Monthly Config Watch #3

November 5, 2026 · MarketMania Research · Monthly series

MONTHLY · CONFIG WATCH · Issue #3 · one calendar month, 83 replays, one survival curve · Issue date November 5, 2026 · Language English · Methodology v1.1 + Config Watch amendments (approved 12.08, hash e66c7e8c864a2233)
Search window: **MONTH** Oct 1-31, 2026 -- the frozen calendar month (Thu Oct 1 00:00 -> Sat Oct 31 23:59 UTC, 31 days), cut-off edge Sun Nov 1 00:00 UTC, so slots dated Nov 1 fall outside it; report cut-off Sun Nov 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 #2 (October 3, 2026, calendar September, Sep 1-30, 2026) -- best mature council +$1,041.5, best mature solo +$923.2, 17 of 53 published cards net-positive on that month; every month-over-month number below is read from its open data. The weekly Config Watch #13 carries the calendar week Oct 26-Nov 2, 2026 and the four-week window Oct 5-Nov 2, 2026 from two collects of its own, and Config Watch #12 the four-week window Sep 28-Oct 26, 2026 that this issue's collect recomputes beside the month; no weekly board is reprinted here.

RESEARCH SNAPSHOT			
10,659 simulations, 0 errors, in one frozen collect	**MONTH** Oct 1-31, 2026	cut-off edge **Sun Nov 1 00:00 UTC**; report cut-off Sun Nov 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 (w373 plan)
walk-forward: **79** past cards + **4** live presets replayed verbatim on the month	**21 of 79** past cards net-positive on the month	stage-2 grid: **3,952** cells on this window	models: **7** (3 lineage axes; cutover inside both windows)
best MONTH council **+$1,119.4** (+111.9%)	best MONTH solo **+$1,311.1** (+131.1%)	zero-trade cells: **1,554** (whole collect)	margin-skip alerts on leaders: **12** rows (MONTH)

KEY FINDING · OBSERVATION (one calendar month)

The best mature council on calendar October makes +$1,119.4 (+111.9% of a $1,000 deposit, 88.5% win rate, 15.5% maxDD, 26 closed), 21 of the 79 cards this series published through Config Watch #12 are net-positive on the same month, and all 7 cards that survived both fresh windows of Config Watch #13 are among them.

WEEKLY EVIDENCE -> MONTHLY VERDICT (rule below)

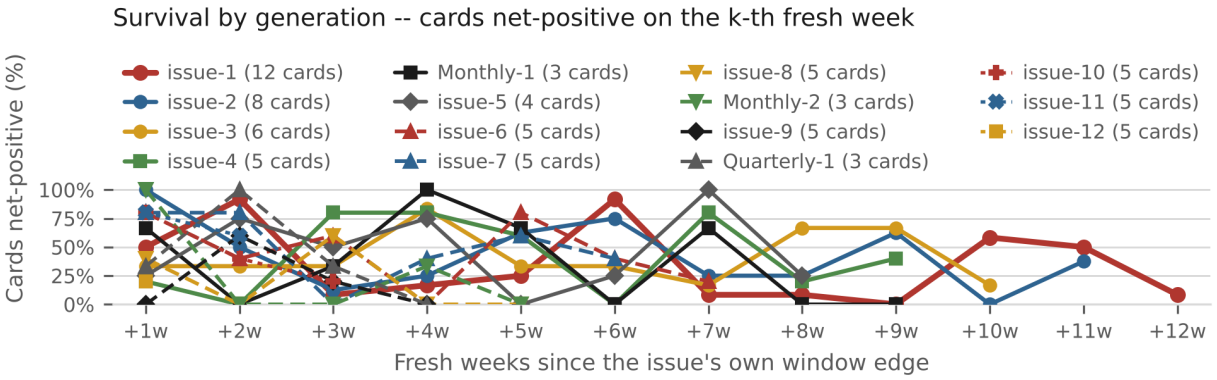
Metric	CW#10 wk	CW#11 wk	CW#12 wk	CW#13 wk	Month	Verdict
Published cards net-positive (majority)	4/64	42/69	29/74	16/79	21/79	Confirmed (3/4)
Best council universe	SOL+XRP	SOL+XRP	SOL+XRP	BTC+ETH+ SOL+XRP	SOL+XRP	Confirmed (3/4)
Calibration helps (matched pairs)	55.7%	27.9%	29.5%	56.8%	34.9%	Inconclusive (2/4)

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-October pull.

Sources, cell by cell. Row 1 counts the replays each issue ran on its own fresh week -- Config Watch #10 4 of 64 (Oct 5-12, 2026), #11 42 of 69 (Oct 12-19, 2026), #12 29 of 74 (Oct 19-26, 2026), #13 16 of 79 (Oct 26-Nov 2, 2026) -- against 21 of 79 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: 1,361 pairs in issue #10, 1,246 in #11, 1,336 in #12 and 1,343 in #13, against 1,338 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 #13's week runs to Nov 2, past the month's right edge, so 6 of its 7 days are October days. The four weeks cover 27 of the month's 31 days; Oct 1-4 lie before Config Watch #10's week.

TL;DR

- OBSERVATION -- **21 of the 79 published cards are net-positive on the calendar month** (calendar September: 17 of 53). Config Watch #13 read all 79 of them on its fresh week (Oct 26-Nov 2, 2026): 16 positive; on its four-week window (Oct 5-Nov 2, 2026): 17. By generation on the month: issue-12 3/5, issue-11 5/5, issue-10 4/5, Quarterly-1 1/3, issue-9 0/5, Monthly-2 0/3, issue-8 0/5, issue-7 0/5, issue-6 1/5, issue-5 1/4, issue-4 1/5, Monthly-1 0/3, issue-3 2/6, issue-2 2/8, issue-1 1/12; live presets 1/4. On the 48 cards published before the month began (issue-7, issue-6, issue-5, issue-4, Monthly-1, issue-3, issue-2 and issue-1) the reading is 8 of 48; the 53 cards Monthly Config Watch #2 read on calendar September (17 of 53 there) read 8 of 53 here. The month holds 6 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 +$1,119.4 on the council board and +$1,311.1 on the solo board** (calendar September: +$1,041.5 and +$923.2). The council card runs at +111.9% of a $1,000 deposit, 88.5% win rate and 15.5% maxDD on 26 closed trades; the solo card at +131.1%, 90.0% win rate and 24.5% maxDD on 20 closed with no unfunded entry. The grid behind them: median $0.0 and 41.9% positive across 3,952 stage-2 cells, 233 of them above +$500, range -$757.8 to +$1,311.1. On the four-week window of the same collect that grid ran 40.8% positive; the month leader has no four-week cell in this collect, and Config Watch #12's four-week leader reads +$959.2 on the month against +$880.9 on the four weeks (section below).
- **The survival curve reaches +12 fresh weeks for the issue-1 generation** (+7 in Monthly Config Watch #2). Share of a generation's cards net-positive on its latest fresh week: issue-1 8.3% (+12); issue-2 37.5% (+11); issue-3 16.7% (+10); issue-4 40.0% (+9); Monthly-1 0.0% (+9); issue-5 25.0% (+8); issue-6 20.0% (+7); issue-7 40.0% (+6); issue-8 0.0% (+5); Monthly-2 0.0% (+5); issue-9 0.0% (+4); Quarterly-1 0.0% (+4); issue-10 20.0% (+3); issue-11 60.0% (+2); issue-12 20.0% (+1); every generation has at least one fresh week. Every cell is read out of the issue that ran that replay, and the points sit on 12 different calendar weeks, from Aug 10-17, 2026 to Oct 26-Nov 2, 2026.
- **The 7 survivors and the 72 faded cards of Config Watch #13, side by side.** The cards that cleared both of its fresh windows against those that did not: median fresh-week maxDD 0.0% vs 29.2%, median win rate 88.9% vs 57.1%, 6 of 7 against 20 of 72 on the 1d horizon, and 0 of 7 (0.0%) vs 14 of 72 (19.4%) calibrated. 7 rows against 72 rows on two windows 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 (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, issue-8 5 cards, Monthly-2 3 cards, issue-9 5 cards, Quarterly-1 3 cards, issue-10 5 cards, issue-11 5 cards, issue-12 5 cards). Each point is one published walk-forward table: issue-1's twelve points are Config Watch #2 to #13, issue-12's one point is #13; the count behind every point is in the survival tables below. The x axis counts fresh weeks, not calendar weeks -- the points sit on 12 different weeks, from Aug 10-17, 2026 to Oct 26-Nov 2, 2026 -- 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 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 15 generations, and the structural comparison of the 7 cards that cleared both fresh windows of Config Watch #13 against the 72 that did not. The month is not an independent test of the weeks: 27 of its 31 days lie in the four published Config Watch weeks it is read against and 25 in Config Watch #12'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 (w373 plan) and printed into the meta of the JSON deliverable: the month run, which sweeps the four-week window of Config Watch #12 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 3,952 cells on each of its two windows.

