

Config Watch #6

September 16, 2026 · MarketMania Research · Weekly series

WEEKLY · CONFIG WATCH · Issue #6 · the ceiling rises again, the board does not · Issue date September 16, 2026 · Language English · Methodology v1.1 + Config Watch amendments (approved 12.08, hash e66c7e8c864a2233)
Search windows: **WK** Sep 7-14, 2026 and **M30** Aug 17-Sep 14, 2026 -- both frozen and Monday-aligned (WK = Mon Sep 7 00:00 -> Sun Sep 13 23:59 UTC; M30 = Mon Aug 17 00:00 -> Sun Sep 13 23:59 UTC, 4 full weeks), sharing the cut-off edge Mon Sep 14 00:00 UTC, so slots dated Sep 7 fall outside both. 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 issue #1 and that the live presets run on.

RESEARCH SNAPSHOT

16,111 simulations, 0 errors, in two frozen collects	WK Sep 7-14, 2026	M30 Aug 17-Sep 14, 2026	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: ≥ 15 closed (WK) / ≥ 20 (M30)	criteria frozen in writing before launch (w244 plan)
walk-forward: 38 past cards + 4 live presets replayed verbatim on both windows	6 of 38 past cards Survive (net-positive on both windows)	stability: 9 configs across WK and M30 (night run)	models: 7 (grok axis = grok-4.6 (incl. 4.5-era))
best WK council +\$609.5 (+61.0%)	best M30 council +\$1,793.8 (+179.4%)	zero-trade cells: 1,121 (weekly) / 1,836 (night run)	margin-skip alerts on leaders: 0 (WK) / 4 (M30) rows


KEY FINDING · OBSERVATION (the ceiling rises again, the board does not)

The ceiling of the fresh search moved from **+\$510.4** to **+\$609.5** on the week while only **8** of the 38 cards this series has published are net-positive on that same week, and under the standing two-window rule 6 of the 38 Survive: 19 clear the month, 8 clear the week.

TL;DR

- OBSERVATION -- The standing board grew to 38 cards.** Issue #5 kept 6 of its 34 past cards under the same rule; this issue replays **38 cards** -- issues #1, #2, #3, #4, #5 and Monthly Config Watch #1 -- on two fresh windows, and **6 of the 38 Survive** (net-positive on both). On the week alone 8 of 38 are positive, on the four-week window 19 of 38. Issue #5's own week champion, replayed verbatim, makes +\$36.3 on the week and -\$578.6 on the month.
- The search ceiling, window by window.** The best mature council on WK makes +\$609.5 (+61.0% of a \$1,000 deposit, 80.5% win rate, 6.4% maxDD, 41 closed) against issue #5's +\$510.4; on M30 it makes +\$1,793.8 (+179.4%, 64.9% win rate, 52.8% maxDD, 168 closed) against issue #5's +\$1,717.1. The grids behind them: median -\$32.2 with 17.5% positive across 3,952 WK cells (10 of them above +\$500), median -\$0.1 with 30.3% positive across 4,080 M30 cells (167 above +\$500).
- Calibration and the universe, window by window.** Across the whole stage-2 grid calibration improves 49.9% of the 1,334 matched WK pairs (mean +\$60.3) and 48.4% of the 1,392 matched M30 pairs (mean +\$24.4); pooled, 49.1% of 2,726. On the cards, **1 of the 2 published WK councils are calibrated cells** and **1 of the 2 M30 cards are calibrated cells**. The mature net top-10 is a single universe on each window: SOL+XRP on WK (10 of 10 slots), SOL+XRP on M30 (10 of 10).
- The stability board, read against itself for the first time.** The night collect computes WK and M30 in one pass, so this issue's cross-window board is measured on the two windows the issue prints -- and **9 configurations** are in the mature net top-40 of both. Issue #5 measured the same pair one week earlier and held 0; issue #3, the other board across the same pair, held 21. The hand-tuned anchor keeps its standing measurement against the whole mature field of each window -- 0 of 1,781 mature WK configs and 465 of 3,265 mature M30 configs beat it on win rate, maxDD and smoothness at the same time. Net PnL stays the headline board (continuity with issues #1-#5).

Stage-2 systematic grid -- net PnL distribution -- WK window (Sep 7-14, 2026) -- n=3,952



Net PnL across all 3,952 stage-2 grid sims, WK window (Sep 7-14, 2026), universe-aware stake throughout. Grid median **-\$32.2**, 17.5% positive (issue #5's week: median \$0.0, 34.6% positive; its month: \$0.0, 47.6% positive). The spike at \$0 is mostly cells whose entry filters produced no trades (1,111 zero-trade sims on this window). Range **-\$798.4 to +\$609.5**; 10 cells finished above **+\$500**. Red marker = best solo default at table stake; green = the published cards. The same board definition one week earlier put its top-5 at **+\$510.4** down to **+\$465.9**; this week it runs **+\$609.5** down to **+\$568.9** -- a property of the week, not evidence about either issue's cards. The M30 histogram appears after the M30 cards. Source: grid_hist_wk.json.

Why it matters

MarketMania publishes default LLM-council trading configs. This series asks one narrower question every week: how much apparent performance can large-scale in-sample search extract -- and how much of it survives out-of-sample? Issue #1 set the in-sample baseline, issue #2 delivered the first walk-forward verdict (brutal), issue #3 the second (the opposite), issue #4 the third on two windows at once, issue #5 the fourth on the widest board the series had had. Issue #6 delivers the fifth, on a board that carries every generation the series has produced: 38 cards from five weekly issues and the first monthly one, each replayed verbatim on the same two fresh windows. That is the point of the exercise -- a rule that is applied to a growing, never-pruned list is the only walk-forward number that cannot be chosen after the fact.

How we searched

Two frozen collects, run to the same plan with the criteria frozen in writing *before* launch (w244 plan) and printed into the meta of every JSON deliverable: the weekly collect over the WK window, and the night collect, which sweeps the M30 window and the same WK week together in one pass. Each stage below prints **weekly + night** and both totals are printed whole; no calendar-month window is computed this issue. **Every WK number in this issue comes from the weekly collect and every M30 number from the night collect** -- the night run's own second pass over the WK week is used for one thing only, the cross-window stability board, and no number from it is printed beside a weekly-collect number.

- **1. Scan (s1)** -- 476 + 952 sims across council compositions and coarse settings.
- **2. Systematic grid (s2)** -- 3,952 + 8,160 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)** -- 118 + 238 sims around the grid winners.
- **4. Preset neighbourhood grid (s2p)** -- 53 + 106 sims: one-knob neighbours of the live presets that have a local grid this issue.
- **5. Per-ticker probes (s5t)** -- 60 + 108 sims: best finalists split onto single coins.
- **6. Dedicated per-ticker (s6t+s6tc)** -- 560 + 1,150 sims: a compacted independent per-coin grid plus calibrate twins (the standing branch from issue #2).
- **7. OOS replays**
(s4cw5+s4cw5o+s4cw4+s4cw4o+s4cwm1+s4cwm1o+s4cw3+s4cw3o+s4cw2+s4cw2o+s4cw1+s4pre) -- 68 + 110 sims: every published card and all four live presets, verbatim, on every fresh window, plus origin re-runs of the issue-5 and issue-4 cards for the drift baseline.

