

Config Watch #8

October 1, 2026 · MarketMania Research · Weekly series

WEEKLY · CONFIG WATCH · Issue #8 · the ceiling slips, and the board with it · Issue date October 1, 2026 · Language English · Methodology v1.1 + Config Watch amendments (approved 12.08, hash e66c7e8c864a2233)

Search windows: **WK** Sep 21-28, 2026 and **M30** Aug 31-Sep 28, 2026 -- both frozen and Monday-aligned (WK = Mon Sep 21 00:00 -> Sun Sep 27 23:59 UTC; M30 = Mon Aug 31 00:00 -> Sun Sep 27 23:59 UTC, 4 full weeks), sharing the cut-off edge Mon Sep 28 00:00 UTC, so slots dated Sep 28 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

16,753 simulations, 0 errors, in two frozen collects	WK Sep 21-28, 2026	M30 Aug 31-Sep 28, 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 (w293 plan)
walk-forward: 48 past cards + 4 live presets replayed verbatim on both windows	10 of 48 past cards Survive (net-positive on both windows)	stability: 4 configs across WK and M30 (night run)	models: 7 (grok axis = grok-4.6 (incl. 4.5-era))
best WK council +\$583.4 (+58.3%)	best M30 council +\$1,018.6 (+101.9%)	zero-trade cells: 913 (weekly) / 1,698 (night run)	margin-skip alerts on leaders: 0 (WK) / 14 (M30) rows


KEY FINDING · OBSERVATION (the ceiling slips, and the board with it)

The ceiling of the fresh search moved from +\$1,467.8 to +\$583.4 on the week and 24 of the 48 cards this series has published are net-positive on that same week, and under the standing two-window rule 10 of the 48 Survive: 16 clear the month, 24 clear the week.

TL;DR

- OBSERVATION -- The standing board grew to 48 cards.** Issue #7 kept 14 of its 43 past cards under the same rule; this issue replays **48 cards** -- issues #1, #2, #3, #4, #5, #6, #7 and Monthly Config Watch #1 -- on two fresh windows, and **10 of the 48 Survive** (net-positive on both). On the week alone 24 of 48 are positive, on the four-week window 16 of 48. Issue #7's own week champion, replayed verbatim, makes +\$102.8 on the week and +\$421.2 on the month.
- The search ceiling, window by window.** The best mature council on WK makes +\$583.4 (+58.3% of a \$1,000 deposit, 73.2% win rate, 22.7% maxDD, 41 closed) against issue #7's +\$1,467.8; on M30 it makes +\$1,018.6 (+101.9%, 72.5% win rate, 48.5% maxDD, 91 closed) against issue #7's +\$1,351.8. The grids behind them: median \$0.0 with 47.7% positive across 4,208 WK cells (102 of them above +\$500), median -\$30.0 with 29.1% positive across 4,208 M30 cells (102 above +\$500).
- Calibration and the universe, window by window.** Across the whole stage-2 grid calibration improves 36.5% of the 1,468 matched WK pairs (mean -\$44.3) and 66.8% of the 1,463 matched M30 pairs (mean +\$242.9); pooled, 51.6% of 2,931. On the cards, **no published WK council is a calibrated cell** and **no published M30 council is a calibrated cell**. The mature net top-10 leans on one universe per window: BNB+XRP on WK (5 of 10 slots), SOL+XRP on M30 (6 of 10).
- The stability board, a fifth reading of the same pair.** 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 **4 configurations** are in the mature net top-40 of both. Issue #7 measured the same pair one week earlier and held 16, issue #6 the week before that 9, issue #5 0 and issue #3 21. The hand-tuned anchor keeps its standing measurement against the whole mature field of each window -- 643 of 2,429 mature WK configs and 992 of 2,925 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-#7).

Stage-2 systematic grid -- net PnL distribution -- WK window (Sep 21-28, 2026) -- n=4,208



Net PnL across all 4,208 stage-2 grid sims, WK window (Sep 21-28, 2026), universe-aware stake throughout. Grid median **\$0.0**, 47.7% positive (issue #7's week: median +\$79.2, 57.0% positive; its month: -\$83.5, 25.7% positive). The spike at \$0 is mostly cells whose entry filters produced no trades (912 zero-trade sims on this window). Range -\$590.6 to +\$842.1; 102 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 +\$1,467.8 down to +\$1,432.0; this week it runs +\$583.4 down to +\$486.6 -- 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, issue #6 the fifth and issue #7 the sixth, each on the widest board the series had had. Issue #8 delivers the seventh, on a board that carries every generation the series has produced: 48 cards from seven 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 (w293 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)** -- 4,208 + 8,416 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)** -- 121 + 243 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)** -- 70 + 130 sims: best finalists split onto single coins.
- **6. Dedicated per-ticker (s6t+s6tc)** -- 585 + 1,165 sims: a compacted independent per-coin grid plus calibrate twins (the standing branch from issue #2).
- **7. OOS replays (s4cw7+s4cw7o+s4cw6+s4cw6o+s4cw5+s4cw5o+s4cw4+s4cw4o+s4cwm1+s4cwm1o+s4cw3+s4cw3o+s4cw2+s4cw2o+s4cw1+s4pre)** -- 88 + 140 sims: every published card and all four live presets, verbatim, on every fresh window, plus origin re-runs of the issue-7 and issue-6 cards for the drift baseline.

