matched pairs)	57.8%	49.9%	25.8%	36.5%	65.6%	Rejected (1/4)


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-September pull.

Sources, cell by cell. Row 1 counts the replays each issue ran on its own fresh week -- Config Watch #5 8 of 34 (Aug 31-Sep 7, 2026), #6 8 of 38 (Sep 7-14, 2026), #7 33 of 43 (Sep 14-21, 2026), #8 24 of 48 (Sep 21-28, 2026) -- against 17 of 53 on the month in the walk-forward table below. 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, over the whole stage-2 grid in every column: 1,403 pairs in issue #5, 1,334 in #6, 1,452 in #7 and 1,468 in #8, against 1,446 whole-grid pairs on the month. No cell reads n/a this issue: all four issues carry a replay table and a calibration axis. Config Watch #5's week opens on Aug 31, before the month, so 6 of its 7 days are September days. The four weeks cover 27 of the month's 30 days; Sep 28-30 lie after Config Watch #8's week.

TL;DR

- **OBSERVATION -- 17 of the 53 published cards are net-positive on the calendar month** (calendar August: 20 of 26). Config Watch #8 read the 48 cards it inherited on its fresh week (Sep 21-28, 2026): 24 positive; on its four-week window (Aug 31-Sep 28, 2026): 16. By generation on the month: issue-8 2/5, issue-7 4/5, issue-6 4/5, issue-5 2/4, issue-4 2/5, Monthly-1 1/3, issue-3 0/6, issue-2 0/8, issue-1 2/12; live presets 0/4. On the 26 cards published before the month began (issue-3, issue-2 and issue-1) the reading is 2 of 26. The month holds 7 of that week's 7 days and 27 of that four-week window's 28, so the readings are not independent tests.
- **The month's ceiling is +\$1,041.5 on the council board and +\$923.2 on the solo board** (calendar August: +\$1,618.4 and +\$1,721.9). The council card runs at +104.2% of a \$1,000 deposit, 71.1% win rate and 37.7% maxDD on 97 closed trades; the solo card at +92.3%, 85.8% win rate and 33.0% maxDD on 204 closed with 3 unfunded entries. The grid behind them: median -\$5.7 and 29.6% positive across 4,176 stage-2 cells, 117 of them above +\$500, range -\$802.6 to +\$1,041.5. On the four-week window of the same collect that grid ran 29.9% positive; the month leader reads +\$943.0 on the four weeks against +\$1,041.5 on the month, and Config Watch #8's four-week leader reads +\$1,021.7 on the month against +\$1,018.6 on the four weeks (section below).
- **The survival curve reaches +7 fresh weeks for the issue-1 generation** (+3 in Monthly Config Watch #1). Share of a generation's cards net-positive on its latest fresh week: issue-1 8.3% (+7); issue-2 75.0% (+6); issue-3 33.3% (+5); issue-4 80.0% (+4); Monthly-1 100.0% (+4); issue-5 50.0% (+3); issue-6 40.0% (+2); issue-7 80.0% (+1); issue-8 has no fresh week yet. Every cell is read out of the issue that ran that replay, and the points sit on 7 different calendar weeks, from Aug 10-17, 2026 to Sep 21-28, 2026.
- **The 10 survivors and the 38 faded cards of Config Watch #8, side by side.** The cards that cleared both of its fresh windows against those that did not: median fresh-week maxDD 19.1% vs 35.2%, median win rate 68.0% vs 61.5%, 4 of 10 against 15 of 38 on the 1d horizon, and 4 of 10 (40.0%) vs 5 of 38 (13.2%) calibrated. 10 rows against 38 rows on two windows is a description of two groups, not a test, and none of the traits is claimed to cause survival.

Survival by generation -- cards net-positive on the k-th fresh week



Share of each generation's cards net-positive on the k-th fresh week after its own window edge (issue-1 12 cards, issue-2 8 cards, issue-3 6 cards, issue-4 5 cards, Monthly-1 3 cards, issue-5 4 cards, issue-6 5 cards, issue-7 5 cards). Each point is one published walk-forward table: issue-1's seven points are Config Watch #2 to #8, issue-7's one point is #8; the count behind every point is in the survival table below. The x axis counts fresh weeks, not calendar weeks -- the points sit on 7 different weeks, from Aug 10-17, 2026 to Sep 21-28, 2026 -- and neither the live presets (no publication date), issue-8 (no fresh week yet), nor the calendar-month column (a 30-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 30-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 standing sections of the series: the survival curve by generation, which places each published card's replay on the k-th fresh week after its own issue and now covers 9 generations, and the structural comparison of the 10 cards that cleared both fresh windows of Config Watch #8 against the 38 that did not. The month is not an independent test of the weeks: 27 of its 30 days lie in the four published Config Watch weeks it is read against and 27 in Config Watch #8's four-week window, and every section that reads them together says so.

How we searched

One frozen collect, run to the plan with the criteria frozen in writing *before* launch (w298 plan) and printed into the meta of the JSON deliverable: the month run, which sweeps the four-week window of Config Watch #8 and the calendar month in one pass. Every stage count below is that run whole. It is a run of its own -- no part of it ships with a weekly issue -- and the stage-2 grid is 4,176 cells on each of its two windows.

- **1. Scan (s1)** -- 952 sims across council compositions and coarse settings.
- **2. Systematic grid (s2)** -- 8,352 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)** -- 240 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)** -- 120 sims: best finalists split onto single coins.
- **6. Dedicated per-ticker (s6t+s6tc)** -- 1,150 sims: a compacted independent per-coin grid plus calibrate twins (the standing branch from Config Watch #2).
- **7. OOS replays (s4cw8+s4cw8o+s4cw7+s4cw7o+s4cw6+s4cw6o+s4cw5+s4cw5o+s4cw4+s4cw4o+s4cw1+s4cwm1+s4cwm1o+s4cw3+s4cw3o+s4cw2+s4cw2o+s4cw1+s4cpre)** -- 155 sims: every published card and all four live presets, verbatim, on both windows of this collect, plus origin re-runs of the cards of issue #8, issue #7, issue #6, issue #5, issue #4, Monthly #1, issue #3 and issue #2 for the drift baseline.