Total 16,111 simulations (5,287 weekly + 10,824 night), 0 errors, 2,957 zero-trade cells (1,121 + 1,836); stage sums reconcile exactly in each collect (counts.by_stage). The axis grid was NOT widened relative to issue #5 (anti-overfit rule): the replay branch grew by one cohort, the search space did not. Equity series, grid histograms, baselines and the calibrate pack are produced inside the same two collects -- there is no separate same-day extras pass this issue.

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 the leaders this issue publishes: 0 on WK (weekly collect) and 4 on M30 (night collect); the night collect raised 14 alerts in total, the rest on its own second run of the WK week, which this issue does not publish. Source: criteria in both collects.

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

TOP councils -- WK window (Sep 7-14, 2026)

#	Council	Coins	FH/TF	Net	Net %	maxDD	Trades	WR
#1	opus, gpt, grok	SOL+XRP	4h/1h	+\$609.5	+61.0%	6.4%	41	80.5%
#2	fable, opus, deepseek	BTC+ETH	4h/1h	+\$135.4	+13.5%	0.2%	23	52.2%
#	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	1/0	reenter	\$450

Against the 3,952-cell stage-2 grid on this window: card #1 lands in the +\$500 to +\$750 bucket (10 of 3,952 cells) and card #2 in the \$0 to +\$250 bucket (1755 cells); 10 cells in the whole grid finished above +\$500 (grid median -\$32.2). Card percentiles exported this issue cover the raw net-board leaders and the replay rows, not the mature cards -- no percentile is claimed for the two cards themselves. 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.


Only two councils are published on this window. All ten slots of the mature net board are the same 2-coin universe (SOL+XRP), so the pairwise-different-universes rule admits exactly one; #2 is the best mature multi-coin council outside that universe on the quality boards -- it comes off the **Calmar** board (+\$135.4 on BTC+ETH, R2 0.795, ulcer 0.06, 0.2% maxDD, 0 unfunded entries) -- disclosed rather than padded. Read the cards as two universes deep, not three.

On the calibrate axis 1 of the 2 cards are calibrated cells, where issues #4 and #5 published WK councils that were all calibrated cells. Card #1 runs the 2-coin table stake of \$450. Neither card carries an unfunded entry.

OPEN IN SANDBOX marketmania.ai/s/cw6-wk1 marketmania.ai/s/cw6-wk2

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

Published cards equity -- WK (Sep 7-14, 2026) -- \$1,000 at 10x -- 41 settled trades



Daily settled equity for card #1 above (\$1,000 start, 10x leverage, frozen window; #1 red -- fixed rank colors across all issues). A curve running past the window end is an open-at-cutoff trade settling at its horizon (the 2-3 day left); the curve's last point reconciles to its config's net to the cent. There is no contrast line on this window: the net-board maximum is the same simulation as card #1 (i=1859, 41 closed, above the >=15 maturity gate), so a second curve would be the same curve. Card #2 has no daily series in this issue's equity pack, and the solo and smoothness series stay in the open data (equity.cards) -- the PDF carries one equity chart per window. Source: equity_wk.json.

TOP councils -- M30 window (Aug 17-Sep 14, 2026)

#	Council	Coins	FH/TF	Net	Net %	maxDD	Trades	WR
#1	opus, deepseek, gemini	SOL+XRP	4h/1h	+\$1,793.8	+179.4%	52.8%	168 *	64.9%
#2	gpt, deepseek, gemini, grok	BTC+ETH	1d/1d	+\$547.8	+54.8%	11.2%	38	89.5%
#	Ladder	BE	Shifts	Conf filter	Source	Entry	Positional	Stake
#1	50:40,100:60	on	TP -10 / SL +75	>=60 (each)	farthest	2/1	reenter	\$450
#2	50:40,100:60	on	TP -10 / SL +75	>=60 (each)	farthest	3/1	reenter	\$450


Against the 4,080-cell stage-2 grid on this window: card #1 is the same simulation as the raw net-board leader (i=1466), so its exported percentile is the card's -- above **100.0%** of the grid; card #2 lands in the +\$500 to +\$750 bucket (115 cells) with 95.9% of the grid finished below it. Grid median -\$0.1, 167 of the 4,080 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.

Only two councils are published on this window too, and it is the same universe the week collapsed onto. 10 of the ten slots of the M30 mature net board are 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 (off the **Smoothness** board: +\$547.8 on BTC+ETH, R2 0.911, 11.2% maxDD, 0 unfunded entries). Read the cards as two universes deep, not three.


On the calibrate axis 1 of the 2 M30 cards are calibrated cells, where 1 of the 2 published WK councils are calibrated cells -- the same axis, read on two windows (Calibrate axis below). Card #1 pays for its +\$1,793.8 with a **52.8% drawdown** on 168 closed trades and 1 unfunded entry, card #2 with 11.2% on 38 and 0. The M30 mature pool is 3,265 configs against 1,781 on WK, because four weeks clear the >=20-trade gate that one week does not.

OPEN IN SANDBOX marketmania.ai/s/cw6-m301 marketmania.ai/s/cw6-m302

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



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); its last point reconciles to the card's net to the cent. There is no contrast line on this window: the net-board maximum is the same simulation as card #1 (i=1466, 168 closed, above the ≥ 20 maturity gate). The solo, Calmar, WR and smoothness series stay in the open data (equity.cards) -- the PDF carries one equity chart per window. Source: equity_night.json.




Net PnL across all 4,080 stage-2 grid sims, M30 window (Aug 17-Sep 14, 2026), universe-aware stake throughout. Grid median **-\$0.1**, 30.3% positive against 17.5% on the week (issue #5's month: median \$0.0, 47.6% positive). The spike at \$0 is mostly cells whose entry filters produced no trades (802 zero-trade sims on this window). Range -\$869.2 to +\$1,793.8; 167 cells finished above +\$500, against 10 on the week. Red marker = best solo default at table stake; green = the published cards. Source: grid_hist_night.json.

Walk-forward: the standing OOS verdict

Every config this series has ever published -- all 38 cards from issues #1, #2, #3, #4, #5 and Monthly Config Watch #1 -- replayed VERBATIM (same knobs, same stake as printed, no re-tuning) on both fresh windows, plus the four hand-tuned presets running live on the replatform, listed first. The monthly issue's three cards are judged by the weekly rule, like every other row. The replays and their origin re-runs cost 68 simulations in the weekly collect and 110 in the night collect, which replays every row on both of its own windows.

Verdict rule, unchanged since issue #2. Survived = net-positive on BOTH fresh windows; **Faded** = net-negative on at least one. The live presets and the monthly cards are judged by the same rule. New WK is pure out-of-sample for every row -- the week begins exactly at issue #5's right edge (Sep 7). New M30 (Aug 17-Sep 14, 2026) shares 21 of its 28 days with issue #5's own M30 window (the shared stretch runs Aug 17-Sep 6), 14 days with issue #4's M30 window and 15 days with Monthly Config Watch #1's calendar-August window (shared stretch Aug 17-31), so for those three cohorts that column is decay context rather than pure out-of-sample: the same asymmetry issues #3, #4 and #5 disclosed for their own second columns.