Total 16,753 simulations (5,601 weekly + 11,152 night), 0 errors, 2,611 zero-trade cells (913 + 1,698); stage sums reconcile exactly in each collect (counts.by_stage). The axis grid was NOT widened relative to issue #7 (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 14 on M30 (night collect); the night collect raised 14 alerts in total, all of them on M30. 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 21-28, 2026)

#	Council	Coins	FH/TF	Net	Net %	maxDD	Trades	WR
#1	gpt, gemini, qwen	SOL+XRP	4h/4h,1h	+\$583.4	+58.3%	22.7%	41	73.2%
#2	gpt, gemini, qwen, grok	BNB+XRP	4h/4h	+\$570.6	+57.1%	37.6%	38	63.2%
#	Ladder	BE	Shifts	Conf filter	Source	Entry	Positional	Stake
#1	50:40,100:60	on	TP -10 / SL +75	>=60 (each)	farthest	3/0	reenter	\$450
#2	50:40,100:60	on	TP -10 / SL +75	>=60 (each)	farthest	3/0	reenter	\$450

Against the 4,208-cell stage-2 grid on this window: card #1 lands in the +\$500 to +\$750 bucket (97 of 4,208 cells) and card #2 in the +\$500 to +\$750 bucket (97 cells); 102 cells in the whole grid finished above +\$500 (grid median \$0.0). 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.


Two councils are published on this window. The mature net board carries more than one universe: BNB+XRP holds 5 of its ten slots, SOL+XRP 3 (card #1 among them) and single-coin XRP 2. Card #2 is the best mature multi-coin council outside card #1's universe (the pairwise-different-universes rule), and it comes off the mature net board itself (+\$570.6 on BNB+XRP, R2 0.002, ulcer 10.89, 37.6% maxDD, 0 unfunded entries).

On the calibrate axis no card is a calibrated cell, where issue #7 published 0 of its 2 WK councils as calibrated cells. Card #1 runs the 2-coin table stake of \$450. Neither card carries an unfunded entry.

OPEN IN SANDBOX marketmania.ai/s/cw8-wk1 marketmania.ai/s/cw8-wk2

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

Published cards equity -- WK (Sep 21-28, 2026) -- \$1,000 at 10x -- 53 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. The muted dashed line is the net-board maximum, drawn for contrast only: at 12 closed trades it sits below the ≥ 15 maturity gate and is not a card. 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 31-Sep 28, 2026)

#	Council	Coins	FH/TF	Net	Net %	maxDD	Trades	WR
#1	fable, opus, gemini, qwen	SOL+XRP	4h/1h	+\$1,018.6	+101.9%	48.5%	91 *	72.5%
#2	fable, opus, gemini, qwen	BNB+XRP	4h/1h	+\$943.0	+94.3%	37.7%	99 *	69.7%
#	Ladder	BE	Shifts	Conf filter	Source	Entry	Positional	Stake
#1	50:40,100:60	on	TP -10 / SL +75	≥ 60 (each)	farthest	2/0	reenter	\$450
#2	50:40,100:60	on	TP -10 / SL +75	≥ 60 (each)	farthest	2/0	reenter	\$450

Against the 4,208-cell stage-2 grid on this window: card #1 is the same simulation as the raw net-board leader ($i=2668$), so its exported percentile is the card's -- above **100.0%** of the grid; card #2 lands in the +\$750 to +\$1,000 bucket (13 cells) with 99.7% of the grid finished below it. Grid median -\$30.0, 102 of the 4,208 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 too. The mature net board carries more than one universe: SOL+XRP holds 6 of its ten slots (card #1 among them) and BNB+XRP 4. Card #2 is the best mature multi-coin council outside card #1's universe (the pairwise-different-universes rule), and it comes off the mature net board itself (+\$943.0 on BNB+XRP, R2 0.836, ulcer 9.48, 37.7% maxDD, 1 unfunded entry).

On the calibrate axis no published M30 council is a calibrated cell, where no published WK council is a calibrated cell -- the same axis, read on two windows (Calibrate axis below). Card #1 pays for its +\$1,018.6 with a **48.5% drawdown** on 91 closed trades and 11 unfunded entries, card #2 with 37.7% on 99 and 1. The M30 mature pool is 2,925 configs against 2,429 on WK, because four weeks clear the ≥ 20 -trade gate that one week does not.

OPEN IN SANDBOX marketmania.ai/s/cw8-m301 marketmania.ai/s/cw8-m302


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

Published councils equity -- M30 (Aug 31-Sep 28, 2026) -- \$1,000 at 10x -- 91 settled trades



Daily settled equity for card #1 above (\$1,000 start, 10x leverage, frozen window; #1 red -- fixed rank colors across all issues). A curve running past the window end is an open-at-cutoff trade settling at its horizon (the 2-3 day luft); 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=2668, 91 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.