Total 11,075 simulations, 0 errors, 1,686 zero-trade cells; stage sums reconcile exactly inside the collect. The stage-2 grid is 4,176 cells on this window against 4,112 on Monthly Config Watch #1's. 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: 14, of the 28 the collect raised over its two windows; the other 14 are on its four-week board, which is not printed here.

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; both windows of this collect (Aug 31-Sep 28, 2026 and Sep 1-30, 2026) lie entirely after the grok-4.5 -> 4.6 cutover [dated Aug 24 in the search plan, and published grok-4.5 configs replayed out-of-sample get the same 4.5 -> 4.6 lineage remap ([remap] rows).

TOP councils -- calendar month (Sep 1-30, 2026)

#	Council	Coins	FH/TF	Net	Net %	maxDD	Trades	WR
#1	fable, opus, gemini, qwen	BNB+XRP	4h/1h	+\$1,041.5	+104.2%	37.7%	97	71.1%
#2	fable, opus, gemini, qwen	SOL+XRP	4h/1h	+\$1,021.7	+102.2%	47.7%	90 *	72.2%
#	Ladder	BE	Shifts	Conf filter	Source	Entry	Positional	Stake
#1	50:40,100:60	on	TP -10 / SL +75	>=60 (each)	farthest	2/0	reenter	\$450
#2	50:40,100:60	on	TP -10 / SL +75	>=60 (each)	farthest	2/0	reenter	\$450

Against the 4,176-cell stage-2 grid on this window: card #1 is the same simulation as the raw net-board leader (i=7089) and the top cell of the grid; card #2 lands in the top bucket of the grid, +\$1,000 to +\$1,250 (2 cells). Grid median -\$5.7, 117 of the 4,176 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. The mature net board's 10 slots carry 2 universes (BNB+XRP 5 and SOL+XRP 5); card #2 is the best mature multi-coin council outside card #1's universe -- it comes off the **mature net** board (+\$1,021.7 on SOL+XRP, 47.7% maxDD, R2 0.821, 10 unfunded entries). Both councils are configurations Config Watch #8 published: card #1 is cw8-m302 and card #2 is cw8-m301, query for query, window aside, and their walk-forward rows below read the same numbers on the month.

The two cards answer the calibrate axis the same way: no card is a calibrated cell (Calibrate axis below). Card #1 runs the \$450 two-coin table stake and card #2 the \$450 two-coin stake; card #1 carries no unfunded entry and card #2 10 unfunded entries, and card #1 draws down 37.7% on 97 closed trades against card #2's 47.7% on 90. The mature pool on this window is 2,967 configs against 2,911 on the four-week window of the same collect (30 days against 28) and 3,265 on calendar August.

OPEN IN SANDBOX marketmania.ai/s/cwm2-m1 marketmania.ai/s/cwm2-m2


Each link opens this exact frozen window; results visible without sign-in (embedded share signature). The codes are minted by the publish script before this file goes live; every card's full frozen query signature is in the open-data JSON next to its simulation index.

Published card equity -- MONTH (Sep 1-30, 2026) -- \$1,000 at 10x -- 97 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 luft); 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=7089), so a second curve would be the same line. The line peaks at +\$1,365.4 on 2026-09-22 and ends the window at +\$1,041.5. Card #2, the solo card and the other 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 (Sep 1-30, 2026) -- n=4,176



Net PnL across all 4,176 stage-2 grid sims, MONTH window (Sep 1-30, 2026), universe-aware stake throughout. Grid median **-\$5.7**, 29.6% positive (calendar August: \$0.0, 41.8% positive); the same grid on the four-week window of this collect ran 29.9% positive and on Config Watch #8's fresh week 47.7%. The \$0 bin holds 1,623 cells; 842 sims on this window produced no trade at all. Range -\$802.6 to +\$1,041.5; 117 cells finished above +\$500 and the maximum of the grid is council card #1. Red marker = best solo default at table stake; green = the published cards.

The 30-day cards against the calendar month

Config Watch #8 published two councils on its four-week M30 window (Aug 31-Sep 28, 2026, 28 days, edge Mon Sep 28 00:00 UTC); this issue publishes two councils and one solo on the calendar month (Sep 1-30, 2026, edge Thu Oct 1 00:00 UTC). This collect replays Config Watch #8's cards on both of its windows, and the two windows share 27 days -- Aug 31 lies in M30 only and Sep 28-30 in the month only. The table measures all five cards on both windows: M30 is the four-week window and Month is this issue's.

Card	Council / solo	Coins	FH/TF	Net M30	Net month	WR M30/Mo	maxDD M30/Mo	Cls
cw8-m301	fable, opus, gemini, qwen	SOL+XRP	4h/1h	+\$1,018.6	+\$1,021.7	72.5% / 72.2%	48.5% / 47.7%	91 / 90
cw8-m302	fable, opus, gemini, qwen	BNB+XRP	4h/1h	+\$943.0	+\$1,041.5	69.7% / 71.1%	37.7% / 37.7%	99 / 97
cwm2-m1	fable, opus, gemini, qwen	BNB+XRP	4h/1h	+\$943.0	+\$1,041.5	69.7% / 71.1%	37.7% / 37.7%	99 / 97
cwm2-m2	fable, opus, gemini, qwen	SOL+XRP	4h/1h	+\$1,018.6	+\$1,021.7	72.5% / 72.2%	48.5% / 47.7%	91 / 90
cwm2-s1	gemini	SOL+XRP	4h/1h	+\$711.8	+\$923.2	83.9% / 85.8%	33.2% / 33.0%	199 / 204

Every paired column reads **M30 first, calendar month second**; Cls = closed trades. **cw8-m301 and cwm2-m2 are one configuration**; **cw8-m302 and cwm2-m1 are one configuration**, so those rows carry the same pair of numbers; the five rows are three distinct configurations. Config Watch #8's rows are this collect's replay of the published query on both windows; Config Watch #8 printed cw8-m301 +\$1,018.6 and cw8-m302 +\$943.0 for them from its own night collect, and this collect's replay reproduces each published figure on the same four weeks. Place on each window's mature net board (councils and solos in one list, as the collect ranked them): cw8-m301 1st on M30 and 2nd on the month; cw8-m302 2nd on M30 and 1st on the month; cwm2-m1 2nd on M30 and 1st on the month; cwm2-m2 1st on M30 and 2nd on the month; cwm2-s1 24th on M30 and 7th on the month. Every cell is in this collect. Boards, the replay set and a pair sample are what the collect exports, not every cell on both windows.