Issue #5's four published cards: home number vs fresh-week replay



Issue #5's 4 published cards: the number its own issue printed (grey, in-sample on its own window) against the same configuration replayed verbatim on this week (green, out-of-sample). 3 of the 4 are net-negative on the fresh week; the M30 card was measured at home on a four-week window and the WK cards on a one-week window, so the grey bars are not all the same kind of number. Source: oos_cw5_wk.json + issue #5 open data.

Config	Family	Coins	Home	New WK (n)	New M30 (n)	Verdict
preset 4h4	fable, qwen, deepseek, opus	ETH+SOL	-- live	+\$63.4 (12)	+\$66.3 (49)	Survived
preset bc2	gemin, gpt, deepseek	SOL+XRP	-- live	-\$74.3 (4) * +	+\$431.6 (24)	Faded
preset sol	deepseek, fable, opus	SOL [conc]	-- live	+\$551.0 (11)	+\$2,159.0 (45)	Survived
preset xrp	gemin, gpt, deepseek	XRP [conc]	-- live	+\$199.2 (2) *	-\$171.0 (1) * +	Faded
cw5-m301	fable, opus, gpt, gemini	SOL+XRP	+\$1,717.1	-\$578.3 (5) * +	+\$489.3 (32)	Faded
cw5-s1	gemin	BNB+XRP	+\$178.7	-\$125.9 (33) +	+\$83.1 (172) +	Faded
cw5-wk1	gpt, gemini	SOL+XRP	+\$510.4	+\$36.3 (50) +	-\$578.6 (51) +	Faded
cw5-wk2	fable, opus, gpt, gemini, grok	BTC+ETH	+\$183.5	-\$24.9 (18) +	-\$202.7 (85) +	Faded
cw4-m301	gpt, deepseek, gemini	SOL+XRP	+\$1,676.0	-\$322.9 (5) * +	+\$840.9 (37)	Faded
cw4-m302	fable, gpt, deepseek, gemini	BNB+XRP	+\$1,155.2	-\$446.8 (7) *	+\$62.2 (27)	Faded
cw4-s1	deepseek	BTC+ETH+ SOL+BNB+XRP	+\$325.0	-\$59.7 (14)	-\$66.5 (70) +	Faded
cw4-wk1	opus, deepseek, gemini, qwen	BTC+ETH+ SOL+BNB+XRP	+\$399.4	-\$47.0 (17) +	+\$299.5 (71)	Faded
cw4-wk2	fable, deepseek, qwen	BNB+XRP	+\$128.2	-\$66.8 (29)	-\$51.8 (96) +	Faded
cwm1-m1	fable, gemini	SOL+XRP	+\$1,618.4	-\$534.9 (6) * +	+\$807.6 (35)	Faded
cwm1-m2	fable, gpt, deepseek, gemini, qwen	ETH+SOL+ BNB+XRP	+\$756.6	-\$263.5 (10)	+\$268.3 (45) +	Faded
cwm1-s1	gemin	SOL+XRP	+\$1,721.9	-\$683.4 (9) * +	+\$598.8 (46)	Faded
cw3-m301	opus, deepseek, gemini	SOL+XRP	+\$2,137.6	+\$130.5 (28) +	+\$1,225.7 (115)	Survived
cw3-m302	fable, opus, deepseek, gemini, qwen	BTC+ETH	+\$412.3	-\$297.0 (4) *	+\$38.1 (22)	Faded
cw3-s1	gemin	SOL+XRP	+\$1,604.3	-\$220.1 (32) +	-\$116.7 (165) +	Faded
cw3-s2	gemin	SOL+XRP	+\$1,449.0	-\$683.4 (9) * +	+\$598.8 (46)	Faded
cw3-wk1	opus, deepseek, gemini	SOL+XRP	+\$2,175.1	+\$130.5 (28) +	+\$1,225.7 (115)	Survived
cw3-wk2	opus, deepseek, gemini, grok	BTC+ETH	+\$837.6	-\$288.4 (16) +	-\$411.4 (67) +	Faded

Config	Family	Coins	Home	New WK (n)	New M30 (n)	Verdict
cw2-30d1	opus, gpt, deepseek, gemini, qwen	BTC+ETH+ SOL+BNB+XRP [conc]	+\$744.7	-\$569.6 (5) * +	-\$115.6 (49) +	Faded
cw2-30d2	opus, deepseek, gemini	ETH+SOL+ BNB+XRP [conc]	+\$655.1	-\$399.0 (8) * +	+\$456.5 (63) +	Faded
cw2-30d3	opus, deepseek, gemini	BNB+XRP	+\$489.1	-\$440.9 (8) * +	-\$236.3 (27) +	Faded
cw2-30ds1	gemini	BNB+XRP	+\$383.8	-\$746.1 (6) * +	-\$626.0 (34) +	Faded
cw2-7d1	fable, opus, deepseek, gemini	SOL+XRP	+\$299.1	+\$326.0 (24) +	+\$910.5 (108)	Survived
cw2-7d2	gpt, deepseek, gemini, qwen	BNB+XRP	+\$223.7	-\$457.3 (7) * +	-\$297.3 (25) +	Faded
cw2-pt1	fable, deepseek, grok [remap]	SOL	+\$698.3	+\$138.6 (16)	+\$914.7 (59)	Survived
cw2-pt2	opus	XRP	+\$543.7	-\$396.3 (3) *	-\$720.6 (12) +	Faded
cw1-30d1	fable, deepseek, opus, qwen	ETH+SOL	+\$868.4	+\$98.2 (12)	+\$89.9 (51)	Survived
cw1-30d2	gemini, gpt, deepseek	SOL+XRP+ETH	+\$590.6	-\$275.9 (8) * +	+\$396.3 (41)	Faded
cw1-30d3	fable, deepseek, opus, qwen	BTC+ETH+SOL	+\$580.5	-\$13.6 (20)	-\$112.3 (77) +	Faded
cw1-30ds1	opus	SOL+XRP	+\$442.8	-\$161.7 (26) +	+\$543.3 (116)	Faded
cw1-30ds2	grok [remap]	SOL+XRP	+\$262.9	-\$0.1 (26) +	-\$4.6 (83) +	Faded
cw1-30ds3	deepseek	XRP+BNB	+\$119.0	-\$44.9 (31) +	-\$107.4 (123) +	Faded
cw1-7d1	fable, deepseek, opus, gpt	SOL+XRP	+\$439.2	+\$401.3 (33) +	+\$1,695.3 (138)	Survived
cw1-7d2	fable, qwen, gemini, opus	XRP+BNB	+\$429.5	-\$281.2 (20) +	-\$266.9 (84) +	Faded
cw1-7d3	gpt, opus	BNB+XRP+SOL	+\$398.6	-\$186.5 (39) +	-\$159.3 (166) +	Faded
cw1-7ds1	gemini	XRP+BNB	+\$417.8	+\$74.2 (50) +	-\$588.8 (97) +	Faded
cw1-7ds2	gpt	XRP+BNB	+\$366.3	-\$236.0 (21) +	-\$122.2 (126) +	Faded
cw1-7ds3	qwen	XRP+BNB	+\$267.0	-\$228.9 (23) +	-\$373.1 (97) +	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 both fresh windows, Faded = net-negative on at least one. Survivors are tinted green.