Stage-2 systematic grid -- net PnL distribution -- M30 window (Aug 31-Sep 28, 2026) -- n=4,208




Net PnL across all 4,208 stage-2 grid sims, M30 window (Aug 31-Sep 28, 2026), universe-aware stake throughout. Grid median **-\$30.0**, 29.1% positive against 47.7% on the week (issue #7's month: median **-\$83.5**, 25.7% positive). The spike at \$0 is mostly cells whose entry filters produced no trades (786 zero-trade sims on this window). Range **-\$829.1** to **+\$1,018.6**; 102 cells finished above **+\$500**, against 102 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 48 cards from issues #1, #2, #3, #4, #5, #6, #7 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 platform, listed first. The monthly issue's 3 cards are judged by the weekly rule, like every other row. The replays and their origin re-runs cost 88 simulations in the weekly collect and 140 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 #7's right edge (Sep 21). New M30 (Aug 31-Sep 28, 2026) shares 21 of its 28 days with issue #7's own M30 window (shared stretch Aug 31-Sep 20), 14 of its 28 days with issue #6's own M30 window (shared stretch Aug 31-Sep 13), 7 of its 28 days with issue #5's own M30 window (shared stretch Aug 31-Sep 6) and 1 of its 28 days with Monthly Config Watch #1's calendar-August window (shared stretch Aug 31), so for those four cohorts that column is decay context rather than pure out-of-sample: the same asymmetry issues #3, #4, #5, #6 and #7 disclosed for their own second columns.

Issue #7's five published cards: home number vs fresh-week replay



Issue #7's 5 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). 1 of the 5 is net-negative on the fresh week; the 2 M30 cards were measured at home on a four-week window and the 3 WK cards were on a one-week window, so the grey bars are not all the same kind of number. Source: oos_cw7_wk.json + issue #7 open data.

Config	Family	Coins	Home	New WK (n)	New M30 (n)	Verdict
preset 4h4	fable, qwen, deepseek, opus	ETH+SOL	-- live	-\$103.8 (18)	-\$158.1 (41) +	Faded
preset bc2	gemin, gpt, deepseek	SOL+XRP	-- live	+\$37.4 (10)	-\$465.9 (19) +	Faded
preset sol	deepseek, fable, opus	SOL [conc]	-- live	-\$193.3 (16)	-\$77.8 (42)	Faded
preset xrp	gemin, gpt, deepseek	XRP [conc]	-- live	+\$402.2 (5) *	-\$387.0 (1) * +	Faded
cw7-m301	fable, gpt, gemini, qwen	SOL+XRP	+\$1,351.8	+\$280.9 (53)	+\$848.1 (136) +	Survived
cw7-m302	opus, gpt, gemini, qwen	BTC+ETH	+\$528.4	+\$42.7 (12)	+\$156.0 (30)	Survived
cw7-s1	gemin	SOL+XRP	+\$1,022.0	+\$94.4 (60)	+\$193.1 (172) +	Survived
cw7-wk1	fable, opus, gpt, gemini	SOL+XRP	+\$1,467.8	+\$102.8 (66)	+\$421.2 (160) +	Survived
cw7-wk2	fable, opus, gpt	BTC+ETH	+\$953.2	-\$203.0 (21)	-\$635.8 (22) +	Faded
cw6-m301	opus, deepseek, gemini	SOL+XRP	+\$1,793.8	-\$24.3 (50)	+\$365.5 (146) +	Faded
cw6-m302	gpt, deepseek, gemini, grok	BTC+ETH	+\$547.8	+\$128.7 (15)	-\$280.9 (33) +	Faded
cw6-s1	grok	SOL+XRP	+\$404.8	+\$85.7 (40)	-\$614.8 (6) * +	Faded
cw6-wk1	opus, gpt, grok	SOL+XRP	+\$609.5	-\$346.6 (50) +	+\$237.3 (143) +	Faded
cw6-wk2	fable, opus, deepseek	BTC+ETH	+\$135.4	-\$63.7 (23)	-\$44.7 (80) +	Faded
cw5-m301	fable, opus, gpt, gemini	SOL+XRP	+\$1,717.1	+\$299.8 (11)	-\$597.0 (9) * +	Faded
cw5-s1	gemin	BNB+XRP	+\$178.7	-\$311.4 (51) +	-\$30.5 (177) +	Faded
cw5-wk1	gpt, gemini	SOL+XRP	+\$510.4	+\$354.9 (78)	+\$1,117.5 (268)	Survived
cw5-wk2	fable, opus, gpt, gemini, grok	BTC+ETH	+\$183.5	-\$92.9 (33)	+\$51.5 (109)	Faded
cw4-m301	gpt, deepseek, gemini	SOL+XRP	+\$1,676.0	+\$684.8 (14)	+\$428.9 (33) +	Survived
cw4-m302	fable, gpt, deepseek, gemini	BNB+XRP	+\$1,155.2	+\$305.4 (11)	-\$581.9 (10) +	Faded
cw4-s1	deepseek	BTC+ETH+ SOL+BNB+XRP	+\$325.0	-\$55.8 (29) +	-\$537.1 (53) +	Faded
cw4-wk1	opus, deepseek, gemini, qwen	BTC+ETH+ SOL+BNB+XRP	+\$399.4	+\$68.1 (32) +	+\$36.8 (78) +	Survived
cw4-wk2	fable, deepseek, qwen	BNB+XRP	+\$128.2	+\$23.2 (22)	+\$316.2 (98)	Survived
cwm1-m1	fable, gemini	SOL+XRP	+\$1,618.4	+\$656.9 (12)	-\$670.3 (16) +	Faded