- **1. Scan (s1)** -- 952 sims across council compositions and coarse settings.
- **2. Systematic grid (s2)** -- 7,904 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)** -- 254 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)** -- 1,090 sims: a compacted independent per-coin grid plus calibrate twins (the standing branch from Config Watch #2).
- **7. OOS replays (s4cw12+s4cw12o+s4cw11+s4cw11o+s4cw10+s4cw10o+s4cwq1+s4cwq1o+s4cw9+s4cw9o+s4cwm2+s4cwm2o+s4cw8+s4cw8o+s4cw7+s4cw7o+s4cw6+s4cw6o+s4cw5+s4cw5o+s4cw4+s4cw4o+s4cwm1+s4cwm1o+s4cw3+s4cw3o+s4cw2+s4cw2o+s4cw1+s4pre)** -- 233 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 #12, issue #11, issue #10, Quarterly #1, issue #9, Monthly #2, issue #8, issue #7, issue #6, issue #5, issue #4, Monthly #1, issue #3 and issue #2 for the drift baseline.

Total 10,659 simulations, 0 errors, 1,554 zero-trade cells; stage sums reconcile exactly inside the collect. The stage-2 grid is 3,952 cells on this window against 4,176 on Monthly Config Watch #2'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: 12, of the 19 the collect raised over its two windows; the other 7 are on its four-week board, which is not printed here.

Models under test (engine IDs as displayed): claude-fable-5, claude-opus-5.5, gpt-6.1-sol, deepseek-v4-pro, gemini-3.1-pro, qwen-3.8-max, grok-4.7. Three axes run as lineages -- **claude-opus-5.5 (incl. opus-5-era), gpt-6.1-sol (incl. 5.6-era), grok-4.7 (incl. 4.6-era)** -- across the model cutover at Oct 17, 2026 09:01 UTC, which lies inside both windows of this collect: day 17 of the

month's 31 and day 20 of the four weeks' 28. Published configs that named claude-opus-5, gpt-5.6-sol, grok-4.6 or grok-4.5 are replayed out-of-sample with the same lineage remap ([remap] rows); the five issue-12 and five issue-11 cards already name the heads. The grok-4.5 -> 4.6 cutover dated Aug 24 in the search plan lies entirely before both windows. The vendors' names are Claude Opus 5.5, GPT-6.1 Sol and Grok 4.7.

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

#	Council	Coins	FH/TF	Net	Net %	maxDD	Trades	WR
#1	gpt, grok	SOL+XRP	1d/1d	**+$1,119.4**	+111.9%	15.5%	26	88.5%
#2	gpt, deepseek	BNB+XRP	1d/1d,4h	**+$613.4**	+61.3%	15.6%	29	93.1%

#	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	3/0	reenter	$450

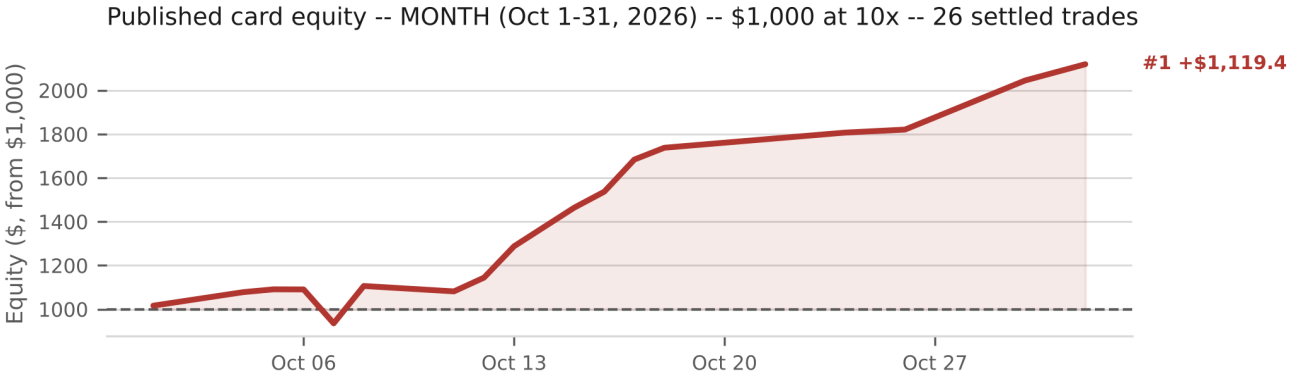
Against the 3,952-cell stage-2 grid on this window: card #1 is the same simulation as the raw net-board leader (i=5862), so its exported percentile is the card's -- above **99.8%** of the grid; card #2 lands in the +$500 to +$750 bucket (180 cells) with 94.1% of the grid finished below it. Grid median $0.0, 233 of the 3,952 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 **WR** board (+$613.4 on BNB+XRP, 15.6% maxDD, R2 0.804, no unfunded entry). Read the cards as two universes deep, not three.

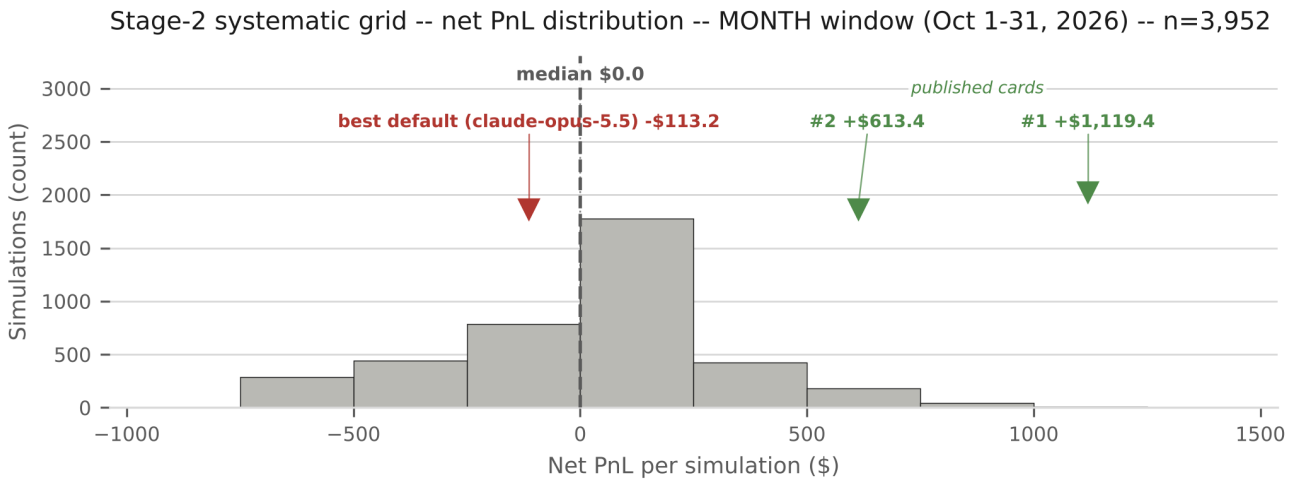
The two cards answer the calibrate axis the same way: both cards are calibrated cells (Calibrate axis below). Card #1 runs the $450 two-coin table stake and card #2 the $450 two-coin stake; neither card carries an unfunded entry, and card #1 draws down 15.5% on 26 closed trades against card #2's 15.6% on 29. The mature pool on this window is 2,950 configs against 2,827 on the four-week window of the same collect (31 days against 28) and 2,967 on calendar September.