All five rows can be read on both windows, and they differ by \$3.2 to \$211.4 between them. cw8-m301 makes +\$1,018.6 on M30 and +\$1,021.7 on the month (+\$3.2, on 90 closed trades against 91); cw8-m302 makes +\$943.0 on M30 and +\$1,041.5 on the month (+\$98.5, on 97 closed trades against 99); cwm2-m1 makes +\$943.0 on M30 and +\$1,041.5 on the month (+\$98.5, on 97 closed trades against 99); cwm2-m2 makes +\$1,018.6 on M30 and +\$1,021.7 on the month (+\$3.2, on 90 closed trades against 91); cwm2-s1 makes +\$711.8 on M30 and +\$923.2 on the month (+\$211.4, on 204 closed trades against 199). 3 of them are larger on the window their own issue published them on.

Walk-forward on the calendar month

Every config this series has ever published -- all 53 cards from Config Watch #1 to #8 and Monthly Config Watch #1 -- 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 155 simulations inside this collect, which replays every row on both of its windows in one pass. Config Watch #8's own five cards are in the table for the first time; their search windows share 27 of 28 and 7 of 7 days with this month (the rule below counts every cohort).

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 to #8 applied. The column is not pure out-of-sample either: the month shares days with the search windows of issue-8, issue-7, issue-6 and issue-5 (days inside the month: issue-8's 28-day window 27 and 7-day window 7; issue-7's 28-day window 20 and 7-day window 7; issue-6's 28-day window 13 and 7-day window 7; issue-5's 28-day window 6 and 7-day window 6), so for those 19 rows the number a card shows here is partly scored on the days it was selected on; the search windows of issue-4, Monthly-1, issue-3, issue-2 and issue-1 share no day with it. Read the column as decay context over a 30-day window.

Config Watch #8's five published cards: home number vs calendar-month replay



Config Watch #8's five 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). 2 of the five are net-positive on the month. The month holds 27 of the 28 days of the m301 and m302 cards' window and 7 of the 7 days of the s1, wk1 and wk2 cards' window, so the green bars are scored mostly on days the cards were selected on.

Config	Family	Coins	Home	New MONTH (n)	Verdict
preset 4h4	fable, qwen, deepseek, opus	ETH+SOL	-- live	-\$20.4 (43) +	Faded
preset bc2	gemini, gpt, deepseek	SOL+XRP	-- live	-\$600.5 (6) * +	Faded
preset sol	deepseek, fable, opus	SOL [conc]	-- live	-\$60.8 (52)	Faded
preset xrp	gemini, gpt, deepseek	XRP [conc]	-- live	-\$387.0 (1) * +	Faded
cw8-m301	fable, opus, gemini, qwen	SOL+XRP	+\$1,018.6	+\$1,021.7 (90) +	Survived
cw8-m302	fable, opus, gemini, qwen	BNB+XRP	+\$943.0	+\$1,041.5 (97)	Survived
cw8-s1	gemini	BNB+XRP	+\$445.4	-\$252.7 (165) +	Faded
cw8-wk1	gpt, gemini, qwen	SOL+XRP	+\$583.4	-\$571.4 (14) +	Faded
cw8-wk2	gpt, gemini, qwen, grok	BNB+XRP	+\$570.6	-\$554.4 (34) +	Faded
cw7-m301	fable, gpt, gemini, qwen	SOL+XRP	+\$1,351.8	+\$960.8 (139) +	Survived
cw7-m302	opus, gpt, gemini, qwen	BTC+ETH	+\$528.4	+\$221.4 (31)	Survived
cw7-s1	gemini	SOL+XRP	+\$1,022.0	+\$773.0 (202) +	Survived
cw7-wk1	fable, opus, gpt, gemini	SOL+XRP	+\$1,467.8	+\$305.1 (161) +	Survived
cw7-wk2	fable, opus, gpt	BTC+ETH	+\$953.2	-\$606.7 (23) +	Faded
cw6-m301	opus, deepseek, gemini	SOL+XRP	+\$1,793.8	+\$544.0 (169) +	Survived
cw6-m302	gpt, deepseek, gemini, grok	BTC+ETH	+\$547.8	-\$229.9 (35) +	Faded
cw6-s1	grok	SOL+XRP	+\$404.8	+\$287.9 (101) +	Survived
cw6-wk1	opus, gpt, grok	SOL+XRP	+\$609.5	+\$411.1 (163) +	Survived
cw6-wk2	fable, opus, deepseek	BTC+ETH	+\$135.4	+\$159.8 (90)	Survived
cw5-m301	fable, opus, gpt, gemini	SOL+XRP	+\$1,717.1	-\$597.0 (9) * +	Faded
cw5-s1	gemini	BNB+XRP	+\$178.7	-\$49.1 (177) +	Faded
cw5-wk1	gpt, gemini	SOL+XRP	+\$510.4	+\$1,085.3 (280)	Survived
cw5-wk2	fable, opus, gpt, gemini, grok	BTC+ETH	+\$183.5	+\$44.3 (112)	Survived
cw4-m301	gpt, deepseek, gemini	SOL+XRP	+\$1,676.0	-\$76.2 (28) +	Faded
cw4-m302	fable, gpt, deepseek, gemini	BNB+XRP	+\$1,155.2	-\$581.9 (10) +	Faded
cw4-s1	deepseek	BTC+ETH+ SOL+BNB+XRP	+\$325.0	-\$289.2 (60) +	Faded