Config	Family	Coins	Home	New WK (n)	New M30 (n)	Verdict
cwm1-m2	fable, gpt, deepseek, gemini, qwen	ETH+SOL+BNB+XRP	+\$756.6	+\$81.8 (28)	-\$18.2 (52) +	Faded
cwm1-s1	gemini	SOL+XRP	+\$1,721.9	+\$646.2 (13)	-\$686.2 (18) +	Faded
cw3-m301	opus, deepseek, gemini	SOL+XRP	+\$2,137.6	-\$81.7 (40)	-\$584.5 (24) +	Faded
cw3-m302	fable, opus, deepseek, gemini, qwen	BTC+ETH	+\$412.3	-\$15.9 (4) *	-\$302.8 (14) +	Faded
cw3-s1	gemini	SOL+XRP	+\$1,604.3	+\$388.6 (66)	-\$557.1 (70) +	Faded
cw3-s2	gemini	SOL+XRP	+\$1,449.0	+\$646.2 (13)	-\$686.2 (18) +	Faded
cw3-wk1	opus, deepseek, gemini	SOL+XRP	+\$2,175.1	-\$81.7 (40)	-\$584.5 (24) +	Faded
cw3-wk2	opus, deepseek, gemini, grok	BTC+ETH	+\$837.6	-\$79.0 (20)	-\$557.9 (26) +	Faded
cw2-30d1	opus, gpt, deepseek, gemini, qwen	BTC+ETH+SOL+BNB+XRP [conc]	+\$744.7	+\$333.8 (15) +	-\$727.3 (12) +	Faded
cw2-30d2	opus, deepseek, gemini	ETH+SOL+BNB+XRP [conc]	+\$655.1	+\$547.2 (17) +	-\$279.0 (25) +	Faded
cw2-30d3	opus, deepseek, gemini	BNB+XRP	+\$489.1	+\$172.8 (10)	+\$56.4 (28) +	Survived
cw2-30ds1	gemini	BNB+XRP	+\$383.8	+\$470.2 (11)	-\$660.8 (14) +	Faded
cw2-7d1	fable, opus, deepseek, gemini	SOL+XRP	+\$299.1	-\$112.2 (38) +	-\$592.7 (17) +	Faded
cw2-7d2	gpt, deepseek, gemini, qwen	BNB+XRP	+\$223.7	+\$399.7 (13)	-\$89.1 (29) +	Faded
cw2-pt1	fable, deepseek, grok [remap]	SOL	+\$698.3	-\$216.3 (22)	-\$50.3 (61)	Faded
cw2-pt2	opus	XRP	+\$543.7	+\$329.6 (6) *	-\$723.8 (7) * +	Faded
cw1-30d1	fable, deepseek, opus, qwen	ETH+SOL	+\$868.4	-\$197.6 (19)	+\$44.4 (45) +	Faded
cw1-30d2	gemini, gpt, deepseek	SOL+XRP+ETH	+\$590.6	-\$119.6 (16)	-\$530.4 (31) +	Faded
cw1-30d3	fable, deepseek, opus, qwen	BTC+ETH+SOL	+\$580.5	-\$153.0 (23)	-\$23.4 (68) +	Faded
cw1-30ds1	opus	SOL+XRP	+\$442.8	-\$87.8 (39)	+\$142.5 (103) +	Faded
cw1-30ds2	grok [remap]	SOL+XRP	+\$262.9	-\$255.4 (28) +	-\$280.9 (75) +	Faded
cw1-30ds3	deepseek	XRP+BNB	+\$119.0	-\$167.0 (36) +	-\$582.0 (14) +	Faded
cw1-7d1	fable, deepseek, opus, gpt	SOL+XRP	+\$439.2	-\$330.3 (45) +	+\$111.5 (145) +	Faded
cw1-7d2	fable, qwen, gemini, opus	XRP+BNB	+\$429.5	-\$211.4 (29) +	-\$69.9 (82) +	Faded
cw1-7d3	gpt, opus	BNB+XRP+SOL	+\$398.6	-\$143.5 (53) +	-\$574.2 (122) +	Faded
cw1-7ds1	gemini	XRP+BNB	+\$417.8	+\$27.9 (68) +	+\$406.9 (212) +	Survived
cw1-7ds2	gpt	XRP+BNB	+\$366.3	-\$487.6 (39) +	-\$501.8 (112) +	Faded
cw1-7ds3	qwen	XRP+BNB	+\$267.0	-\$340.7 (35) +	-\$555.2 (83) +	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: 10 of 48 past cards **Survive** both fresh windows (issue #7 on its two windows: 14 of 43) -- by cohort 4 of 5 issue-7 cards, 0 of 5 issue-6 cards, 1 of 4 issue-5 cards, 3 of 5 issue-4 cards, 0 of 3 Monthly-1 cards, 0 of 6