OPEN IN SANDBOX **marketmania.ai/s/cwm3-m1** **marketmania.ai/s/cwm3-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.



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=5862), so a second curve would be the same line. The line peaks at +$1,119.4 on 2026-11-01 and ends the window at +$1,119.4. Card #2, the solo card and the other board leaders keep their daily series in the open data -- the PDF carries one equity chart.



Net PnL across all 3,952 stage-2 grid sims, MONTH window (Oct 1-31, 2026), universe-aware stake throughout. Grid median **$0.0**, 41.9% positive (calendar September: -$5.7, 29.6% positive); the same grid on the four-week window of this collect ran 40.8% positive and on Config Watch #13's fresh week 22.8%. The $0 bin holds 1,778 cells; 778 sims on this window produced no trade at all. Range -$757.8 to +$1,311.1; 233 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 #12 published two councils on its four-week M30 window (Sep 28-Oct 26, 2026, 28 days, edge Mon Oct 26 00:00 UTC); this issue publishes two councils and one solo on the calendar month (Oct 1-31, 2026, edge Sun Nov 1 00:00 UTC). This collect replays Config Watch #12's cards on both of its windows, and the two windows share 25 days -- Sep 28-30 lies in M30 only and Oct 26-31 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
cw12-m301	opus, gpt	SOL+XRP	1d/1d	+$880.9	+$959.2	73.5% / 75.0%	36.8% / 36.8%	34 / 32
cw12-m302	gpt, deepseek	BNB+XRP	1d/1d,4h	+$821.0	+$751.1	88.0% / 85.7%	21.4% / 21.4%	25 / 28
cwm3-m1	gpt, grok	SOL+XRP	1d/1d	n/a	+$1,119.4	n/a / 88.5%	n/a / 15.5%	n/a / 26
cwm3-m2	gpt, deepseek	BNB+XRP	1d/1d,4h	n/a	+$613.4	n/a / 93.1%	n/a / 15.6%	n/a / 29
cwm3-s1	grok	SOL+XRP	1d/4h	n/a	+$1,311.1	n/a / 90.0%	n/a / 24.5%	n/a / 20

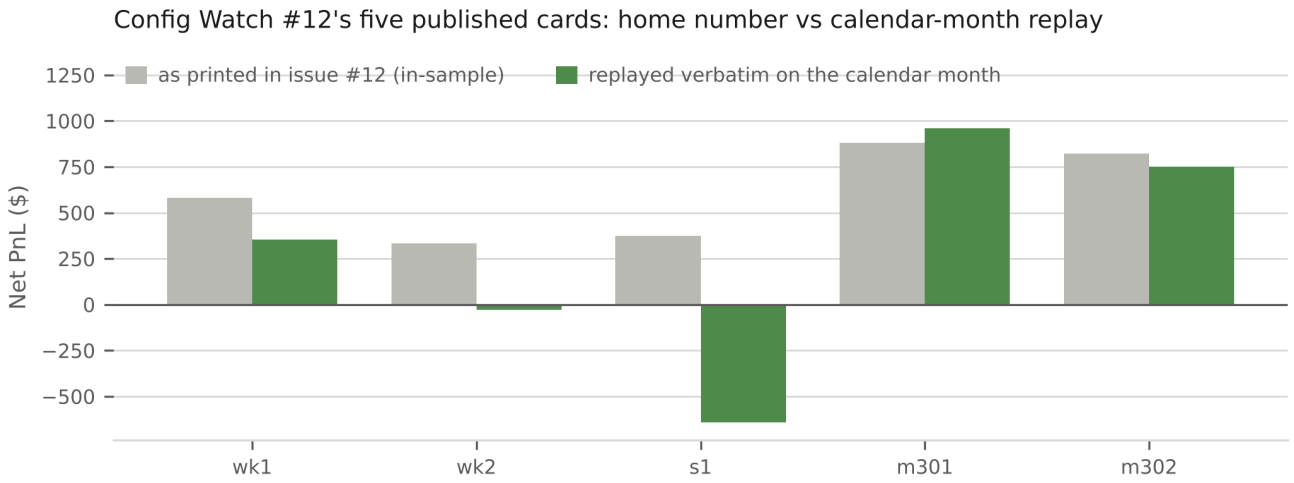
Every paired column reads **M30 first, calendar month second**; Cls = closed trades. The five rows are five distinct configurations. Config Watch #12's rows are this collect's replay of the published query on both windows; Config Watch #12 printed cw12-m301 +$880.9 and cw12-m302 +$821.0 for them from its own night collect, and this collect's replay reproduces each published figure on the same four weeks, where the replay is in-sample (the window the cards were selected on). No row reaches the mature net top-40 of both windows, so no row carries a rank pair. cwm3-m1 has no four-week cell in this collect (no board, replay or exported pair carries it), so it reads n/a; cwm3-m2 has no four-week cell in this collect (no board, replay or exported pair carries it), so it reads n/a; cwm3-s1 has no four-week cell in this collect (no board, replay or exported pair carries it), so it reads n/a. Boards, the replay set and a pair sample are what the collect exports, not every cell on both windows.

Two of the five rows can be read on both windows, and they differ by $69.9 to $78.3 between them. cw12-m301 makes +$880.9 on M30 and +$959.2 on the month (+$78.3, on 32 closed trades against 34); cw12-m302 makes +$821.0 on M30 and +$751.1 on the month (-$69.9, on 28 closed trades against 25). 1 of them is larger on the window their own issue published it on.

Walk-forward on the calendar month

Every config this series has ever published through Config Watch #12 -- all 79 cards from Config Watch #1 to #12, Monthly Config Watch #1, Monthly Config Watch #2 and Quarterly 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 233 simulations inside this collect, which replays every row on both of its windows in one pass. Config Watch #12's own five cards are in a calendar-month table for the first time; their search windows share 25 of 28 and 7 of 7 days with this month (the rule below counts every cohort). Config Watch #13's five cards (published November 4, 2026) are not in this collect: it replays the cards published through Config Watch #12.

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 #13 applied. The column is not pure out-of-sample either: the month shares days with the search windows of issue-12, issue-11, issue-10 and issue-9 (days inside the month: issue-12's 28-day window 25 and 7-day window 7; issue-11's 28-day window 18 and 7-day window 7; issue-10's 28-day window 11 and 7-day window 7; issue-9's 28-day window 4 and 7-day window 4), so for those 20 rows the number a card shows here is partly scored on the days it was selected on; the search windows of Quarterly-1, Monthly-2, issue-8, issue-7, issue-6, issue-5, issue-4, Monthly-1, issue-3, issue-2 and issue-1 share no day with it. Read the column as decay context over a 31-day window.