Counted: 6 of 38 past cards **Survive** both fresh windows (issue #5 on its two windows: 6 of 34) -- by cohort 0 of 4 issue-5 cards, 0 of 5 issue-4 cards, 0 of 3 Monthly-1 cards, 2 of 6 issue-3 cards, 2 of 8 issue-2 cards, 2 of 12 issue-1 cards. Taken a window at a time: 8 of 38 are positive on the week and 19 of 38 on the four weeks, and the two columns do not nest: a row can clear one and fail the other, so both windows remove cards. The survival curve by generation, each on its own first, second, third, fourth and fifth fresh week, reads: issue-1 50.0% -> 91.7% -> 8.3% -> 16.7% -> **25.0%** now (3 of 12); issue-2 100.0% -> 50.0% -> 12.5% -> **25.0%** now (2 of 8); issue-3 33.3% -> 33.3% -> **33.3%** now (2 of 6); issue-4 20.0% -> **0.0%** now (0 of 5); Monthly-1 66.7% -> **0.0%** now (0 of 3); issue-5 **25.0%** now (1 of 4). Live presets: 2 of 4 Survive, 3 of 4 positive on the week, 3 of 4 on the month.

The four live presets lead the table by standing rule. On the week: preset 4h4 +\$63.4 on 12 closed, preset bc2 -\$74.3 on 4 closed, preset sol +\$551.0 on 11 closed, preset xrp +\$199.2 on 2 closed -- 3 of the four net-positive. On the four-week window: preset 4h4 +\$66.3 on 49 closed, preset bc2 +\$431.6 on 24 closed, preset sol +\$2,159.0 on 45 closed, preset xrp -\$171.0 on 1 closed -- 3 of four. 2 Survive both. **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. Passports for all four are in the open data (presets_recon_wk.json and presets_recon_night.json).

Origin re-runs (drift baseline). Re-simulating the issue-5 cards on their OWN origin windows today lands within -\$497.3...-\$98.9 of the printed issue-5 numbers and the issue-4 cards within -\$211.2...+\$60.3 of theirs; 1 of the 9 re-runs reproduce to the cent. The

widest deviation is cw5-m301 (+\$1,717.1 printed, +\$1,219.8 today) -- today's engine reads a forecast history that keeps growing, which is the one thing a verbatim replay cannot freeze. Snapshot principle: no past issue is restated; the full drift table is in the open data (oos_walk_forward.drift_baseline).

Honest read. The fresh search and its own back catalogue moved in opposite directions this week. On the fresh week 30 of 38 past cards are net-negative -- issue #5 counted 26 of 34 on its own week -- while the ceiling of the fresh search rose to +\$609.5 from +\$510.4. On the fresh four weeks 19 of 38 are net-positive. Same replay machinery, same frozen knobs; what changed is the window and the size of the board. Three things keep this from being a statement about search. The board is not a fixed sample: it grows every issue and nothing is ever pruned, so a share computed on 38 cards is not comparable with one computed on 34 without saying which generations joined -- 9 of the 38 rows are new this issue. The month column is not clean evidence for three of the six cohorts: it shares 14 of its 28 days with issue #5's own M30 window and 15 with Monthly Config Watch #1's calendar-August window, so those cards are partly being scored on their own data. And the size of the in-sample number is not what decides the give-back on this board: the row that lost most on the week is cw2-30ds1 (-\$746.1 on the week against +\$383.8 printed at home), while the largest home number on the board, cw3-wk1 at +\$2,175.1, gave back +\$130.5. Five fresh weeks of verdicts now: brutal, the opposite, split by window, the fourth, and this one. That is a series of weathers, not a strategy.

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

Matched pairs -- identical knobs, calibration OFF vs ON -- measured on net PnL. The table is the **whole stage-2 grid** of each window, and the **Both** row pools the two windows' pairs and nothing else. The board-leader slice is a second, separate measurement (14 WK pairs and 13 M30 pairs whose OFF or ON leg reached the top of a board); it is reported in the paragraph below and in the open data, and the two are never pooled together.

Window	Pairs	Calibrate wins	Mean delta	Best delta	Worst delta
WK	1,334	49.9%	+\$60.3	+\$862.2	-\$706.7
M30	1,392	48.4%	+\$24.4	+\$950.1	-\$1,354.9
Both	2,726	49.1%	+\$42.0	+\$950.1	-\$1,354.9

The two windows answer the axis, and the whole-grid reading and the board-leader slice are separate measurements. On the week calibration improves 49.9% of the 1,334 matched pairs (mean +\$60.3, median \$0.0) and 51.0% of the 453 mature pairs (mean -\$9.6); on the four weeks it improves 48.4% of 1,392 (mean +\$24.4, median \$0.0) and 59.5% of the 879 mature pairs (mean +\$49.7). Pooled over both windows the axis improves 49.1% of 2,726 pairs with a mean of +\$42.0 -- a number that describes neither window on its own. The board-leader slices are separate: 0.0% of the 14 WK slice pairs improve (mean -\$568.6, median -\$582.1, OFF legs already at a median of +\$554.7), and 15.4% of the 13 M30 slice pairs (mean -\$670.5, OFF legs at +\$1,469.7). Issue #5 measured a 12-pair WK slice (8 improving) and a 14-pair M30 slice (0). Where the axis reached the cards is visible above: 1 of the 2 published WK councils are calibrate=ON cells and 21 of the 54 WK finalists run it, while 1 of the 2 M30 cards are and 34 of the 53 M30 finalists do.

The open-data pack also carries a 400-pair export (calibrate_effect) per collect. Its delta distribution is not the population's (median +\$246.3 against the WK grid's \$0.0), 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 (>=15 closed on WK, >=20 on M30), one config per model.

Solo top -- WK (Sep 7-14, 2026)

#	Model	Coins	FH/TF	Net	Net %	maxDD	Trades	WR
#1	grok	SOL+XRP	4h/4h,1h	+\$404.8	+40.5%	4.6%	31	80.6%
#2	gpt	SOL+XRP	4h/4h,1h	+\$314.8	+31.5%	9.7%	16	87.5%
#	Ladder	BE	Shifts	Conf filter	Source	Entry	Positional	Stake
#1	50:40,100:60	on	TP -10 / SL +75	>=60 (each)	median	1/0	reenter	\$450
#2	50:40,100:60	on	TP -10 / SL +75	>=70 (each)	median	1/0	reenter	\$450

2 solo cards on this window. Both qualifying solos run the 2-coin table stake of \$450. No single-coin row outranks them on the mature net board this issue; single-coin results are excluded from the cards by the >=2-coin rule either way and appear under

Per-ticker bests. Source: summary_wk.json top_mature_by_window_branch.

Solo top -- M30 (Aug 17-Sep 14, 2026)

#	Model	Coins	FH/TF	Net	Net %	maxDD	Trades	WR
#1	opus	SOL+XRP	4h/1h	+\$955.5	+95.6%	22.5%	122	91.0%
#	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

One solo card, not two on M30. The mature solo board collapses onto opus (SOL+XRP), so the one-config-per-model rule admits exactly one; the card carries 0 unfunded entries. Source: top_mature_by_window_branch in the night collect.