issue-3 cards, 1 of 8 issue-2 cards, 1 of 12 issue-1 cards. Taken a window at a time: 24 of 48 are positive on the week and 16 of 48 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, fifth, sixth and seventh fresh week, reads: issue-1 50.0% -> 91.7% -> 8.3% -> 16.7% -> 25.0% -> 91.7% -> **8.3%** now (1 of 12); issue-2 100.0% -> 50.0% -> 12.5% -> 25.0% -> 62.5% -> **75.0%** now (6 of 8); issue-3 33.3% -> 33.3% -> 33.3% -> 83.3% -> **33.3%** now (2 of 6); issue-4 20.0% -> 0.0% -> 80.0% -> **80.0%** now (4 of 5); Monthly-1 66.7% -> 0.0% -> 33.3% -> **100.0%** now (3 of 3); issue-5 25.0% -> 75.0% -> **50.0%** now (2 of 4); issue-6 80.0% -> **40.0%** now (2 of 5); issue-7 **80.0%** now (4 of 5). Live presets: 0 of 4 Survive, 2 of 4 positive on the week, 0 of 4 on the month.

The four live presets lead the table by standing rule. On the week: preset 4h4 -\$103.8 on 18 closed, preset bc2 +\$37.4 on 10 closed, preset sol -\$193.3 on 16 closed, preset xrp +\$402.2 on 5 closed -- 2 of the four net-positive. On the four-week window: preset 4h4 -\$158.1 on 41 closed, preset bc2 -\$465.9 on 19 closed, preset sol -\$77.8 on 42 closed, preset xrp -\$387.0 on 1 closed -- 0 of four. 0 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-7 cards on their OWN origin windows today lands within -\$228.8..\$0.0 of the printed issue-7 numbers and the issue-6 cards within -\$158.7..\$0.0 of theirs; 4 of the 10 re-runs reproduce to the cent. The widest deviation is cw7-m302 (+\$528.4 printed, +\$299.6 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 the same direction this week. On the fresh week 24 of 48 past cards are net-positive -- issue #7 counted 33 of 43 on its own week -- and the ceiling of the fresh search moved to +\$583.4 from +\$1,467.8. On the fresh four weeks 16 of 48 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 48 cards is not comparable with one computed on 43 without saying which generations joined -- 5 of the 48 rows are new this issue. The month column is not clean evidence for four of the eight cohorts: it shares 21 of its 28 days with issue #7's own M30 window (shared stretch Aug 31-Sep 20), 14 of its 28 days with issue #6's own M30 window (shared stretch Aug 31-Sep 13), 7 of its 28 days with issue #5's own M30 window (shared stretch Aug 31-Sep 6) and 1 of its 28 days with Monthly Config Watch #1's calendar-August window (shared stretch Aug 31), 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 cw1-7ds2 (-\$487.6 on the week against +\$366.3 printed at home), while the largest home number on the board, cw3-wk1 at +\$2,175.1, made -\$81.7. Seven fresh weeks of verdicts now: brutal, the opposite, split by window, the fourth, the fifth, the sixth, 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 (22 WK pairs and 14 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	1,468	36.5%	-\$44.3	+\$503.9	-\$797.7
M30	1,463	66.8%	+\$242.9	+\$1,204.1	-\$1,001.7
Both	2,931	51.6%	+\$99.1	+\$1,204.1	-\$1,001.7

The two windows answer the axis, and the whole-grid reading and the board-leader slice are separate measurements. On the week calibration improves 36.5% of the 1,468 matched pairs (mean -\$44.3, median \$0.0) and 54.0% of the 648 mature pairs (mean +\$21.2); on the four weeks it improves 66.8% of 1,463 (mean +\$242.9, median +\$214.3) and 72.2% of the 706 mature pairs (mean +\$213.8). Pooled over both windows the axis improves 51.6% of 2,931 pairs with a mean of +\$99.1 -- a number that describes neither window on its own. The board-leader slices are separate: 0.0% of the 22 WK slice pairs improve (mean -\$475.3, median -\$537.4, OFF legs already at a median of +\$742.3), and 14.3% of the 14 M30 slice pairs (mean -\$441.2, OFF legs at +\$857.0). Issue #7 measured a 14-pair WK slice (0 improving) and a 12-pair M30 slice (0). Where the axis reached the cards is visible above: 0 of the 2 published WK councils are calibrate=ON cells and 28 of the 54 WK finalists run it, while 0

of the 2 published M30 councils are and 42 of the 60 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 -\$256.5 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 21-28, 2026)

#	Model	Coins	FH/TF	Net	Net %	maxDD	Trades	WR
#1	gemini	BNB+XRP	4h/4h	+\$445.4	+44.5%	35.6%	50	60.0%
#	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 on this window, not two. Every other row of the mature solo board is either the same model or a single-coin universe, so the one-config-per-model and ≥ 2 -coin rules admit exactly one; it runs the 2-coin table stake of \$450 on 50 closed trades with 0 unfunded entries. The one single-coin row that outranks it on the mature net board (gemini on XRP at +\$496.9) is excluded by the ≥ 2 -coin rule and appears under Per-ticker bests. Source: summary_wk.json top_mature_by_window_branch.