Config Watch #12'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). 3 of the five are net-positive on the month. The month holds 25 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 [remap]	ETH+SOL	-- live	-$383.4 (32) +	Faded
preset bc2	gemini, gpt, deepseek [remap]	SOL+XRP	-- live	+$73.9 (23) +	**Survived**
preset sol	deepseek, fable, opus [remap]	SOL [conc]	-- live	-$937.9 (39)	Faded
preset xrp	gemini, gpt, deepseek [remap]	XRP [conc]	-- live	-$404.3 (4) * +	Faded
cw12-m301	opus, gpt	SOL+XRP	+$880.9	+$959.2 (32)	**Survived**
cw12-m302	gpt, deepseek	BNB+XRP	+$821.0	+$751.1 (28)	**Survived**
cw12-s1	gemini	SOL+XRP	+$373.1	-$643.4 (169) +	Faded
cw12-wk1	fable, gemini	SOL+XRP	+$580.5	+$353.7 (168) +	**Survived**
cw12-wk2	gpt, gemini, qwen	BTC+ETH	+$333.4	-$29.6 (149)	Faded
cw11-m301	gpt, deepseek	SOL+XRP	+$1,484.4	+$785.3 (38) +	**Survived**
cw11-m302	fable, opus, gpt, gemini, qwen	BNB+XRP	+$894.5	+$437.4 (15)	**Survived**
cw11-s1	gpt	BTC+ETH+ SOL+BNB+XRP	+$747.1	+$618.0 (106) +	**Survived**

Config	Family	Coins	Home	New MONTH (n)	Verdict
cw11-wk1	opus, gpt, grok	SOL+XRP	+$822.1	+$614.9 (51) +	**Survived**
cw11-wk2	gpt, grok	BTC+ETH+ SOL+BNB+XRP	+$767.3	+$558.1 (85) +	**Survived**
cw10-m301	fable, gpt, qwen [remap]	SOL+XRP	+$1,023.9	+$193.1 (62)	**Survived**
cw10-m302	fable, gpt, qwen [remap]	BNB+XRP	+$972.4	-$76.2 (56)	Faded
cw10-s1	gpt [remap]	SOL+XRP	+$415.2	+$662.4 (64) +	**Survived**
cw10-wk1	fable, opus, gpt, gemini, qwen [remap]	SOL+XRP	+$528.8	+$705.5 (79) +	**Survived**
cw10-wk2	qwen, grok [remap]	BTC+ETH	+$176.5	+$15.1 (126)	**Survived**
cwq1-q1	fable, opus, gpt, gemini [remap]	SOL+XRP	+$2,271.8	-$513.2 (112) +	Faded
cwq1-q2	fable, opus, deepseek, gemini [remap]	BNB+XRP	+$1,107.1	+$109.0 (26)	**Survived**
cwq1-s1	gemini	SOL+XRP	+$1,491.0	-$577.1 (167) +	Faded
cw9-m301	fable, opus, gpt, qwen, grok [remap]	BNB+XRP	+$797.6	-$350.0 (41) +	Faded
cw9-m302	fable, opus, qwen, grok [remap]	SOL+XRP	+$784.9	-$534.4 (52) +	Faded
cw9-s1	gemini	SOL+XRP	+$266.9	-$577.1 (167) +	Faded
cw9-wk1	deepseek, gemini, grok [remap]	SOL+XRP	+$250.4	-$280.3 (120) +	Faded
cw9-wk2	opus, gpt, deepseek, gemini [remap]	BTC+ETH	+$216.9	-$218.9 (80) +	Faded
cwm2-m1	fable, opus, gemini, qwen [remap]	BNB+XRP	+$1,041.5	-$366.5 (49) +	Faded
cwm2-m2	fable, opus, gemini, qwen [remap]	SOL+XRP	+$1,021.7	-$650.3 (54) +	Faded
cwm2-s1	gemini	SOL+XRP	+$923.2	-$490.6 (150) +	Faded
cw8-m301	fable, opus, gemini, qwen [remap]	SOL+XRP	+$1,018.6	-$650.3 (54) +	Faded
cw8-m302	fable, opus, gemini, qwen [remap]	BNB+XRP	+$943.0	-$366.5 (49) +	Faded
cw8-s1	gemini	BNB+XRP	+$445.4	-$514.6 (132) +	Faded
cw8-wk1	gpt, gemini, qwen [remap]	SOL+XRP	+$583.4	-$280.5 (82) +	Faded
cw8-wk2	gpt, gemini, qwen, grok [remap]	BNB+XRP	+$570.6	-$287.4 (88) +	Faded
cw7-m301	fable, gpt, gemini, qwen [remap]	SOL+XRP	+$1,351.8	-$559.6 (103) +	Faded
cw7-m302	opus, gpt, gemini, qwen [remap]	BTC+ETH	+$528.4	-$271.3 (13) +	Faded
cw7-s1	gemini	SOL+XRP	+$1,022.0	-$484.2 (195) +	Faded
cw7-wk1	fable, opus, gpt, gemini [remap]	SOL+XRP	+$1,467.8	-$513.2 (112) +	Faded
cw7-wk2	fable, opus, gpt [remap]	BTC+ETH	+$953.2	-$182.2 (39) +	Faded
cw6-m301	opus, deepseek, gemini [remap]	SOL+XRP	+$1,793.8	-$551.9 (81) +	Faded
cw6-m302	gpt, deepseek, gemini, grok [remap]	BTC+ETH	+$547.8	-$363.2 (29) +	Faded
cw6-s1	grok [remap]	SOL+XRP	+$404.8	+$13.1 (110) +	**Survived**
cw6-wk1	opus, gpt, grok [remap]	SOL+XRP	+$609.5	-$579.5 (70) +	Faded
cw6-wk2	fable, opus, deepseek [remap]	BTC+ETH	+$135.4	-$163.5 (59) +	Faded
cw5-m301	fable, opus, gpt, gemini [remap]	SOL+XRP	+$1,717.1	+$400.1 (18) +	**Survived**
cw5-s1	gemini	BNB+XRP	+$178.7	-$281.0 (150) +	Faded
cw5-wk1	gpt, gemini [remap]	SOL+XRP	+$510.4	-$322.2 (177) +	Faded
cw5-wk2	fable, opus, gpt, gemini, grok [remap]	BTC+ETH	+$183.5	-$168.4 (87) +	Faded
cw4-m301	gpt, deepseek, gemini [remap]	SOL+XRP	+$1,676.0	-$633.3 (29) +	Faded
cw4-m302	fable, gpt, deepseek, gemini [remap]	BNB+XRP	+$1,155.2	-$601.2 (13) +	Faded
cw4-s1	deepseek	BTC+ETH+ SOL+BNB+XRP	+$325.0	+$552.8 (74)	**Survived**

Config	Family	Coins	Home	New MONTH (n)	Verdict
cw4-wk1	opus, deepseek, gemini, qwen [remap]	BTC+ETH+ SOL+BNB+XRP	+$399.4	-$139.6 (52)	Faded
cw4-wk2	fable, deepseek, qwen	BNB+XRP	+$128.2	-$599.2 (79) +	Faded
cwm1-m1	fable, gemini	SOL+XRP	+$1,618.4	-$124.9 (28) +	Faded
cwm1-m2	fable, gpt, deepseek, gemini, qwen [remap]	ETH+SOL+ BNB+XRP	+$756.6	-$118.8 (40) +	Faded
cwm1-s1	gemini	SOL+XRP	+$1,721.9	-$416.6 (29) +	Faded
cw3-m301	opus, deepseek, gemini [remap]	SOL+XRP	+$2,137.6	-$231.7 (71) +	Faded
cw3-m302	fable, opus, deepseek, gemini, qwen [remap]	BTC+ETH	+$412.3	+$22.6 (11)	**Survived**
cw3-s1	gemini	SOL+XRP	+$1,604.3	+$64.6 (184) +	**Survived**
cw3-s2	gemini	SOL+XRP	+$1,449.0	-$416.6 (29) +	Faded
cw3-wk1	opus, deepseek, gemini [remap]	SOL+XRP	+$2,175.1	-$231.7 (71) +	Faded
cw3-wk2	opus, deepseek, gemini, grok [remap]	BTC+ETH	+$837.6	-$575.4 (25) +	Faded
cw2-30d1	opus, gpt, deepseek, gemini, qwen [remap]	BTC+ETH+ SOL+BNB+XRP [conc]	+$744.7	+$467.7 (30) +	**Survived**
cw2-30d2	opus, deepseek, gemini [remap]	ETH+SOL+ BNB+XRP [conc]	+$655.1	-$65.0 (36) +	Faded
cw2-30d3	opus, deepseek, gemini [remap]	BNB+XRP	+$489.1	+$54.6 (28) +	**Survived**
cw2-30ds1	gemini	BNB+XRP	+$383.8	-$162.7 (27) +	Faded
cw2-7d1	fable, opus, deepseek, gemini [remap]	SOL+XRP	+$299.1	-$583.6 (46) +	Faded
cw2-7d2	gpt, deepseek, gemini, qwen [remap]	BNB+XRP	+$223.7	-$13.8 (21) +	Faded
cw2-pt1	fable, deepseek, grok [remap]	SOL	+$698.3	-$387.9 (62)	Faded
cw2-pt2	opus [remap]	XRP	+$543.7	-$134.5 (9) *	Faded
cw1-30d1	fable, deepseek, opus, qwen [remap]	ETH+SOL	+$868.4	-$274.5 (33) +	Faded
cw1-30d2	gemini, gpt, deepseek [remap]	SOL+XRP+ETH	+$590.6	+$132.8 (41) +	**Survived**
cw1-30d3	fable, deepseek, opus, qwen [remap]	BTC+ETH+SOL	+$580.5	-$334.1 (48) +	Faded
cw1-30ds1	opus [remap]	SOL+XRP	+$442.8	-$411.7 (93) +	Faded
cw1-30ds2	grok [remap]	SOL+XRP	+$262.9	-$144.5 (95) +	Faded
cw1-30ds3	deepseek	XRP+BNB	+$119.0	-$556.6 (114) +	Faded
cw1-7d1	fable, deepseek, opus, gpt [remap]	SOL+XRP	+$439.2	-$255.4 (87) +	Faded
cw1-7d2	fable, qwen, gemini, opus [remap]	XRP+BNB	+$429.5	-$319.1 (90) +	Faded
cw1-7d3	gpt, opus [remap]	BNB+XRP+SOL	+$398.6	-$701.1 (90) +	Faded
cw1-7ds1	gemini	XRP+BNB	+$417.8	-$591.3 (73) +	Faded
cw1-7ds2	gpt [remap]	XRP+BNB	+$366.3	-$386.1 (129) +	Faded
cw1-7ds3	qwen	XRP+BNB	+$267.0	-$364.7 (94) +	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] = a superseded model id (claude-opus-5, gpt-5.6-sol, grok-4.6 or grok-4.5) replayed on its lineage head; none of the issue-12 and issue-11 rows carries it. Verdict rule: Survived = net-positive on the calendar month. Survivors are tinted green.