OPEN IN SANDBOX marketmania.ai/s/cw6-s1

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

Council vs solo. On the week the best solo sits below the best council (+\$404.8 vs +\$609.5 on the mature board, 31 closed against 41, 4.6% drawdown against 6.4%). On the four weeks the best solo sits below the best council (+\$955.5 vs +\$1,793.8, 122 closed against 168, 22.5% against 52.8%). Same in-sample caveats apply to both sides.

Quality boards

Net PnL stays the primary board (continuity with issues #1-#4); 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 (>=15 closed on WK, >=20 on M30); **Smoothness** = R2 of a linear fit through the daily closed-equity series, mature rows with net > 0 and at least 5 days, ties broken by ulcer index; **Quality composite** = mean rank over net, Calmar-like, win rate and R2, lower is better. Top-3 of each board on each window below; all 15 rows of all eight boards are in the open-data JSON.

Board	Config	Coins	FH/TF	Net	WR	maxDD	Cls	R2
WK Calmar	opus, grok	SOL	4h/1h	+\$345.4	87.5%	0.0%	16	0.957
WK Calmar	opus, grok	SOL	4h/1h	+\$345.4	87.5%	0.0%	16	0.957
WK Calmar	fable, opus, deepseek [cal]	BTC+ETH	4h/1h	+\$135.4	52.2%	0.2%	23	0.795
WK WR	opus, grok [cal]	SOL	4h/1h	+\$119.3	93.8%	0.0%	16	0.918
WK WR	opus, grok [cal]	SOL	4h/1h	+\$119.3	93.8%	0.0%	16	0.918
WK WR	fable, opus [cal]	SOL+XRP	4h/4h,1h	+\$147.1	92.6%	10.7%	27	0.448
WK Smoothness	opus, grok	SOL	4h/1h	+\$345.4	87.5%	0.0%	16	0.957
WK Smoothness	opus, grok	SOL	4h/1h	+\$345.4	87.5%	0.0%	16	0.957
WK Smoothness	opus, grok [cal]	BTC	4h/1h	+\$14.1	53.3%	5.2%	15	0.949
WK Quality	opus, grok	SOL	4h/1h	+\$345.4	87.5%	0.0%	16	0.957
WK Quality	opus, grok	SOL	4h/1h	+\$345.4	87.5%	0.0%	16	0.957
WK Quality	fable, opus, gpt, grok	SOL+XRP	4h/1h	+\$593.2	82.9%	6.4%	35	0.909
M30 Calmar	fable, deepseek, qwen [cal]	SOL	4h/4h,1h	+\$294.9	100.0%	0.0%	22	0.941
M30 Calmar	fable, deepseek, qwen [cal]	SOL	4h/4h,1h	+\$245.7	100.0%	0.0%	22	0.941
M30 Calmar	fable, deepseek, qwen [cal]	SOL	4h/4h,1h	+\$196.6	100.0%	0.0%	22	0.941
M30 WR	fable, deepseek, qwen [cal]	SOL	4h/4h,1h	+\$294.9	100.0%	0.0%	22	0.941

Board	Config	Coins	FH/TF	Net	WR	maxDD	CIs	R2
M30 WR	fable, deepseek, qwen [cal]	SOL	4h/4h,1h	+\$245.7	100.0%	0.0%	22	0.941
M30 WR	fable, deepseek, qwen [cal]	SOL	4h/4h,1h	+\$196.6	100.0%	0.0%	22	0.941
M30 Smoothness	opus, deepseek, gemini [cal]	SOL	4h/1h	+\$711.4	89.3%	15.0%	84	0.967
M30 Smoothness	gpt, deepseek, gemini, grok [cal]	SOL	4h/1h	+\$462.7	90.4%	5.7%	73	0.965
M30 Smoothness	gpt, deepseek, gemini, grok [cal]	SOL	4h/1h	+\$578.4	90.4%	7.1%	73	0.965
M30 Quality	gpt, deepseek, gemini, grok [cal]	SOL	4h/1h	+\$694.1	90.4%	8.5%	73	0.965
M30 Quality	opus [cal]	SOL	4h/1h	+\$784.4	96.7%	9.9%	61	0.846
M30 Quality	opus [cal]	SOL	4h/1h	+\$784.4	96.7%	9.9%	61	0.846

[cal] = calibrate=ON cell; the board column names the window. Both Calmar boards are degenerate by construction -- their top rows sit at 0.0% (WK) and 0.0% (M30) maxDD, so the score collapses onto net PnL; read them as 'net among configs that barely drew down'. The Smoothness boards are the ones that can leave a window's dominant universe: the WK board's top row is a SOL config at R2 0.957 on a net of +\$345.4, the M30 board's a SOL config at R2 0.967 on +\$711.4. That trade -- straightness bought with size -- is exactly what the board exists to expose. The M30 boards are drawn from a mature pool of 3,265 configs against 1,781 on WK, because four weeks clear a >=20-trade gate that one week does not.

Stability, and what it is across this issue. The board is defined as **mature net top-40 on BOTH windows** of the run that computes it. The run that produced it here is the night collect, whose two windows are **WK and M30** -- the same two windows this issue prints, computed in a second execution of the WK plan -- so this issue's board reads **stable across WK and M30. 9 configurations qualify**. All of them SOL+XRP, 9 councils and 1 clean (zero unfunded entries on both windows); the top three by rank sum are council fable, opus, deepseek, gemini; council fable, opus, deepseek, gemini; council fable, gpt, deepseek, gemini. The shared board is narrower than either window's own: 19 of the 20 rows exported from its WK board are SOL+XRP; 14 of the 20 rows exported from its M30 board are SOL+XRP, and no row exported from one window's mature board appears on the other's. The count is not comparable with issue #4's 25 (its board was measured across M30 and the calendar month); the boards measured across WK and M30 were issue #5's, at 0, and issue #3's, at 21. Full list in the open data (stability).

Hand-tuned baseline vs the grid

Standing section, third issue. 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 (referred to below as 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 each window

Window	Config	Coins	FH/TF	Net	WR	maxDD	R2	Ulcer	CIs
WK	Hand-tuned anchor (1d_3_llm_BC2_v02)	SOL+XRP	1d/1d	-\$74.3	75.0%	16.6%	n/a	0.00	4
M30	Hand-tuned anchor (1d_3_llm_BC2_v02)	SOL+XRP	1d/1d	+\$431.6	70.8%	56.0%	0.232	13.58	24

0 of the 1,781 mature WK configs (0.0%) clear all three bars at once, and 465 of the 3,265 mature M30 configs (14.2%). The anchor's own two windows are different rows: on the week -\$74.3 net at 75.0% win rate with a 16.6% drawdown, an R2 of n/a and 4 closed trades; on the four weeks +\$431.6 at 70.8%, a 56.0% drawdown, an R2 of 0.232 and 24 closed. The week row is thin against the >=15 gate the search rows must pass (the anchor is exempt from that gate by construction: it is the reference, not a candidate). Read against net PnL, no WK config clears all three bars, so there is no net comparison to make on that window; all 10 of the 10 best M30 winners by net also beat it on net PnL. The share itself moves with the window (0.0% of the WK field, 14.2% of the M30 field),

so read it as one field at a time, not as a property of the anchor.