Solo top -- M30 (Aug 31-Sep 28, 2026)

#	Model	Coins	FH/TF	Net	Net %	maxDD	Trades	WR
#1	grok	BNB+XRP	4h/4h,1h	+\$802.0	+80.2%	29.9%	105 *	76.2%
#2	gemini	SOL+XRP	4h/1h	+\$616.2	+61.6%	36.2%	199	83.4%
#	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	≥ 60 (each)	median	0/0	reenter	\$450

Two solo cards on M30. The mature solo board carries two models, grok on BNB+XRP and gemini on SOL+XRP, so the one-config-per-model rule admits two. Unfunded entries on the cards: #1 4, #2 0 (flagged * on Trades); the best row on the board with 0 skips is card #2 itself, a calibrate=ON cell. Source: top_mature_by_window_branch in the night collect.

OPEN IN SANDBOX marketmania.ai/s/cw8-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.

Council vs solo. On the week the best solo sits below the best council (+\$445.4 vs +\$583.4 on the mature board, 50 closed against 41, 35.6% drawdown against 22.7%). On the four weeks the best solo sits below the best council (+\$802.0 vs +\$1,018.6, 105 closed against 91, 29.9% against 48.5%). Same in-sample caveats apply to both sides.

Quality boards

Net PnL stays the primary board (continuity with issues #1-#7); 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	CIs	R2
WK Calmar	fable, gpt, deepseek, gemini, grok [cal]	BNB+XRP	4h/4h,1h	+\$385.8	93.8%	0.0%	16	0.880
WK Calmar	fable, gpt, deepseek, gemini, grok [cal]	BNB+XRP	4h/4h,1h	+\$385.8	93.8%	0.0%	16	0.880
WK Calmar	fable, gpt, deepseek, gemini, grok [cal]	BNB+XRP	4h/4h,1h	+\$385.8	93.8%	0.0%	16	0.880

Board	Config	Coins	FH/TF	Net	WR	maxDD	CIs	R2
WK WR	gpt, gemini, qwen, grok [cal]	BTC+ETH	4h/4h,1h	+\$194.2	100.0%	0.0%	16	0.905
WK WR	gpt, gemini, qwen, grok [cal]	BTC+ETH	4h/4h,1h	+\$194.2	100.0%	0.0%	16	0.905
WK WR	gpt, gemini, qwen, grok [cal]	BTC+ETH	4h/4h,1h	+\$161.8	100.0%	0.0%	16	0.905
WK Smoothness	gpt, deepseek, gemini, qwen, grok [cal]	BNB+XRP	4h/1h	+\$307.8	94.1%	0.0%	17	0.965
WK Smoothness	gpt, deepseek, gemini, qwen, grok [cal]	BNB+XRP	4h/1h	+\$317.4	90.0%	0.0%	20	0.963
WK Smoothness	gpt, gemini, qwen, grok	SOL+XRP	4h/1h	+\$57.6	66.7%	46.1%	45	0.956
WK Quality	gpt, deepseek, gemini, qwen, grok [cal]	BNB+XRP	4h/1h	+\$307.8	94.1%	0.0%	17	0.965
WK Quality	gpt, deepseek, gemini, qwen, grok [cal]	BNB+XRP	4h/1h	+\$317.4	90.0%	0.0%	20	0.963
WK Quality	gpt, deepseek, gemini, qwen [cal]	BNB+XRP	4h/4h,1h	+\$300.7	93.3%	0.0%	15	0.952
M30 Calmar	fable, opus, gpt, qwen [cal]	SOL+XRP	1d/1d	+\$535.5	90.9%	0.0%	22	0.958
M30 Calmar	fable, opus, gpt, qwen [cal]	SOL+XRP	1d/1d	+\$535.5	90.9%	0.0%	22	0.958
M30 Calmar	fable, opus, gpt, qwen [cal]	SOL+XRP	1d/1d	+\$446.3	90.9%	0.0%	22	0.958
M30 WR	gemini, gpt, deepseek [cal]	SOL+XRP	1d/1d	+\$572.9	94.9%	13.5%	39	0.879
M30 WR	gemini, gpt, deepseek [cal]	SOL+XRP	1d/1d	+\$477.5	94.9%	11.2%	39	0.879
M30 WR	gemini, gpt, deepseek [cal]	SOL+XRP	1d/1d	+\$382.0	94.9%	9.0%	39	0.879
M30 Smoothness	fable, opus, gpt, qwen [cal]	SOL+XRP	1d/1d	+\$535.5	90.9%	0.0%	22	0.958
M30 Smoothness	fable, opus, gpt, qwen [cal]	SOL+XRP	1d/1d	+\$535.5	90.9%	0.0%	22	0.958
M30 Smoothness	fable, opus, gpt, qwen [cal]	SOL+XRP	1d/1d	+\$357.0	90.9%	0.0%	22	0.958
M30 Quality	fable, opus, gpt, qwen [cal]	SOL+XRP	1d/1d	+\$535.5	90.9%	0.0%	22	0.958
M30 Quality	fable, opus, gpt, qwen [cal]	SOL+XRP	1d/1d	+\$535.5	90.9%	0.0%	22	0.958
M30 Quality	fable, opus, gpt, qwen [cal]	SOL+XRP	1d/1d	+\$446.3	90.9%	0.0%	22	0.958