Counted: 21 of 79 past cards are net-positive on the calendar month (calendar September: 17 of 53) -- by cohort 3 of 5 issue-12 cards, 5 of 5 issue-11 cards, 4 of 5 issue-10 cards, 1 of 3 Quarterly-1 cards, 0 of 5 issue-9 cards, 0 of 3 Monthly-2 cards, 0 of 5 issue-8 cards, 0 of 5 issue-7 cards, 1 of 5 issue-6 cards, 1 of 4 issue-5 cards, 1 of 5 issue-4 cards, 0 of 3 Monthly-1 cards, 2 of 6 issue-3 cards, 2 of 8 issue-2 cards, 1 of 12 issue-1 cards. Config

Watch #13 read all 79 of these cards on two windows that lie mostly inside this one: 16 of 79 on its fresh week, 17 of 79 on its four weeks, 7 surviving both. All 7 of those survivors are net-positive here, and so are 14 of its 72 faded cards. On this collect's own four-week window the 79 past cards read 29. The best replay is cw12-m301 at +$959.2 on 32 closed trades and the worst cw1-7d3 at -$701.1 on 90. Live presets: 1 of 4 net-positive on the month (calendar September: 0/4), 0 of 4 survived both of Config Watch #13's windows.

The four live presets lead the table by standing rule. On the month: sol -$937.9 on 39 closed, xrp -$404.3 on 4 closed, bc2 +$73.9 on 23 closed, 4h4 -$383.4 on 32 closed -- 1 of the four net-positive, where Config Watch #13 recorded 2/4 on its fresh week and Monthly Config Watch #2 0/4 on calendar September. **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 +$73.9 is the same simulation on both sides of the report.

Origin re-runs (drift baseline). Re-simulating the issue-12 cards on their OWN origin windows in this collect lands within -$122.6..$0.0 of the numbers issue #12 printed: 4 of the 5 reproduce to the cent. The issue-11 cards drift -$139.7..$0.0 (4 of 5 to the cent). The largest gap is cw11-m301 (+$1,484.4 printed, +$1,344.6 on the re-run). Cards with positions still open when their origin window closed re-run to their settled result; issue #12 printed open positions for cw12-wk1 and issue #11 for cw11-m301, and those are the cards that do not reproduce to the cent. Snapshot principle: no past issue is restated; the full drift table is in the open data.

What this column is. 21 of 79 past cards are net-positive over 31 days. Three measured facts bound the reading. The window overlaps the cards' own: four of the fifteen cohorts searched on days inside it (20 rows). The rows Config Watch #13 dropped on its two-window rule did not all lose here: 14 of its 72 faded cards are net-positive on the month, and 7 of its 7 survivors. And the fresh search moved with them: the month's best mature council is +$1,119.4 (the grid's maximum, +$1,311.1, is the mature solo board's leader) against +$880.9 on the four-week window of this collect and +$333.6 on Config Watch #13'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 fifteen card rows of this table that have a fresh week, as shares; it is printed once, there. Twelve fresh weeks and the calendar month do not fit one table width, so the table prints in two parts: +1 to +6 weeks with the calendar month, then +7 to +12 weeks for the generations that reach them.

Generation	Cards	+1 week	+2 weeks	+3 weeks	+4 weeks	+5 weeks	+6 weeks	Calendar Oct
Config Watch #1	12	6/12	11/12	1/12	2/12	3/12	11/12	1/12
Config Watch #2	8	8/8	4/8	1/8	2/8	5/8	6/8	2/8
Config Watch #3	6	2/6	2/6	2/6	5/6	2/6	2/6	2/6
Config Watch #4	5	1/5	0/5	4/5	4/5	3/5	0/5	1/5
Monthly Config Watch #1	3	2/3	0/3	1/3	3/3	2/3	0/3	0/3
Config Watch #5	4	1/4	3/4	2/4	3/4	0/4	1/4	1/4
Config Watch #6	5	4/5	2/5	3/5	0/5	4/5	2/5	1/5
Config Watch #7	5	4/5	4/5	0/5	2/5	3/5	2/5	0/5
Config Watch #8	5	2/5	0/5	3/5	0/5	0/5	--	0/5
Monthly Config Watch #2	3	3/3	0/3	0/3	1/3	0/3	--	0/3
Config Watch #9	5	0/5	3/5	1/5	0/5	--	--	0/5
Quarterly Config Watch #1	3	1/3	3/3	1/3	0/3	--	--	1/3
Config Watch #10	5	4/5	2/5	1/5	--	--	--	4/5
Config Watch #11	5	4/5	3/5	--	--	--	--	5/5
Config Watch #12	5	1/5	--	--	--	--	--	3/5
Live presets	4	2/2	2/4	1/4	3/4	3/4	2/4	1/4

Generation	Cards	+7 weeks	+8 weeks	+9 weeks	+10 weeks	+11 weeks	+12 weeks
Config Watch #1	12	1/12	1/12	0/12	7/12	6/12	1/12
Config Watch #2	8	2/8	2/8	5/8	0/8	3/8	--
Config Watch #3	6	1/6	4/6	4/6	1/6	--	--
Config Watch #4	5	4/5	1/5	2/5	--	--	--
Monthly Config Watch #1	3	2/3	0/3	0/3	--	--	--
Config Watch #5	4	4/4	1/4	--	--	--	--
Config Watch #6	5	1/5	--	--	--	--	--
Live presets	4	0/4	0/4	1/4	0/4	2/4	--

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), Config Watch #9 (Sep 28-Oct 5, 2026), Config Watch #10 (Oct 5-12, 2026), Config Watch #11 (Oct 12-19, 2026), Config Watch #12 (Oct 19-26, 2026), Config Watch #13 (Oct 26-Nov 2, 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 and Monthly-2 at #9, issue-9 and Quarterly-1 at #10, issue-10 at #11, issue-11 at #12, issue-12 at #13. The live presets are not a generation and have no publication date: their eleven cells are the eleven issues that replayed them -- Config Watch #3 replayed 2 of the 4, #4 to #13 all 4 -- and that row is not drawn on the chart. Config Watch #13's own cards are not a generation of this table: this collect does not replay them (walk-forward above). Calendar Oct is this issue's window (Oct 1-31, 2026) for every row; Monthly Config Watch #2 read the issue-1 row at 2/12, the issue-2 row at 0/8, the issue-3 row at 0/6, the issue-4 row at 2/5, the Monthly-1 row at 1/3, the issue-5 row at 2/4, the issue-6 row at 4/5, the issue-7 row at 4/5, the issue-8 row at 2/5 and the presets row at 0/4 on calendar September. 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-12's is Oct 26-Nov 2, 2026), and issue-1's +1 week overlaps its own search window (Aug 5-12, 2026) by 2 days; of the monthly and quarterly generations, Monthly-1's overlaps calendar August by 1 day (Aug 31), Monthly-2's overlaps calendar September by 3 days (Sep 28-30) and Quarterly-1's shares no day with the quarter (Jul 1-Sep 30, 2026).