Config	Family	Coins	Home	New MONTH (n)	Verdict
cw4-wk1	opus, deepseek, gemini, qwen	BTC+ETH+SOL+BNB+XRP	+\$399.4	+\$49.3 (85) +	Survived
cw4-wk2	fable, deepseek, qwen	BNB+XRP	+\$128.2	+\$397.2 (97)	Survived
cwm1-m1	fable, gemini	SOL+XRP	+\$1,618.4	-\$670.3 (16) +	Faded
cwm1-m2	fable, gpt, deepseek, gemini, qwen	ETH+SOL+BNB+XRP	+\$756.6	+\$17.6 (55) +	Survived
cwm1-s1	gemini	SOL+XRP	+\$1,721.9	-\$567.3 (13) +	Faded
cw3-m301	opus, deepseek, gemini	SOL+XRP	+\$2,137.6	-\$279.2 (100) +	Faded
cw3-m302	fable, opus, deepseek, gemini, qwen	BTC+ETH	+\$412.3	-\$247.2 (14) +	Faded
cw3-s1	gemini	SOL+XRP	+\$1,604.3	-\$763.1 (66) +	Faded
cw3-s2	gemini	SOL+XRP	+\$1,449.0	-\$567.3 (13) +	Faded
cw3-wk1	opus, deepseek, gemini	SOL+XRP	+\$2,175.1	-\$279.2 (100) +	Faded
cw3-wk2	opus, deepseek, gemini, grok	BTC+ETH	+\$837.6	-\$505.2 (64) +	Faded
cw2-30d1	opus, gpt, deepseek, gemini, qwen	BTC+ETH+SOL+BNB+XRP [conc]	+\$744.7	-\$727.3 (12) +	Faded
cw2-30d2	opus, deepseek, gemini	ETH+SOL+BNB+XRP [conc]	+\$655.1	-\$343.6 (27) +	Faded
cw2-30d3	opus, deepseek, gemini	BNB+XRP	+\$489.1	-\$168.4 (31) +	Faded
cw2-30ds1	gemini	BNB+XRP	+\$383.8	-\$660.8 (14) +	Faded
cw2-7d1	fable, opus, deepseek, gemini	SOL+XRP	+\$299.1	-\$73.3 (104) +	Faded
cw2-7d2	gpt, deepseek, gemini, qwen	BNB+XRP	+\$223.7	-\$124.7 (31) +	Faded
cw2-pt1	fable, deepseek, grok [remap]	SOL	+\$698.3	-\$67.1 (63)	Faded
cw2-pt2	opus	XRP	+\$543.7	-\$559.8 (5) * +	Faded
cw1-30d1	fable, deepseek, opus, qwen	ETH+SOL	+\$868.4	-\$19.8 (46) +	Faded
cw1-30d2	gemini, gpt, deepseek	SOL+XRP+ETH	+\$590.6	-\$377.9 (32) +	Faded
cw1-30d3	fable, deepseek, opus, qwen	BTC+ETH+SOL	+\$580.5	-\$73.2 (71) +	Faded
cw1-30ds1	opus	SOL+XRP	+\$442.8	+\$192.4 (106) +	Survived
cw1-30ds2	grok [remap]	SOL+XRP	+\$262.9	-\$276.5 (77) +	Faded
cw1-30ds3	deepseek	XRP+BNB	+\$119.0	-\$176.7 (108) +	Faded
cw1-7d1	fable, deepseek, opus, gpt	SOL+XRP	+\$439.2	+\$103.8 (150)	Survived
cw1-7d2	fable, qwen, gemini, opus	XRP+BNB	+\$429.5	-\$257.4 (89) +	Faded
cw1-7d3	gpt, opus	BNB+XRP+SOL	+\$398.6	-\$706.8 (128) +	Faded
cw1-7ds1	gemini	XRP+BNB	+\$417.8	-\$94.1 (214) +	Faded
cw1-7ds2	gpt	XRP+BNB	+\$366.3	-\$506.7 (116) +	Faded
cw1-7ds3	qwen	XRP+BNB	+\$267.0	-\$594.6 (90) +	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: 17 of 53 past cards are net-positive on the calendar month (calendar August: 20 of 26) -- by cohort 2 of 5 issue-8 cards, 4 of 5 issue-7 cards, 4 of 5 issue-6 cards, 2 of 4 issue-5 cards, 2 of 5 issue-4 cards, 1 of 3 Monthly-1 cards, 0 of 6 issue-3 cards, 0 of 8 issue-2 cards, 2 of 12 issue-1 cards. Config Watch #8 read 48 of these cards (all but its own 5) on two windows that lie mostly inside this one: 24 of 48 on its fresh week, 16 of 48 on its

four weeks, 10 surviving both. 7 of the 10 of those survivors are net-positive here, and so are 8 of its 38 faded cards. On this collect's own four-week window the 53 past cards read 20. The best replay is cw5-wk1 at +\$1,085.3 on 280 closed trades and the worst cw3-s1 at -\$763.1 on 66. Live presets: 0 of 4 net-positive on the month (calendar August: 4/4), 0 of 4 survived both of Config Watch #8's windows.

The four live presets lead the table by standing rule. On the month: sol -\$60.8 on 52 closed, xrp -\$387.0 on 1 closed, bc2 -\$600.5 on 6 closed, 4h4 -\$20.4 on 43 closed -- none of the four net-positive, where Config Watch #8 recorded 2/4 on its fresh week and Monthly Config Watch #1 4/4 on calendar August. **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 -\$600.5 is the same simulation on both sides of the report.

Origin re-runs (drift baseline). Re-simulating the issue-8 cards on their OWN origin windows in this collect lands within -\$162.0..\$0.0 of the numbers issue #8 printed: 3 of the 5 reproduce to the cent. The issue-7 cards drift -\$175.6..\$0.0 (1 of 5 to the cent). The largest gap is cw7-m302 (+\$528.4 printed, +\$352.8 on the re-run). Snapshot principle: no past issue is restated; the full drift table is in the open data.

What this column is. 17 of 53 past cards are net-positive over 30 days. Three measured facts bound the reading. The window overlaps the cards' own: four of the nine cohorts searched on days inside it (19 rows). The rows Config Watch #8 dropped on its two-window rule did not all lose here: 8 of its 38 faded cards are net-positive on the month, and 7 of its 10 survivors. And the fresh search moved with them: the search grid's month ceiling is +\$1,041.5 against +\$1,018.6 on the four-week window of this collect and +\$583.4 on Config Watch #8's 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 counts that generation's cards that finished net-positive on the k-th fresh week after the issue's own window edge, out of all of its cards, 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 eight card rows of this table that have a fresh week, as shares; it is printed once, there.