[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, and this issue neither does: the WK board's top row is a BNB+XRP config at R2 0.965 on a net of +\$307.8, the M30 board's a SOL+XRP config at R2 0.958 on +\$535.5, each inside the universe its window's mature net top-10 leans on. That trade -- straightness bought with size -- is exactly what the board exists to expose. The M30 boards are drawn from a mature pool of 2,925 configs against 2,429 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. 4**

configurations qualify. All of them BNB+XRP, 4 councils and 4 clean (zero unfunded entries on both windows); the top three by rank sum are council fable, opus, gpt, gemini, qwen; council fable, opus, gpt, gemini, qwen; council fable, opus, gpt, gemini, qwen. The shared board is narrower than either window's own: 8 of the 20 rows exported from its WK board are SOL+XRP; 10 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 #7's, at 16, issue #6's, at 9, 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, fifth 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	Cls
WK	Hand-tuned anchor (1d_3_llm_BC2_v02)	SOL+XRP	1d/1d	+\$37.4	70.0%	41.4%	0.019	12.37	10
M30	Hand-tuned anchor (1d_3_llm_BC2_v02)	SOL+XRP	1d/1d	-\$465.9	63.2%	55.2%	0.301	30.33	19

643 of the 2,429 mature WK configs (26.5%) clear all three bars at once, and 992 of the 2,925 mature M30 configs (33.9%). The anchor's own two windows are different rows: on the week +\$37.4 net at 70.0% win rate with a 41.4% drawdown, an R2 of 0.019 and 10 closed trades; on the four weeks -\$465.9 at 63.2%, a 55.2% drawdown, an R2 of 0.301 and 19 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, all 10 of the 10 best WK winners by net also beat it on net PnL; all 10 of the 10 best M30 winners by net also beat it on net PnL. The share itself moves with the window (26.5% of the WK field, 33.9% of the M30 field), so read it as one field at a time, not as a property of the anchor.

Top-10 by net among the 643 WK configs that beat the anchor on win rate, maxDD and smoothness at once

#	Config	Coins	FH/TF	Stake	Net	WR	maxDD	R2
#1	gpt, gemini, qwen	SOL+XRP	4h/4h,1h	\$450	+\$583.4	73.2%	22.7%	0.023
#2	gemini, qwen	XRP	4h/4h	\$450	+\$499.7	76.9%	22.8%	0.029
#3	fable, gpt, gemini, qwen, grok	SOL+XRP	4h/1h	\$450	+\$472.6	70.4%	11.0%	0.274
#4	gpt, gemini, qwen [cal]	BNB+XRP	4h/4h	\$450	+\$414.6	88.6%	11.2%	0.483
#5	fable, gpt, gemini, grok [cal]	BNB+XRP	4h/1h	\$450	+\$408.9	83.3%	7.7%	0.562
#6	gpt, deepseek, gemini, qwen [cal]	SOL+XRP	4h/1h	\$450	+\$399.2	89.1%	7.7%	0.733
#7	gpt, deepseek, gemini, qwen [cal]	SOL+XRP	4h/1h	\$450	+\$399.2	89.1%	7.7%	0.733
#8	gpt, gemini, qwen [cal]	BNB+XRP	4h/4h	\$450	+\$391.8	88.4%	25.2%	0.039
#9	gpt, gemini, qwen [cal]	BNB+XRP	4h/4h	\$450	+\$391.8	88.4%	25.2%	0.039
#10	fable, gpt, gemini, qwen, grok [cal]	BNB+XRP	4h/4h	\$450	+\$389.6	88.9%	11.9%	0.459

Top-5 by net among the 992 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, gemini, qwen	SOL+XRP	4h/1h	\$450	+\$1,018.6	72.5%	48.5%	0.824
#2	fable, opus, gpt, gemini, qwen	SOL+XRP	4h/1h	\$450	+\$968.8	71.9%	49.0%	0.819
#3	fable, opus, gemini, qwen	BNB+XRP	4h/1h	\$450	+\$943.0	69.7%	37.7%	0.836
#4	fable, opus, gemini, qwen	SOL+XRP	4h/1h	\$450	+\$908.9	69.9%	54.4%	0.770
#5	fable, opus, gemini, qwen	BNB+XRP	4h/4h	\$450	+\$899.8	64.8%	36.4%	0.765

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 585 + 1,165 sims -- a compacted grid, one best per coin per window; the branch tag in brackets says which kind won.

WK (Sep 21-28, 2026)

Coin	Best config (branch)	FH/TF	Net (trades)
BTC	gpt [solo]	4h/1h	+\$164.3 (12)
ETH	gemini, gpt, deepseek [council]	1d/1d	+\$156.3 (3) *
SOL	gpt, deepseek, gemini [council]	1d/1d,4h	+\$271.6 (5) *
BNB	gemini, gpt, deepseek [council]	1d/1d	+\$123.4 (3) *
XRP	fable, gpt, gemini, grok [council]	1d/1d	+\$590.9 (5) *

M30 (Aug 31-Sep 28, 2026)

Coin	Best config (branch)	FH/TF	Net (trades)
BTC	gpt [solo]	4h/4h,1h	+\$244.6 (26)
ETH	gemini, grok [council]	4h/4h	+\$357.8 (68)
SOL	fable, opus, gpt, gemini [council]	4h/1h	+\$649.5 (90)
BNB	grok [solo]	4h/4h,1h	+\$402.1 (55)
XRP	fable, gemini, qwen [council]	4h/4h	+\$626.8 (50)

n in parentheses; * = fewer than 10 closed trades (thin -- direction only). 1 of the ten winners runs calibrate ON; 4 of the five WK winners and 3 of the five M30 winners are councils.