Survivors vs faded: structure

Config Watch #13 replayed 79 past cards on two fresh windows (Oct 26-Nov 2, 2026 and Oct 5-Nov 2, 2026) and its standing rule kept 7 of them (net-positive on both) and dropped 72. 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: 16 cards cleared the week, 17 the four weeks and 7 both. 7 rows against 72 rows: this is a description of two groups, not a test, and no trait here is claimed to cause survival.

Trait	Survivors (7)	Faded (72)
FH/TF	1d/4h (3), 1d/1d (2), 1d/1d,4h (1), 4h/4h (1)	4h/1h (34), 1d/1d (15), 4h/4h,1h (9), 4h/4h (8), 1d/1d,4h (4), 1d/4h (1), 1h/1h (1)
Council size	2 models (3), solo (3), 4 models (1)	4 models (22), solo (19), 3 models (18), 5 models (7), 2 models (6)
Universe size	2 coins (4), 5 coins (3)	2 coins (63), 3 coins (3), 1 coin (2), 4 coins (2), 5 coins (2)
Calibrate ON	0 of 7 (0.0%)	14 of 72 (19.4%)
Min confidence	60 (6), 70 (1)	60 (64), 0 (3), 50 (3), 55 (2)
Table stake	$200 (3), $450 (3), $375 (1)	$450 (66), $300 (3), $200 (1), $250 (1), $375 (1)
Median maxDD	0.0%	29.2%
Median win rate	88.9%	57.1%
Median smoothness R2	0.679	0.582
Model membership	gpt (6), deepseek (2), opus (2), fable (1), gemini (1), grok (1)	gemini (49), opus (35), fable (29), gpt (29), deepseek (27), qwen (26), grok (13)

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 #13 fresh-week (Oct 26-Nov 2, 2026) replay -- one of the two windows the verdict needed, and the R2 median is taken over the rows that have one (5 of 7 survivors, 55 of 72 faded: a replay with too few days carries no fit). On Config Watch #13's four-week window (Oct 5-Nov 2, 2026) the same medians read maxDD 36.8% against 52.2% and win rate 74.2% against 58.4%. On this issue's calendar month the same two groups read 7 of 7 and 14 of 72 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)	1,338	34.9%	-$43.2	+$605.3	-$757.2
-- mature rows only	724	42.5%	-$52.5	+$605.3	-$757.2
Board-leader slice (top of board)	14	28.6%	-$149.4	+$421.1	-$710.4

On this window the axis is negative in the population and negative at the top. Calibration improves 34.9% of the 1,338 matched pairs on the whole grid (mean -$43.2, median -$3.9) and 42.5% of the 724 mature pairs (mean -$52.5). On the board-leader slice it improves 28.6% of 14 pairs (mean -$149.4, median -$49.8): the OFF legs there are at a median of +$1,024.8 and the ON legs at +$940.3, 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 40.2% of its pairs, on Config Watch #13's fresh week 56.8% and on calendar September (Monthly Config Watch #2) 65.6%. Where the axis reached the cards is visible above: both council cards are calibrated cells and neither solo card is, and 17 of the 41 finalists run it.

The open-data pack also carries a 176-pair export for this window. Its delta distribution is not the population's (median -$354.2 against the grid's -$3.9), 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 (Oct 1-31, 2026)

#	Model	Coins	FH/TF	Net	Net %	maxDD	Trades	WR
#1	grok	SOL+XRP	1d/4h	**+$1,311.1**	+131.1%	24.5%	20	90.0%
#2	gpt	SOL+XRP	1d/1d	**+$1,162.6**	+116.3%	44.7%	40 *	82.5%

#	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	0/0	reenter	$450

Two solo cards, one per model (grok and gpt); only #1 carries a short link. No single-coin row outranks it on that board this window. The card carries no unfunded entry.

Council vs solo. On this window the best solo sits **above** the best council (+$1,311.1 against +$1,119.4 on the mature boards), on 20 closed trades against 26, with a 24.5% drawdown against 15.5% and 0 unfunded entries against 0. Config Watch #13 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/cwm3-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** = 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. 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	gpt	ETH+SOL+ BNB+XRP	1d/1d	+$799.4	96.0%	0.5%	25	0.938
MONTH Calmar	gpt	ETH+SOL+ BNB+XRP	1d/1d	+$799.4	96.0%	0.5%	25	0.938
MONTH Calmar	gpt	ETH+SOL+ BNB+XRP	1d/1d	+$639.5	96.0%	0.4%	25	0.938
MONTH WR	gpt [cal]	ETH+SOL+ BNB+XRP	1d/1d	+$625.7	96.2%	0.5%	26	0.974
MONTH WR	gpt [cal]	ETH+SOL+ BNB+XRP	1d/1d	+$625.7	96.2%	0.5%	26	0.974
MONTH WR	gpt	ETH+SOL+ BNB+XRP	1d/1d	+$799.4	96.0%	0.5%	25	0.938
MONTH Smoothness	gpt [cal]	ETH+SOL+ BNB+XRP	1d/1d	+$625.7	96.2%	0.5%	26	0.974
MONTH Smoothness	gpt [cal]	ETH+SOL+ BNB+XRP	1d/1d	+$625.7	96.2%	0.5%	26	0.974
MONTH Smoothness	grok	SOL+XRP	1d/4h	+$1,311.1	90.0%	24.5%	20	0.971
MONTH Quality	grok	SOL+XRP	1d/4h	+$1,311.1	90.0%	24.5%	20	0.971
MONTH Quality	grok [cal]	SOL+XRP	1d/4h	+$940.3	95.0%	13.1%	20	0.951
MONTH Quality	grok	SOL+XRP	1d/4h	+$1,311.1	90.0%	24.5%	20	0.971

[cal] = calibrate=ON cell; the board column names the window. The Calmar board is degenerate by construction -- its top rows sit at 0.5% 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 5 of them are calibrated. The pool is 2,950 mature configs against 2,827 on the four-week window of the same collect and 2,967 on calendar September.