No WK table here: 0 of the 1,781 mature WK configs clear all three bars at once, so there is no top-10 to rank. The anchor's own WK row carries a smoothness R2 of n/a on 4 closed trades, and the smoothness bar cannot be cleared against a row that has none.

Top-5 by net among the 465 M30 configs that beat the anchor on win rate, maxDD and smoothness at once

#	Config	Coins	FH/TF	Stake	Net	WR	maxDD	R2
#1	fable, opus, deepseek	SOL	4h/1h	\$450	+\$1,153.3	78.7%	11.5%	0.630
#2	opus, deepseek	SOL	4h/1h	\$450	+\$1,132.4	75.4%	23.5%	0.466
#3	fable, gpt, deepseek, gemini [cal]	SOL+XRP	1d/1d	\$450	+\$1,124.0	84.8%	20.1%	0.810
#4	fable, opus	SOL	4h/1h	\$450	+\$1,069.9	79.0%	25.9%	0.436
#5	fable, gpt, deepseek, gemini [cal]	SOL+XRP	1d/1d	\$450	+\$1,063.4	84.4%	20.1%	0.815

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 this window, 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 560 + 1,150 sims -- a compacted grid, one best per coin per window; the branch tag in brackets says which kind won.

WK (Sep 7-14, 2026)

Coin	Best config (branch)	FH/TF	Net (trades)
BTC	gpt [solo]	4h/4h,1h	+\$103.0 (10)
ETH	deepseek [solo]	1d/1d,4h	+\$174.1 (5) *
SOL	opus, gpt, grok [council]	4h/1h	+\$360.0 (19)
BNB	opus, deepseek [council]	4h/1h	+\$136.4 (17)
XRP	fable, deepseek, gemini [council]	4h/1h	+\$281.1 (21)

M30 (Aug 17-Sep 14, 2026)

Coin	Best config (branch)	FH/TF	Net (trades)
BTC	opus, grok [council]	1d/1d	+\$319.3 (9) *
ETH	deepseek [solo]	1d/1d	+\$570.6 (21)
SOL	fable, opus, deepseek [council]	4h/1h	+\$1,153.3 (61)
BNB	opus, deepseek, qwen [council]	4h/4h	+\$440.8 (46)
XRP	opus, deepseek, gemini [council]	4h/1h	+\$780.7 (84)

n in parentheses; * = fewer than 10 closed trades (thin -- direction only). 0 of the ten winners run calibrate ON; 3 of the five WK winners and 4 of the five M30 winners are councils.

The reads. On the week SOL leads at +\$360.0 and 1 of the five rows are thin -- 5 to 21 closed trades against a 15-trade gate. On the four weeks 1 of the five rows are thin: 9 to 84 closed, SOL leading at +\$1,153.3, and the same coin leads both windows. The 1d horizon wins 1 of the five WK coins and 2 of the five M30 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 next issue.

Baselines

Three independent reference points, none of them search output, re-simulated fresh on both of this issue's frozen windows inside the same two collects (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 issue #1. The tables show the universe-aware **table-stake** run -- this issue's canonical mode (\$200 for the 5-coin defaults; A and B already at their \$300 3-coin cap; C moves \$400 -> \$450 on 2 coins) -- with the best and worst of the seven solo defaults on each window; all eleven rows per window and both stake modes are in the open data.

WK (Sep 7-14, 2026)

Baseline	Net PnL	Trades	WR	maxDD
Best solo default (grok-4.6)	-\$28.7 *	204	61%	3.9%
Worst solo default (claude-fable-5)	-\$92.9 *	247	52%	9.3%
Council-of-7 default	-\$16.2 *	137	63%	2.9%
Hand-tuned config A (pre-search)	-\$702.1 *	7	14%	74.7%
Hand-tuned config B (pre-search)	-\$81.6 *	39	67%	23.4%
Hand-tuned config C (pre-search)	+\$98.2	12	75%	11.5%

M30 (Aug 17-Sep 14, 2026)

Baseline	Net PnL	Trades	WR	maxDD
Best solo default (claude-opus-5)	-\$62.4 *	775	62%	13.1%
Worst solo default (deepseek-v4-pro)	-\$197.1 *	1203	59%	19.7%
Council-of-7 default	-\$57.4 *	513	61%	10.8%
Hand-tuned config A (pre-search)	-\$355.3	31	61%	101.1%
Hand-tuned config B (pre-search)	-\$75.6 *	172	60%	90.2%
Hand-tuned config C (pre-search)	+\$89.9	51	65%	68.5%

* = non-zero skipped_nofunds at table stake. On the 5-coin defaults those counts are the stake table doing its job: at \$200 a \$1,000 deposit funds at most five concurrent positions while the defaults fire hundreds of entries per window (135 to 317 skipped each on the week, 458 to 1,194 on the four weeks). The other five solo defaults run between the two printed rows on each window. **The conclusions do not change between modes:** on WK they do so for every row, and on M30 for every row but one, printed as it stands: solo default gemini-3.1-pro is +\$35.6 at the config's own stake and -\$102.2 at the table stake. 2 of the three hand-tuned configs are already at their own table cap, so their two modes are identical on both windows. Source: baselines_wk.json (22 sims) and baselines_night.json (44 sims, of which 22 are the M30 half), inside this issue's two collects.

Every engine default lost money on both windows; the hand-tuned set did not. On WK the seven solo defaults run -\$28.7 to -\$92.9 at table stake, the council-of-7 -\$16.2, and the three hand-tuned configs A -\$702.1, B -\$81.6, C +\$98.2 -- 1 of the three positive. On M30 every default is negative again (-\$62.4 to -\$197.1 solo, council-of-7 -\$57.4) and the three hand-tuned configs A -\$355.3, B -\$75.6, C +\$89.9 -- 1 of the three positive. The upper half of the standing ordering (search > hand-tuned > defaults) holds on both windows, in-sample as ever -- the fresh search makes +\$609.5 on the week against a best baseline of +\$98.2, and +\$1,793.8 on the four weeks against +\$89.9. The live presets are a separate row of evidence: 3 of the four are net-positive on the week and 3 of four on the month (walk-forward table above), and 2 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 (presets_recon_wk.json and presets_recon_night.json). 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 per window); the two single-coin presets do not (0 cells -- the local grid is not built for them, so no neighbour claim is made about either).

Live preset	Window	As configured	Best neighbour in local grid	Cells	Verdict
preset bc2	WK	-\$74.3	+\$83.4 (1d/1d, SOL+XRP+ETH, 7 cls)	34	neighbour +\$157.7
preset bc2	M30	+\$431.6	+\$617.7 (1d/1d, SOL+XRP, 32 cls)	34	neighbour +\$186.1
preset 4h4	WK	+\$63.4	+\$102.8 (4h/1h, SOL+XRP, 12 cls)	19	neighbour +\$39.5
preset 4h4	M30	+\$66.3	+\$632.7 (4h/1h, ETH+SOL, 89 cls)	19	neighbour +\$566.4

In **4 of the 4** preset-window cells with a grid, a one-knob neighbour beat the configuration as it is actually running. The largest single edit is preset 4h4's 4h/1h -> 4h/1h horizon move on M30, which turns +\$66.3 into +\$632.7. The presets are not at a local optimum -- a testable, low-risk edit, unlike adopting a search champion wholesale. Cell counts are small (34 and 19 per window) and every neighbour is an in-sample maximum of its own little grid.