Generation	Cards	+1 week	+2 weeks	+3 weeks	+4 weeks	+5 weeks	+6 weeks	+7 weeks	Calendar Sep
Config Watch #1	12	6/12	11/12	1/12	2/12	3/12	11/12	1/12	2/12
Config Watch #2	8	8/8	4/8	1/8	2/8	5/8	6/8	--	0/8
Config Watch #3	6	2/6	2/6	2/6	5/6	2/6	--	--	0/6
Config Watch #4	5	1/5	0/5	4/5	4/5	--	--	--	2/5
Monthly Config Watch #1	3	2/3	0/3	1/3	3/3	--	--	--	1/3
Config Watch #5	4	1/4	3/4	2/4	--	--	--	--	2/4
Config Watch #6	5	4/5	2/5	--	--	--	--	--	4/5
Config Watch #7	5	4/5	--	--	--	--	--	--	4/5
Config Watch #8	5	--	--	--	--	--	--	--	2/5
Live presets	4	2/2	2/4	1/4	3/4	3/4	2/4	--	0/4

Cell by cell: every +k cell is the walk-forward of the issue whose fresh week it is -- Config Watch #2 (Aug 10-17, 2026), Config Watch #3 (Aug 17-24, 2026), Config Watch #4 (Aug 24-31, 2026), Config Watch #5 (Aug 31-Sep 7, 2026), Config Watch #6 (Sep 7-14, 2026), Config Watch #7 (Sep 14-21, 2026), Config Watch #8 (Sep 21-28, 2026) -- and a generation's +1 week is the first of those weeks after its own window edge: issue-1 at #2, issue-2 at #3, issue-3 at #4, issue-4 and Monthly-1 at #5, issue-5 at #6, issue-6 at #7, issue-7 at #8. issue-8's cards have no fresh week yet and only the Calendar Sep cell. The live presets are not a generation and have no publication date: their six cells are the six issues that replayed them -- Config Watch #3 replayed 2 of the 4, #4 to #8 all 4 -- and that row is not drawn on the chart. Calendar Sep is this issue's window (Sep 1-30, 2026) for every row; Monthly Config Watch #1 read the issue-1 row at 8/12, the issue-2 row at 8/8, the issue-3 row at 4/6 and the presets row at 4/4 on calendar August. 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-7's is Sep 21-28, 2026), and issue-1's +1 week overlaps its own search window (Aug 5-12, 2026) by 2 days, Monthly-1's overlaps calendar August by 1 day (Aug 31).

Survivors vs faded: structure

Config Watch #8 replayed 48 past cards on two fresh windows (Sep 21-28, 2026 and Aug 31-Sep 28, 2026) and its standing rule kept 10 of them (net-positive on both) and dropped 38. 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. Neither window decided alone: 24 cards cleared the week, 16 the four weeks and 10 both. 10 rows against 38 rows: this is a description of two groups, not a test, and no trait here is claimed to cause survival.

Trait	Survivors (10)	Faded (38)
FH/TF	1d/1d (4), 4h/1h (4), 1h/1h (1), 4h/4h,1h (1)	4h/1h (17), 1d/1d (10), 4h/4h (4), 1d/1d,4h (3), 1d/4h (2), 4h/4h,1h (2)
Council size	4 models (4), 3 models (3), solo (2), 2 models (1)	solo (13), 4 models (10), 3 models (9), 5 models (4), 2 models (2)
Universe size	2 coins (9), 5 coins (1)	2 coins (29), 3 coins (3), 1 coin (2), 4 coins (2), 5 coins (2)
Calibrate ON	4 of 10 (40.0%)	5 of 38 (13.2%)
Min confidence	60 (9), 0 (1)	60 (32), 0 (2), 50 (2), 55 (2)
Table stake	\$450 (8), \$200 (1), \$375 (1)	\$450 (33), \$300 (3), \$200 (1), \$250 (1)
Median maxDD	19.1%	35.2%
Median win rate	68.0%	61.5%
Median smoothness R2	0.359	0.825
Model membership	gemini (9), gpt (5), deepseek (4), opus (4), qwen (4), fable (3)	gemini (22), deepseek (20), opus (20), fable (14), gpt (13), qwen (8), grok (7)

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 #8 fresh-week (Sep 21-28, 2026) replay -- one of the two windows the verdict needed, and the R2 median is taken over the rows that have one (10 of 10 survivors, 37 of 38 faded: a replay with too few days carries no fit). On Config Watch #8's four-week window (Aug 31-Sep 28, 2026) the same medians read maxDD 39.7% against 58.0% and win rate 66.4% against 52.0%. On this issue's calendar month the same two groups read 7 of 10 and 8 of 38 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)	1,446	65.6%	+\$218.4	+\$1,263.5	-\$1,024.2
-- mature rows only	719	75.2%	+\$197.8	+\$1,104.4	-\$1,024.2
Board-leader slice (top of board)	14	14.3%	-\$413.3	+\$200.6	-\$670.1

On this window the axis is positive in the population and negative at the top. Calibration improves 65.6% of the 1,446 matched pairs on the whole grid (mean +\$218.4, median +\$162.3) and 75.2% of the 719 mature pairs (mean +\$197.8). On the board-leader slice it improves 14.3% of 14 pairs (mean -\$413.3, median -\$508.0): the OFF legs there are at a median of +\$918.3 and the ON legs at +\$455.7, 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 65.1% of its pairs, on Config Watch #8's fresh week 36.5% and on calendar August (Monthly Config Watch #1) 37.9%. Where the axis reached the cards is visible above: no council card is a calibrated cell while solo card #1 is, and 31 of the 49 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 +\$766.5 against the grid's +\$162.3), 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 (Sep 1-30, 2026)

#	Model	Coins	FH/TF	Net	Net %	maxDD	Trades	WR
#1	gemini	SOL+XRP	4h/1h	+\$923.2	+92.3%	33.0%	204 *	85.8%
#2	grok	BNB+XRP	4h/4h,1h	+\$878.5	+87.9%	29.9%	107	76.6%
#	Ladder	BE	Shifts	Conf filter	Source	Entry	Positional	Stake
#1	50:40,100:60	on	TP -10 / SL +75	≥ 60 (each)	median	0/0	reenter	\$450
#2	50:40,100:60	on	TP -10 / SL +75	≥ 60 (each)	median	1/0	reenter	\$450

Two solo cards, one per model (gemini and grok); only #1 carries a short link. No single-coin row outranks it on that board this window. The card carries 3 unfunded entries (flagged * on Trades); the best mature solo row with no unfunded entry is solo card #2 (grok on BNB+XRP, +\$878.5 on 107 closed).