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 Sep 28-Oct 26, 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 25 days (Oct 1-25), Sep 28-30 lies in the four-week window only and Oct 26-31 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. 20 configs qualify (under the collect's row cap of 25), on 2 universes (SOL+XRP 18 and SOL 2), 8 of them councils and 6 clean (zero unfunded entries on both windows). The top three by rank sum are solo gpt (+$1,157.5); solo gpt (+$1,157.5); solo gpt (+$1,162.6). Monthly Config Watch #2 read the same pair of windows on calendar September from its own collect and held 22 and Monthly Config Watch #1 on calendar August 25, the collect's row cap -- the comparable boards; Config Watch #13's board (0 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	+$73.9	69.6%	53.7%	0.214	23.03	23

818 of the 2,950 mature configs (27.7%) clear all three bars at once. The anchor's own month is +$73.9 net at 69.6% win rate with a 53.7% drawdown, an R2 of 0.214 and 23 closed trades. Of the ten best of those 818 by net, all ten beat the anchor on net PnL as well. 4 of the ten are calibrated cells, 2 are councils, their stakes are $375 and $450 and their universes are 2 coins. On the four-week window of the same collect the same rule admitted 1,657 of 2,827 configs (58.6%), and on calendar September (Monthly Config Watch #2) 679 of 2,967 (22.9%), 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 818 configs that beat the anchor on win rate, maxDD and smoothness at once

#	Config	Coins	FH/TF	Stake	Net	WR	maxDD	R2
#1	grok	SOL+XRP	1d/4h	$450	+$1,311.1	90.0%	24.5%	0.971
#2	grok	SOL+XRP	1d/4h	$450	+$1,311.1	90.0%	24.5%	0.971
#3	gpt	SOL+XRP	1d/1d	$450	+$1,162.6	82.5%	44.7%	0.802
#4	gpt	SOL+XRP	1d/1d	$450	+$1,162.6	82.5%	44.7%	0.802
#5	gpt [cal]	SOL+XRP	1d/1d	$450	+$1,157.5	88.1%	22.4%	0.922
#6	gpt [cal]	SOL+XRP	1d/1d	$450	+$1,157.5	88.1%	22.4%	0.922
#7	gpt	SOL+XRP	1d/1d	$375	+$1,155.6	81.4%	37.2%	0.851
#8	gpt	SOL+XRP	1d/1d	$375	+$1,155.6	81.4%	37.2%	0.851
#9	gpt, grok [cal]	SOL+XRP	1d/1d	$450	+$1,119.4	88.5%	15.5%	0.916
#10	gpt, grok [cal]	SOL+XRP	1d/1d	$450	+$1,119.4	88.5%	15.5%	0.916

Anchor passport (verbatim, as stored): members=gemini-3.1-pro,gpt-6.1-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 1,090 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	opus, gpt, grok [council]	1d/4h	+$191.0 (10)
ETH	gpt [solo]	1d/1d,4h	+$742.0 (28)
SOL	gpt [solo]	1d/1d	+$959.7 (25)
BNB	gemini [solo]	4h/4h,1h	+$359.1 (37)
XRP	opus, grok [council]	4h/4h	+$644.6 (38)

n in parentheses; * = fewer than 10 closed trades (thin -- direction only). 1 of the five winners run calibrate ON and 2 of the five are councils.

The reads. Every coin is net-positive over the month and no row is thin: 10 to 38 closed trades, SOL leading at +$959.7 and BTC lowest at +$191.0. 1 of the five sits below the >=20-trade maturity gate the council boards apply (BTC 10), so that row is 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 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 (claude-opus-5.5)	-$113.2 *	755	63%	11.3%
Worst solo default (deepseek-v4-pro)	-$243.4 *	1161	55%	24.8%
Council-of-7 default	-$69.1 *	575	64%	7.7%
Hand-tuned config A (pre-search)	+$25.4 *	27	74%	42.4%
Hand-tuned config B (pre-search)	-$731.3 *	101	52%	83.1%
Hand-tuned config C (pre-search)	-$274.5 *	33	58%	51.6%

* = 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 (735 to 1,478 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 -$277.9 at its own stake against -$274.5 at the table stake.

Over the month the standing ordering holds in one half only. The seven solo defaults run -$113.2 to -$243.4 at table stake and the council-of-7 -$69.1 -- every default lost money at table stake; the three pre-search hand-tuned configs read +$25.4 / -$731.3 / -$274.5, and the fresh search's +$1,119.4 is above the best reference point in the table (+$25.4). 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: 1 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	+$73.9	+$661.8 (1d/1d, SOL+XRP, 43 cls)	34	neighbour +$587.9
preset 4h4	-$383.4	+$109.1 (1d/1d, ETH+SOL, 3 cls)	19	neighbour +$492.5
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 +$587.9 (+$73.9 -> +$661.8; same_side update -> reenter); 4h4 by +$492.5 (-$383.4 -> +$109.1; fh 4h -> 1d, tf 1h -> 1d). 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; 41 unique configs, 17 of them councils and 24 solo).

Knob	MONTH finalists mode
Forecast horizon	1d (41/41)
Ladder shape (steps)	50:40,100:60 (41/41)
Break-even stop (BE)	on (41/41)
Min confidence	60 (31/41)
SL shift	75 (41/41)
TP shift	-10 (41/41)
TP/SL source	farthest (29/41)
Calibrate ON	17/41
Universe	SOL+XRP (25/41)

The mechanical core: 50:40,100:60 ladder (41/41), break-even on (41/41), conf 60 (31/41), SL +75 / TP -10 (41/41), farthest source (29/41) -- the same core Monthly Config Watch #2 printed for calendar September. The horizon mode is **1d** (41/41; calendar September: 4h (33/49)). **Calibrate=ON** runs in 17 of the 41 finalists, against 24 of 49 on the four-week window of this collect and 31 of 49 on calendar September. The universe is SOL+XRP in 25 of the 41 finalists, and 1 of them is a single-coin row. Membership across the 17 council finalists: gpt 17, grok 8, opus 6, deepseek 5. 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-79 survival on Config Watch #13's two windows and 21-of-79 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 15.5% (council #1), 15.6% (council #2) and 24.5% (the solo card), and the replay table is deeper: 61 of the 83 replayed rows drew down more than 40% over the month and 41 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 (Sun Nov 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 #8 through Monthly Config Watch #2's table and Config Watch #9 to #13 directly) and the structure table is computed from the frozen queries and the replay statistics of the cards Config Watch #13 already replayed. Each monthly issue appends the fresh weeks published since and the generations 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 (Oct 1-31, 2026) on a single regime; the month-over-month figures set it beside the previous month (calendar September, Monthly Config Watch #2), and no past issue is restated. Trades settle past the window edge at their horizon; open = a position still open when the collect ran (Nov 5, 10:36 UTC, 106 hours after the window edge): 1 on the four published cards, 13 of the 166 replay cells (83 rows x 2 windows) carry 17 open positions between them -- listed in the open data.
- **This month overlaps the weekly windows, so it is not independent of them.** 27 of its 31 days lie in the four published Config Watch weeks it is read against (Config Watch #10 Oct 5-12, 2026; Config Watch #11 Oct 12-19, 2026; Config Watch #12 Oct 19-26, 2026; Config Watch #13 Oct 26-Nov 2, 2026) and 25 in Config Watch #12'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-12, issue-11, issue-10 and issue-9 overlap this month, so 20 of the 79 replayed cards are partly scored on the days they were selected on; the older cohorts' windows end before it. Publication dates bound it too: 31 of the 79 cards were published inside or after the month (issue-12 on Oct 28, issue-11 on Oct 22, issue-10 on Oct 14, Quarterly-1 on Oct 9, issue-9 on Oct 7, Monthly-2 on Oct 3 and issue-8 on Oct 1), so for them the October days before publication are not a forward test of a published card; only the 48 cards of issue-7, issue-6, issue-5, issue-4, Monthly-1, issue-3, issue-2 and issue-1 were published before Oct 1 (issue-7 on Sep 23, issue-6 on Sep 16, issue-5 on Sep 9, issue-4 and Monthly-1 on Sep 4, issue-3 on Aug 26, issue-2 on Aug 19 and issue-1 on Aug 13), and they read 8 of 48 here. The 79 cards are 73 distinct configurations window aside (cwq1-q1 = cw7-wk1; cwq1-s1 = cw9-s1; cwm2-m1 = cw8-m302; cwm2-m2 = cw8-m301; 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, September, October and November, the Calendar Oct 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 that share 25 of their days -- so it measures persistence across largely overlapping windows; it is comparable with the boards of Monthly Config Watch #2 and Monthly Config Watch #1, read across the same pair, and not with the weekly issues' boards built across WK and M30.
- **Model lineage is a splice.** Three model axes run as lineages (claude-opus-5.5 (incl. opus-5-era), gpt-6.1-sol (incl. 5.6-era), grok-4.7 (incl. 4.6-era)) across the cutover at Oct 17, 2026 09:01 UTC, which lies inside the calendar month (Oct 1-31, 2026, day 17 of 31) and inside the four-week window of this collect (Sep 28-Oct 26,

2026, day 20 of 28): every number in this issue joins the two builds of each line under one id -- on the month 16 days on the old builds and 15 from the cutover on. Every published card that named a superseded id is replayed on the lineage head ([remap] rows: 57 of the 83 rows -- 53 of the 79 cards (by cohort issue-10 5 of 5, Quarterly-1 2 of 3, issue-9 4 of 5, Monthly-2 2 of 3, issue-8 4 of 5, issue-7 4 of 5, issue-6 5 of 5, issue-5 3 of 4, issue-4 3 of 5, Monthly-1 1 of 3, issue-3 4 of 6, issue-2 7 of 8 and issue-1 9 of 12) and all four live presets), and a remapped replay is not the same simulation the origin issue ran. The grok 4.5 -> 4.6 cutover of Aug 24, 2026 lies entirely before both windows.