Patterns: what the finalists look like

Knob modes across each window's unique finalists (mature net top-10 per branch plus all four quality boards, deduplicated by full signature; 54 unique configs on WK, 44 of them councils; 53 on M30, 43 councils).

Knob	WK finalists mode	M30 finalists mode
Forecast horizon	4h (54/54)	4h (44/53)
Ladder shape (steps)	50:40,100:60 (54/54)	50:40,100:60 (53/53)
Break-even stop (BE)	on (54/54)	on (53/53)
Min confidence	60 (53/54)	60 (53/53)
SL shift	75 (54/54)	75 (53/53)
TP shift	-10 (54/54)	-10 (53/53)
TP/SL source	farthest (47/54)	farthest (48/53)
Calibrate ON	21/54	34/53
Universe	SOL+XRP (41/54)	SOL+XRP (27/53)

The mechanical core is the same on both windows and unchanged for the sixth issue running: 50:40,100:60 ladder, break-even on, conf 60, SL +75 / TP -10, farthest source. The horizon is **4h** on WK (54/54) and **4h** on M30 (44/53). **Calibrate=ON** runs in 21 of 54 WK finalists against 34 of 53 on M30, where issue #5 measured 52/62 on its week and 27/47 on its month. The universe: SOL+XRP takes 41 of the 54 WK finalists while SOL+XRP takes 27 of the 53 M30 finalists. Membership concentrates too -- opus leads the WK council finalists (42 of 44) and deepseek the M30 ones (38 of 43). Standing caveats: refinement seeds from the same grid winners, so part of the convergence is search-design echo; and a knob core that produced 6-of-34 survival one issue ago and 6-of-38 this issue is describing the weather, 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. This issue the published cards draw down 6.4% and 0.2% on the week and 52.8% and 11.2% on the month, and the replay table is deeper: 12 replayed rows drew down more than 40% on the week and 35 on the month, 5 and 32 of them 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 week.

Watch amendments (this issue)

Amendment 3 (standing) -- smoothness and a quality composite. Every simulation stores its daily closed-equity series and the three numbers read off it (n_days , $smooth_r2$, $ulcer$); the smoothness board ranks mature rows with $net > 0$ and at least 5 days by R2 with ulcer as tie-break, and the quality composite ranks by the mean of the net, Calmar-like, win-rate and R2 ranks. Both still cost zero additional simulations and the headline board is still net PnL.

Amendment 4 (standing) -- the monthly split. The weekly Config Watch keeps two windows, the calendar week (WK) and the four calendar weeks that end at the same edge (M30, Aug 17-Sep 14, 2026); the calendar-month board belongs to the Monthly Config Watch line and no calendar-month number is printed in a weekly issue. This issue's night collect computes no calendar-month window at all: Monthly Config Watch #2 covers September and is produced after 01.10.

Amendment 5 (standing) -- the standing OOS board grows with every issue. Every card this series publishes joins the walk-forward table and is never removed: this issue replays **38 cards** -- issue #5's 4 join issue #4's 5, Monthly Config Watch #1's 3 and issues #1, #2 and #3's 26 -- plus the four live presets, on both fresh windows. **Monthly cards are judged by the weekly rule** (net-positive on both of this issue's fresh windows), the same rule every other row is judged by; their own calendar-month origin window is the Home column and nothing else. One consequence for the counts: a survival share is computed on a board that changes size every issue, so it is reported with the cohort breakdown beside it and never as a trend on its own. **Stability** is defined inside the run that computes it: this issue's night collect computes WK and M30, so the board reads across WK and M30 -- the issue-#3 definition, the one issue #5 used as well -- where issue #4's read across M30 and the calendar month.

All three amendments are standing, 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

- **Two windows, one regime, and they overlap.** M30 (Aug 17-Sep 14, 2026) contains WK (Sep 7-14, 2026) -- 7 of the week's 7 days are inside it -- and it shares 21 of its 28 days with issue #5's own M30 window, 14 with issue #4's and 15 with Monthly Config Watch #1's calendar-August window. The two columns of this issue are therefore not two independent tests: agreement between them is partly arithmetic, and for the issue-5, issue-4 and Monthly-1 cohorts the M30 replay column is decay context rather than pure out-of-sample. Do not read them as a cross-validation.
- **The two windows come from two collects, and the two runs of the same week do not agree.** WK is the weekly run (generated 2026-09-15T20:56:34Z), M30 the night run (2026-09-16T09:13:44Z -- boards computed at 08:58:12Z, files written after the baselines pass), which recomputed the same WK week in the same pass on its own grid: the night run's WK ceiling is +\$534.6 (fable, opus, deepseek, gemini on SOL+XRP) against the weekly collect's +\$609.5. Every WK number this issue prints is the weekly collect's; the night run's WK half is used only for the cross-window stability board. Same engine, same criteria, same collect code, but not the same execution -- a number is comparable across the two windows only as far as that is.
- **Stability is defined inside the run that computes it.** The board comes from the night run, whose two windows are WK and M30, so it reads across those two -- and 9 configurations are in the mature net top-40 of both. Issue #5 read the same pair one week earlier and held 0, and issue #3's board (the other one across the same pair) held 21; issue #4's 25 is not comparable at all, having been measured across M30 and the calendar month. A board this size is one observation per issue and no claim is made from a single reading.
- **Model lineage is a splice.** The grok axis runs as grok-4.6 (incl. 4.5-era) -- one lineage across a mid-series model cutover dated 2026-08-24 in the search plan both collects ran to. This issue's WK window (Sep 7-14, 2026) lies entirely after that cutover; the M30 window (Aug 17-Sep 14, 2026) contains it. Issue-1 and issue-2 configs that named grok-4.5 are replayed on the 4.6 lineage ([remap] rows: cw2-pt1, cw1-30ds2), and a remapped replay is not the same simulation the origin issue ran.
- **Concentrated rows are published, not hidden.** 4 rows run stakes above this issue's cap for their universe size and are marked [conc] (preset_sol, preset_xrp, cw2-30d1, cw2-30d2); their printed economics assume funding the current rule would not grant. They stay in the table because removing them would flatter the preset row.
- **The stake rule changes what runs, not only its size.** Capping the stake changes WHICH entries get funded: the seven table-stake solo defaults skip 135 to 317 entries each, and 30 replay rows carry unfunded entries (up to 23 on cw5-wk1). Where skipped_nofunds is non-zero the printed economics are not the economics that ran;

the rows are flagged, never silently pooled.