Council vs solo. On this window the best solo sits **below** the best council (+\$923.2 against +\$1,041.5 on the mature boards), on 204 closed trades against 97, with a 33.0% drawdown against 37.7% and 3 unfunded entries against 0. Config Watch #8 measured the same comparison on its two windows: the council above the solo on the week, the council above the solo on the four weeks. Same in-sample caveats apply to both sides.

OPEN IN SANDBOX marketmania.ai/s/cwm2-s1

Each link opens this exact frozen window; results visible without sign-in (embedded share signature). The codes are minted by the publish script before this file goes live; 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** = $\text{net} / \max(\text{maxDD}, 1.0)$, 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. R2 and the ulcer index are computed by the collect on the engine's daily settled-equity series: one point per day that settled a trade, from the first such day (the \$1,000 start is not a point), index-spaced, the ulcer index over the running peak of those points -- so days without a settlement are compressed and a loss on the first settled day does not enter the ulcer index. 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	Cls	R2
MONTH Calmar	fable, opus, qwen [cal]	SOL+XRP	1d/1d	+\$582.4	93.3%	0.0%	30	0.967
MONTH Calmar	fable, opus, qwen [cal]	SOL+XRP	1d/1d	+\$582.4	93.3%	0.0%	30	0.967
MONTH Calmar	fable, qwen [cal]	SOL+XRP	1d/1d	+\$578.3	93.9%	0.0%	33	0.977
MONTH WR	fable, gpt, qwen, grok [cal]	BTC+ETH	4h/4h,1h	+\$289.7	95.0%	10.8%	40	0.792
MONTH WR	fable, gpt, qwen, grok [cal]	BTC+ETH	4h/4h,1h	+\$241.4	95.0%	9.0%	40	0.792
MONTH WR	fable, gpt, qwen, grok [cal]	BTC+ETH	4h/4h,1h	+\$193.1	95.0%	7.2%	40	0.792
MONTH Smoothness	fable, qwen [cal]	SOL+XRP	1d/1d	+\$578.3	93.9%	0.0%	33	0.977
MONTH Smoothness	fable, qwen [cal]	SOL+XRP	1d/1d	+\$578.3	93.9%	0.0%	33	0.977
MONTH Smoothness	fable, qwen [cal]	SOL+XRP	1d/1d	+\$385.5	93.9%	0.0%	33	0.977
MONTH Quality	fable, qwen [cal]	SOL+XRP	1d/1d	+\$578.3	93.9%	0.0%	33	0.977
MONTH Quality	fable, qwen [cal]	SOL+XRP	1d/1d	+\$578.3	93.9%	0.0%	33	0.977
MONTH Quality	fable, opus, qwen [cal]	SOL+XRP	1d/1d	+\$582.4	93.3%	0.0%	30	0.967

[cal] = calibrate=ON cell; the board column names the window. The Calmar board is degenerate by construction -- its top rows sit at 0.0% maxDD, so the score collapses onto net PnL; read it as 'net among configs that barely drew down'. None of the 12 printed rows is a single-coin config, and all 12 are calibrated. The pool is 2,967 mature configs against 2,911 on the four-week window of the same collect and 3,265 on calendar August.

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 31-Sep 28, 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; they share 27 days (Sep 1-27), Aug 31 lies in the four-week window only and Sep 28-30 in the month only, so neither contains the other and this is a persistence check across two largely overlapping windows, not a cross-period one. 22 configs qualify (under the collect's row cap of 25), on 2 universes (BNB+XRP 12 and SOL+XRP 10), 18 of them councils and 5 clean (zero unfunded entries on both windows). The top three by rank sum are council fable, opus, gemini, qwen (+\$1,021.7); council fable, opus, gemini, qwen (+\$1,041.5); council fable, opus, gemini, qwen (+\$970.2). Monthly Config Watch #1 read the same pair of windows on calendar August from its own collect and held 25, the collect's row cap -- the comparable board; Config Watch #8's board (4 configurations) reads across WK and M30 and is not comparable. 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	-\$600.5	33.3%	60.0%	0.610	35.04	6

679 of the 2,967 mature configs (22.9%) clear all three bars at once. The anchor's own month is -\$600.5 net at 33.3% win rate with a 60.0% drawdown, an R2 of 0.610 and 6 closed trades. Of the ten best of those 679 by net, all ten beat the anchor on net PnL as well. 2 of the ten are calibrated cells, 8 are councils, all ten run the \$450 stake and their universes are 2 coins. On the four-week window of the same collect the same rule admitted 1,083 of 2,911 configs (37.2%), and on calendar August (Monthly Config Watch #1) 12 of 3,265 (0.4%), 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 679 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, gemini, qwen	BNB+XRP	4h/1h	\$450	+\$1,041.5	71.1%	37.7%	0.797
#2	fable, opus, gemini, qwen	SOL+XRP	4h/1h	\$450	+\$1,021.7	72.2%	47.7%	0.821
#3	fable, opus, gemini, qwen	SOL+XRP	4h/1h	\$450	+\$970.2	71.1%	52.6%	0.760
#4	fable, gpt, gemini, qwen	SOL+XRP	4h/1h	\$450	+\$960.8	69.8%	37.2%	0.763
#5	fable, opus, gemini, qwen	SOL+XRP	4h/1h	\$450	+\$937.9	70.8%	52.6%	0.757
#6	fable, opus, gpt, grok	BNB+XRP	4h/1h	\$450	+\$937.4	78.6%	18.7%	0.823
#7	gemini [cal]	SOL+XRP	4h/1h	\$450	+\$923.2	85.8%	33.0%	0.757
#8	gemini [cal]	SOL+XRP	4h/1h	\$450	+\$923.2	85.8%	33.0%	0.757
#9	fable, opus, gpt, qwen	BNB+XRP	4h/1h	\$450	+\$923.0	79.5%	28.1%	0.822
#10	fable, gpt, gemini, qwen	SOL+XRP	4h/1h	\$450	+\$913.6	69.3%	37.2%	0.755

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 1,150 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	gpt [solo]	4h/4h,1h	+\$187.1 (28)
ETH	deepseek [solo]	1d/1d,4h	+\$376.0 (18)
SOL	fable, opus, gpt, gemini [council]	4h/1h	+\$720.5 (94)
BNB	grok [solo]	4h/4h,1h	+\$371.1 (54)
XRP	fable, opus, gemini, qwen [council]	4h/1h	+\$828.1 (50)

n in parentheses; * = fewer than 10 closed trades (thin -- direction only). All five winners run calibrate OFF and 2 of the five are councils.