- **The cutover hours.** The two roll-out slots of Oct 17 (07:01 and 08:01 UTC; invalid rows for the three superseded ids, council of four in those hours) fall inside the calendar month and inside the four-week window.
- **Calibrated replays are pinned to their origin point.** 13 replay rows carry calibrate=1 with the calibration point fixed in their stored query (calibrate_asof: 2026-08-25, 2026-09-01, 2026-09-07, 2026-09-14, 2026-09-15, 2026-09-21, 2026-10-01, 2026-10-06, 2026-10-20, 2026-10-26), so their TP/SL multipliers on the month are the ones in force when the origin issue was computed -- including the issue-12 card cw12-wk2, pinned to Oct 26, 2026 at publication; up to Config Watch #11 a calibrated replay used the live point of the day it ran, and the replay signatures of Monthly Config Watch #2 (calendar September) carried no pinned point. The rows: cw11-m302, cw12-wk2, cw3-m302, cw4-wk2, cw5-wk1, cw5-wk2, cw6-m302, cw6-wk2, cw7-m302, cw9-wk2, cwm1-m2, cwm2-s1, cwq1-q2.
- **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 82,334 resolved forecasts at the monthly cutoff (Sun Nov 1, 2026, 16:00 UTC), where Config Watch #13 printed 82,895 at its own cutoff (Nov 2, 2026 16:00 UTC) and Monthly Config Watch #2 59,755 at its own (Thu Oct 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.** 10,659 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 2 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)	1,090
Systematic grid (s2)	7,904	OOS replays (s4cw12+s4cw12o+s4cw11+s4cw11o+s4cw10+s4cw10o+s4cwq1+s4cwq1o+s4cw9+s4cw9o+s4cwm2+s4cwm2o+s4cw8+s4cw8o+s4cw7+s4cw7o+s4cw6+s4cw6o+s4cw5+s4cw5o+s4cw4+s4cw4o+s4cwm1+s4cwm1o+s4cw3+s4cw3o+s4cw2+s4cw2o+s4cw1+s4pre)	233
Refinement: ladder + stake sweep (s2b+s3b)	254	**Total**	**10,659**
Preset neighbourhood grid (s2p)	106	Errors / zero-trade	0 / 1,554
Per-ticker probes (s5t)	120	Engine	1.1

Scale note. This issue's collect is a run of its own: 10,659 simulations over its two windows, against 11,075 in Monthly Config Watch #2's collect. Config Watch #13 ran two separate collects (5,392 weekly + 10,573 night = 15,965), and so did the three weekly issues before it that the verdict block reads (Config Watch #10 16,283, Config Watch #11 15,740 and Config Watch #12 16,016); this issue uses none of them. The per-window stage-2 grid is 3,952 cells here and 3,952 on Config Watch #13's fresh week. The replay branch is 83 rows -- 79 past cards and 4 live presets -- for 233 simulations, because every row is replayed on both windows of this collect and the cards of fourteen 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. **Model lines in these windows.** The collect's lineage probe reads 10 raw model ids

in the four-week window: claude-opus-5.5 joins claude-opus-5 (4,095 rows to Oct 17 08:01 UTC) and claude-opus-5.5 (1,785 from Oct 17 09:01); gpt-6.1-sol joins gpt-5.6-sol (4,095 rows to Oct 17 08:01 UTC) and gpt-6.1-sol (1,785 from Oct 17 09:01); grok-4.7 joins grok-4.6 (4,095 rows to Oct 17 08:01 UTC) and grok-4.7 (1,785 from Oct 17 09:01) -- spliced at the cutover of Oct 17, 2026 09:01 UTC, inside both windows; the other four lines run one id each.

Source files: summary_month.json (the w373 collect) with grid_hist_month, equity_month, baselines_month, presets_recon_month and the fifteen oos_*_month packs beside it; the published open data of Monthly Config Watch #2, Quarterly Config Watch #1 and Config Watch #4 to #13 for every number this issue quotes from them. Generated at: 2026-11-05T10:36:34Z · 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: 10,659 simulations (0 errors) in one frozen collect, of which 233 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 Nov 1, 2026 16:00 UTC): 82,334 resolved since Jul 11 · 7 models tracked (current line-up; earlier versions folded into their successors' lineage) · 5 assets · 5 horizons · hourly · 70 published reports

Counter pinned from wave 3 onward and read at this issue's cutoff (Sun Nov 1, 2026, 16:00 UTC), where Config Watch #13 read 82,895 at its own cutoff and Monthly Config Watch #2 59,755 at its own; the model line counts the current line-up, with earlier versions folded into their successors' lineage. Monthly Config Watch #3 is the 70th published report: Monthly Consensus Watch #3, Monthly Calibration #3 and Monthly Model Watch #3 are 67, 68 and 69, and Config Watch #13 is 66.

Issue #3. 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 October -- a best mature council of +$1,119.4 (September: +$1,041.5) and a best mature solo of +$1,311.1 (September: +$923.2), 21 of 79 published cards net-positive on the month (September: 17 of 53), a survival curve that reaches +12 fresh weeks for the issue-1 generation (September: +7), and a structural table of Config Watch #13's 7 survivors against its 72 faded cards. The calendar week and the four-week window stay in the weekly Config Watch.

What we're testing next

- **Monthly Config Watch #4** -- the November calendar month (2026-11-01 to 2026-12-01) 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 and for Config Watch #13's.
- **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 #4 prints those columns for the November weeks, and the Config Watch #13 generation joins the curve with its first fresh week.
- **Quarterly Config Watch #2** -- the fourth calendar quarter as one window: the weekly and monthly issues of the quarter on one axis, with the survival curve as its spine, as Quarterly Config Watch #1 read the third.

Related research

Monthly Consensus Watch #3 -- https://marketmania.ai/research/reports/consensus-watch-monthly-2026-10.pdf
Monthly Calibration #3 -- https://marketmania.ai/research/reports/calibration-monthly-2026-10.pdf
Monthly Model Watch #3 -- https://marketmania.ai/research/reports/model-watch-monthly-2026-10.pdf
Config Watch #13 -- https://marketmania.ai/research/reports/config-watch-2026-10-28.pdf
Monthly Config Watch #2 -- https://marketmania.ai/research/reports/config-watch-monthly-2026-09.pdf

The three monthly reports publish together as one issue each month; this issue joins them on the same cadence. Config Watch #13 is the newest weekly issue: it carries the calendar week Oct 26-Nov 2, 2026 and the four-week window Oct 5-Nov 2, 2026 from collects of its own, and no board of either is reprinted here; Config Watch #12's four-week window (Sep 28-Oct 26, 2026) is the one this issue's collect recomputes. Monthly Config Watch #2 is the previous issue of this series and the source of every month-over-month number.

CITE THIS REPORT

```
@techreport{mm_configwatch_monthly_2026m10,
  title        = {Monthly Config Watch #3},
  author       = {{MarketMania Research}},
  institution  = {MarketMania},
  year         = {2026},
  month        = nov,
  day          = {5},
  type         = {Monthly Research Report},
  series       = {Monthly Config Watch},
  number       = {3},
  note         = {Methodology v1.1 + Config Watch amendments, hash e66c7e8c864a2233},
  url          = {https://marketmania.ai/research/reports/config-watch-monthly-2026-10.pdf}
}
```