- **Snapshot principle.** This report is a frozen snapshot taken at the generated-at timestamps; numbers are not updated retroactively and past issues are not restated. The drift baseline exists precisely because today's engine and a longer forecast history do not reproduce every past number (cw5-m301: +\$1,717.1 printed, +\$1,219.8 today).
- **Research-to-date counter is pinned.** The counter below is pinned from wave 3 onward and is read at the Sep 14, 2026 16:00 UTC cut-off; issue #5 printed 42,905 resolved and 26 published reports against 47,373 and 30 here. The model line keeps the definition issue #4 introduced (7 tracked, current line-up; earlier versions folded into their successors' lineage), so the three issues' model counts are comparable and issue #3's is not.
- **No calendar month in this issue.** The night collect computes M30 and WK and nothing else, so no claim about a calendar month is made anywhere here. Monthly Config Watch #2 covers September and is produced after 01.10; Monthly Config Watch #1's 3 cards appear in this issue only as replay rows, judged by the weekly rule.
- **Multiple testing.** 16,111 simulations across the two collects (5,287 weekly + 10,824 night); at this scale some winners are expected from chance alone. Antidotes: axes frozen before launch, the walk-forward table, independent baselines, in-sample labeling. No formal correction yet.
- **Winner's curse and thin cells.** Every card, board row and per-ticker best is the maximum of a search, biased high by selection alone. On the week the per-ticker winners closed 5-21 trades and 16 of the 42 replay rows are under 10 closed; on the four weeks 1 replay rows are thin and 1 per-ticker winners are. Flagged with *, reported for completeness.
- Research output, not financial advice.

Counters & lineage

Weekly collect (WK window)

Stage	Sims	Stage	Sims
Scan (s1)	476	Dedicated per-ticker (s6t+s6tc)	560
Systematic grid (s2)	3,952	OOS replays (s4cw5+s4cw5o+s4cw4+s4cw4o+s4cwm1+s4cwm1o+s4cw3+s4cw3o+s4cw2+s4cw2o+s4cw1+s4pre)	68
Refinement: ladder + stake sweep (s2b+s3b)	118	Total	5,287
Preset neighbourhood grid (s2p)	53	Errors / zero-trade	0 / 1,121
Per-ticker probes (s5t)	60	Engine	1.1

Night collect (M30 and a second execution of the WK week; this issue publishes M30 from it, and its WK half feeds the cross-window stability board only)

Stage	Sims	Stage	Sims
Scan (s1)	952	Dedicated per-ticker (s6t+s6tc)	1,150
Systematic grid (s2)	8,160	OOS replays (s4cw5+s4cw5o+s4cw4+s4cw4o+s4cwm1+s4cwm1o+s4cw3+s4cw3o+s4cw2+s4cw2o+s4cw1+s4pre)	110
Refinement: ladder + stake sweep (s2b+s3b)	238	Total	10,824
Preset neighbourhood grid (s2p)	106	Errors / zero-trade	0 / 1,836
Per-ticker probes (s5t)	108	Engine	1.1

Scale note (vs issue #5). Issue #5 ran 16,398 sims over two windows in two collects; issue #6 runs 16,111 over two windows in two -- 5,287 in the weekly collect (WK) and 10,824 in the night collect (M30 and the same WK week together). The per-window grid is close to unchanged: 3,952 WK cells and 4,080 M30 cells here against 4,080 per window there. The per-coin branch is 560 and 1,150 against 580. The replay branch grew in rows (42 configs against 38) and in sims (68 and 110 against 60 and 98), because one more cohort joined and every row is replayed on every window of its collect. The two selection boards still cost zero simulations.

No extras pass in either collect. Equity series, the grid histograms (0 extra sims), the baselines and the calibrate packs are produced by the same two collects that wrote the boards. The baseline runs are verification, outside both search totals: 22 sims on

WK (11 configs x 2 stake modes) and 44 in the night collect, of which the 22 M30 rows belong to this issue. 0 equity replays were needed.

Generated at: 2026-09-15T20:56:34Z (weekly collect) and 2026-09-16T09:13:44Z (night collect) · 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 timestamps above; numbers are not updated retroactively. Each issue re-runs the pipeline fresh over that issue's windows.

Research to date

RESEARCH SNAPSHOT

THIS REPORT -- search effort: 16,111 simulations (0 errors) across two frozen collects, of which 178 walk-forward replays and origin re-runs · 66 verification sims (baseline runs inside the same collects; 0 equity replays needed)

MARKETMANIA RESEARCH TO DATE (as of cutoff): 47,373 directional forecasts resolved since Jul 11 (as of Sep 14, 2026 16:00 UTC cutoff) · 7 models tracked (current line-up; earlier versions folded into their successors' lineage) · 5 assets · 5 horizons · hourly · 30 published reports

Counter pinned from wave 3 onward and read at the Sep 14, 2026 16:00 UTC cut-off: issue #5 printed 42,905 resolved and 26 published reports. The model line keeps the definition issue #4 introduced (7 tracked, current line-up; earlier versions folded into their successors' lineage), so the three issues' model counts are comparable. Config Watch #6 is the 30th published report.

Issue #6. The standing walk-forward board reaches 38 cards -- issue #5's 4 join it, and monthly cards are judged by the weekly rule (Amendment 5). Two fresh windows again, the calendar week and the four calendar weeks ending at the same edge, both computed in the night collect as well as the week: the table keeps 6 of 38 past cards where the week alone keeps 8 and the month alone 19. The cross-window stability board is measured on this issue's own two windows (WK and M30) and comes back with 9 configurations. Formalization of the amendments is still slated for methodology v1.2.

What we're testing next

- **The first fresh week of this issue's own cards.** Config Watch #7 replays the 5 cards published here that carry a short link (both WK councils, the WK solo card, both M30 councils; the M30 solo card has no link and is not replayed, as issue #5's was not) on its own fresh week, alongside the 38 already on the board. The number to check against: 25.0% of issue #5's 4 linked cards were net-positive on their first fresh week, which is this issue's week.
- **Whether the WK-and-M30 stability board reads the same twice.** 9 configurations are in the mature net top-40 of both windows of this issue's night run, against 0 on issue #5's board and 21 on issue #3's, both across the same pair. Config Watch #7's night run computes the same pair one week on, so the question is whether the count moves when the windows do.
- **Whether the calibrate share holds for a fourth window in a row.** Calibration improves 49.9% of the 1,334 matched WK pairs and 48.4% of the 1,392 M30 pairs this issue, against 57.8% and 37.8% in issue #5. The test is the same two measurements on the next pair of windows -- the whole grid and the board-leader slice, reported separately and never pooled.

Related research

Consensus Watch #6 -- <https://marketmania.ai/research/reports/consensus-watch-2026-09-07.pdf>

Weekly Calibration #6 -- <https://marketmania.ai/research/reports/weekly-calibration-2026-09-07.pdf>

Weekly Model Watch #6 -- <https://marketmania.ai/research/reports/model-watch-2026-09-07.pdf>

Config Watch #5 -- <https://marketmania.ai/research/reports/config-watch-2026-09-02.pdf>

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

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

The three weekly reports publish together as one issue each week; Config Watch follows its own cycle. Market-regime figures, where this issue refers to them, come from that wave and are not re-tabulated here. Monthly Config Watch #1 is listed because its 3 cards are replayed in the walk-forward table above.

CITE THIS REPORT

```
@techreport{mm_configwatch_2026w38,
  title      = {Config Watch #6},
  author     = {{MarketMania Research}},
  institution = {MarketMania},
  year       = {2026},
  month      = sep,
  day        = {16},
  type       = {Weekly Research Report},
  series     = {Config Watch},
  number     = {6},
  note       = {Methodology v1.1 + Config Watch amendments, hash e66c7e8c864a2233},
  url        = {https://marketmania.ai/research/reports/config-watch-2026-09-09.pdf}
}
```