The reads. Every coin is net-positive over the month and no row is thin: 18 to 94 closed trades, XRP leading at +\$828.1 and BTC lowest at +\$187.1. 1 of the five sits below the ≥ 20 -trade maturity gate the council boards apply (ETH 18), so that row is direction with a small sample rather than a measured edge. The 1d horizon wins 1 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 on 5 coins for the 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 (grok-4.6)	-\$91.4 *	797	61%	9.2%
Worst solo default (deepseek-v4-pro)	-\$194.4 *	1260	59%	19.6%
Council-of-7 default	-\$69.8 *	570	61%	7.1%
Hand-tuned config A (pre-search)	-\$815.9 *	9	22%	81.6%
Hand-tuned config B (pre-search)	-\$195.8 *	181	63%	71.8%
Hand-tuned config C (pre-search)	-\$19.8 *	46	65%	58.7%

* = 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 (486 to 1,384 skipped each over the month). 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 C moves \$400 -> \$450 on 2 coins and reads -\$40.1 at its own stake against -\$19.8 at the table stake.

Over the month the standing ordering holds in one half only. The seven solo defaults run -\$91.4 to -\$194.4 at table stake and the council-of-7 -\$69.8 -- every default lost money at table stake; the three pre-search hand-tuned configs read -\$815.9 / -\$195.8 / -\$19.8, and the fresh search's +\$1,041.5 is above the best reference point in the table (-\$19.8). Search > hand-tuned > defaults is only partly what this month measures: the search maximum is above every reference point; a default is at or above a hand-tuned config. The live presets are a separate row of evidence: none of the four is 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); sol and xrp 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	-\$600.5	+\$555.3 (1d/1d, SOL+XRP, 18 cls)	34	neighbour +\$1,155.8
preset 4h4	-\$20.4	+\$273.9 (4h/1h, SOL+XRP, 44 cls)	19	neighbour +\$294.3
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 on this window: bc2 by +\$1,155.8 (-\$600.5 -> +\$555.3; tp_sl_source nearest -> median); 4h4 by +\$294.3 (-\$20.4 -> +\$273.9; tokens ETH,SOL -> SOL,XRP). 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; 49 unique configs, 39 of them councils and 10 solo).

Knob	MONTH finalists mode
Forecast horizon	4h (33/49)
Ladder shape (steps)	50:40,100:60 (49/49)
Break-even stop (BE)	on (49/49)
Min confidence	60 (49/49)
SL shift	75 (49/49)
TP shift	-10 (49/49)
TP/SL source	farthest (42/49)
Calibrate ON	31/49
Universe	SOL+XRP (32/49)

The mechanical core: 50:40,100:60 ladder (49/49), break-even on (49/49), conf 60 (49/49), SL +75 / TP -10 (49/49), farthest source (42/49) -- the same core Monthly Config Watch #1 printed for calendar August. The horizon mode is **4h** (33/49; calendar August: 1d (46/62)). **Calibrate=ON** runs in 31 of the 49 finalists, against 37 of 55 on the four-week window of this collect and 37 of 62 on calendar August. The universe is SOL+XRP in 32 of the 49 finalists, and none of them is a single-coin row. Membership across the 39 council finalists: fable 39, qwen 32, opus 24, gpt 18, gemini 11, grok 10. Standing caveats: refinement seeds from the same grid winners, so part of the convergence is search-design echo, and a knob core that produced 10-of-48 survival on Config Watch #8's two windows and 17-of-53 net-positive on this month is describing 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 37.7% (council #1), 47.7% (council #2) and 33.0% (the solo card), and the replay table is deeper: 45 of the 57 replayed rows drew down more than 40% over the month and 32 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 (standing). 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 (Thu Oct 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; this issue's collect is a run of its own that sweeps the four-week window beside the month, and it 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 (standing). 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 to #4 through Monthly Config Watch #1's table and Config Watch #5 to #8 directly) and the structure table is computed from the frozen queries and the replay statistics of the cards Config Watch #8 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 30-day window (Sep 1-30, 2026) on a single regime; the month-over-month figures set it beside one earlier month (calendar August, Monthly Config Watch #1), and no past issue is restated. Trades settle past the window edge at their horizon; open = a position still open when the collect ran (Oct 2, 17:14 UTC, 41 hours after the window edge): 0 on the three published cards, 8 of the 114 replay cells (57 rows x 2 windows) carry one each -- listed in the open data.
- **This month overlaps the weekly windows, so it is not independent of them.** 27 of its 30 days lie in the four published Config Watch weeks it is read against (Config Watch #5 Aug 31-Sep 7, 2026; Config Watch #6 Sep 7-14, 2026; Config Watch #7 Sep 14-21, 2026; Config Watch #8 Sep 21-28, 2026) and 27 in Config Watch #8's 28-day four-week window; the month, the weeks and the four weeks are readings of overlapping data, and agreement between them is partly arithmetic.
- **The replay column is decay context, not clean out-of-sample.** The search windows of issue-8, issue-7, issue-6 and issue-5 overlap this month, so 19 of the 53 replayed cards are partly scored on the days they were selected on; the older cohorts' windows end before it. Publication dates bound it too: 27 of the 53 cards were published inside or after the month (issue-8 on Oct 1, issue-7 on Sep 23, issue-6 on Sep 16, issue-5 on Sep 9 and issue-4 and Monthly-1 on Sep 4), so for them the September days before publication are not a forward test of a published card; only the 26 cards of issue-3, issue-2 and issue-1 were published before Sep 1 (issue-3 on Aug 26, issue-2 on Aug 19 and issue-1 on Aug 13) -- the 26 cards Monthly Config Watch #1 read at 20 of 26 on calendar August, and they read 2 of 26 here. The 53 cards are 51 distinct configurations window aside (cwm1-s1 = cw3-s2; cw3-m301 = cw3-wk1). 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 and September, the Calendar Sep column is a 30-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 that share 27 of their days -- so it measures persistence across largely overlapping windows; it is comparable with Monthly Config Watch #1's board, read across the same pair, and not with the weekly issues' boards built across WK and M30.
- **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 59,755 resolved forecasts at the monthly cutoff (Thu Oct 1, 2026, 16:00 UTC), where Config Watch #8 printed 57,929 at its own cutoff (Sep 28, 2026 16:00 UTC) and Monthly Config Watch #1 39,073 at its own (Mon Sep 1, 2026, 16:00 UTC): the same counter at three cutoffs, not three measurements. The model line counts 7 current models with earlier versions folded into their successors' lineage.
- **Multiple testing, winner's curse and thin cells.** 11,075 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 4 replay rows are under 10 closed trades and 1 of the five per-ticker winners is 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

Month collect (the four-week and calendar-month windows together; the four-week boards are not printed)

Stage	Sims	Stage	Sims
Scan (s1)	952	Dedicated per-ticker (s6t+s6tc)	1,150
Systematic grid (s2)	8,352	OOS replays (s4cw8+s4cw8o+s4cw7+s4cw7o+s4cw6+s4cw6o+s4cw5+s4cw5o+s4cw4+s4cw4o+s4cwm1+s4cwm1o+s4cw3+s4cw3o+s4cw2+s4cw2o+s4cw1+s4pre)	155
Refinement: ladder + stake sweep (s2b+s3b)	240	Total	11,075
Preset neighbourhood grid (s2p)	106	Errors / zero-trade	0 / 1,686
Per-ticker probes (s5t)	120	Engine	1.1

Scale note. This issue's collect is a run of its own: 11,075 simulations over its two windows, against 10,889 in Monthly Config Watch #1's collect. Config Watch #8 ran two separate collects (5,601 weekly + 11,152 night = 16,753), which this issue does not use. The per-window stage-2 grid is 4,176 cells here and 4,208 on Config Watch #8's fresh week. The replay branch is 57 rows -- 53 past cards and 4 live presets -- for 155 simulations, because every row is replayed on both windows of this collect and the cards of eight cohorts are re-run on their origin windows. The selection boards and the survival, structure and 30-day tables 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 are printed here and the other 22 (its four-week window) stay in the open data. 0 equity replays were needed.

Source files: summary_month.json (the w298 collect) with grid_hist_month, equity_month, baselines_month, presets_recon_month and the nine oos_*_month packs beside it; the published open data of Monthly Config Watch #1 and Config Watch #4 to #8 for every number this issue quotes from them. Generated at: 2026-10-02T17:14:05Z · 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: 11,075 simulations (0 errors) in one frozen collect, of which 155 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 Oct 1, 16:00 UTC): 59,755 resolved since Jul 11 · 7 models tracked (current line-up; earlier versions folded into their successors' lineage) · 5 assets · 5 horizons · hourly · 42 published reports

Counter pinned from wave 3 onward and read at this issue's cutoff (Thu Oct 1, 2026, 16:00 UTC), where Config Watch #8 read 57,929 at its own cutoff and Monthly Config Watch #1 39,073 at its own; the model line counts the current line-up, with earlier versions folded into their successors' lineage. Monthly Config Watch #2 is the 42nd published report: Monthly Consensus Watch #2, Monthly Calibration #2 and Monthly Model Watch #2 are 39, 40 and 41, and Config Watch #8 is 38.

Issue #2. 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 issue appends calendar September -- a best mature council of +\$1,041.5 (August: +\$1,618.4) and a best mature solo of +\$923.2 (August: +\$1,721.9), 17 of 53 published cards net-positive on the month (August: 20 of 26), a survival curve that reaches +7 fresh weeks for the issue-1 generation (August: +3), and a structural table of Config Watch #8's 10 survivors against its 38 faded cards. The calendar week and the four-week window stay in the weekly Config Watch.

What we're testing next

- **Monthly Config Watch #3** -- the October calendar month (2026-10-01 to 2026-11-01) from the same engine, the same criteria and the same collect code, read against the weekly issues whose weeks lie in that month. It carries the first calendar-month replay rows for this issue'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 #3 prints those columns for the October weeks, which is the first point at which the issue-8 generation has a fresh-week reading.
- **The quarterly** -- the first cross-issue aggregate of the walk-forward table: the weekly and monthly issues of the quarter on one axis, with the survival curve as its spine.

Related research

Monthly Consensus Watch #2 -- <https://marketmania.ai/research/reports/consensus-watch-monthly-2026-09.pdf>

Monthly Calibration #2 -- <https://marketmania.ai/research/reports/calibration-monthly-2026-09.pdf>

Monthly Model Watch #2 -- <https://marketmania.ai/research/reports/model-watch-monthly-2026-09.pdf>

Config Watch #8 -- <https://marketmania.ai/research/reports/config-watch-2026-09-23.pdf>

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

The three monthly reports publish together as one issue each month; this issue joins them on the same cadence. Config Watch #8 is the newest weekly issue: it carries the calendar week Sep 21-28, 2026 and the four-week window Aug 31-Sep 28, 2026 from collects of its own, and no board of either is reprinted here. Monthly Config Watch #1 is the previous issue of this series and the source of every month-over-month number.

CITE THIS REPORT

```
@techreport{mm_configwatch_monthly_2026m09,
  title      = {Monthly Config Watch #2},
  author     = {{MarketMania Research}},
  institution = {MarketMania},
  year       = {2026},
  month      = oct,
  day        = {3},
  type       = {Monthly Research Report},
  series     = {Monthly Config Watch},
  number     = {2},
  note       = {Methodology v1.1 + Config Watch amendments, hash e66c7e8c864a2233},
  url        = {https://marketmania.ai/research/reports/config-watch-monthly-2026-09.pdf}
}
```