The reads. On the week XRP leads at +\$590.9 and 4 of the five rows are thin -- 3 to 12 closed trades against a 15-trade gate. On the four weeks no row is thin: 26 to 90 closed, SOL leading at +\$649.5, and no one coin leads both windows. The 1d horizon wins 4 of the five WK coins and 0 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 21-28, 2026)

Baseline	Net PnL	Trades	WR	maxDD
Best solo default (claude-fable-5)	-\$3.1 *	306	65%	5.6%
Worst solo default (deepseek-v4-pro)	-\$81.9 *	365	55%	10.6%
Council-of-7 default	-\$20.4 *	181	63%	3.9%
Hand-tuned config A (pre-search)	+\$546.3	12	92%	7.5%
Hand-tuned config B (pre-search)	-\$327.6 *	59	59%	54.0%
Hand-tuned config C (pre-search)	-\$197.6	19	58%	45.8%

M30 (Aug 31-Sep 28, 2026)

Baseline	Net PnL	Trades	WR	maxDD
Best solo default (grok-4.6)	-\$72.6 *	760	62%	8.2%
Worst solo default (deepseek-v4-pro)	-\$193.5 *	1203	59%	19.3%
Council-of-7 default	-\$62.8 *	541	62%	6.3%
Hand-tuned config A (pre-search)	-\$108.1 *	25	64%	72.4%
Hand-tuned config B (pre-search)	-\$31.3 *	168	64%	54.3%
Hand-tuned config C (pre-search)	+\$44.4 *	45	67%	45.8%

* = non-zero skipped_nofunds at table stake. On the 5-coin defaults those counts are the stake table doing its job: at \$200 a \$1,000 deposit funds at most five concurrent positions while the defaults fire hundreds of entries per window (101 to 361 skipped each on the week, 486 to 1,360 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 they do so for every row. 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 split on both windows. On WK the seven solo defaults run -\$3.1 to -\$81.9 at table stake, the council-of-7 -\$20.4, and the three hand-tuned configs A +\$546.3, B -\$327.6, C -\$197.6 -- 1 of the three positive. On M30 the defaults run -\$72.6 to -\$193.5 solo (council-of-7 -\$62.8) and the three hand-tuned configs A -\$108.1, B -\$31.3, C +\$44.4 -- 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 +\$583.4 on the week against a best baseline of +\$546.3, and +\$1,018.6 on the four weeks against +\$44.4. The live presets are a separate row of evidence: 2 of the four are net-positive on the week and 0 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	+\$37.4	+\$802.6 (1d/1d, SOL+XRP, 7 cls)	34	neighbour +\$765.2
preset bc2	M30	-\$465.9	+\$572.9 (1d/1d, SOL+XRP, 39 cls)	34	neighbour +\$1,038.8
preset 4h4	WK	-\$103.8	+\$172.8 (4h/4h,1h, ETH+SOL, 13 cls)	19	neighbour +\$276.5
preset 4h4	M30	-\$158.1	+\$197.4 (4h/1h, SOL+XRP, 45 cls)	19	neighbour +\$355.5

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 bc2's calibrate switch (off -> on) on M30, which turns -\$465.9 into +\$572.9. 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, 42 of them councils; 60 on M30, 50 councils).

Knob	WK finalists mode	M30 finalists mode
Forecast horizon	4h (50/54)	4h (38/60)
Ladder shape (steps)	50:40,100:60 (54/54)	50:40,100:60 (57/60)
Break-even stop (BE)	on (54/54)	on (60/60)
Min confidence	60 (52/54)	60 (59/60)
SL shift	75 (54/54)	75 (57/60)
TP shift	-10 (54/54)	-10 (57/60)
TP/SL source	farthest (48/54)	farthest (51/60)
Calibrate ON	28/54	42/60
Universe	BNB+XRP (24/54)	SOL+XRP (34/60)

The mechanical core is the same on both windows and unchanged for the eighth issue running: 50:40,100:60 ladder, break-even on, conf 60, SL +75 / TP -10, farthest source. The horizon is **4h** on WK (50/54) and **4h** on M30 (38/60). **Calibrate=ON** runs in 28 of 54 WK finalists against 42 of 60 on M30, where issue #7 measured 11/62 on its week and 25/51 on its month. The universe: BNB+XRP takes 24 of the 54 WK finalists while SOL+XRP takes 34 of the 60 M30 finalists. Membership concentrates too -- gpt leads the WK council finalists (41 of 42) and fable the M30 ones (41 of 50). Standing caveats: refinement seeds from the same grid winners, so part of the convergence is search-design echo; and a knob core that produced 14-of-43 survival one issue ago and 10-of-48 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 22.7% and 37.6% on the week and 48.5% and 37.7% on the month, and the replay table is deeper: 23 replayed rows drew down more than 40% on the week and 39 on the month, 9 and 27 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 31-Sep 28, 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 **48 cards** -- issue #7's 5 join issue #6's 5, issue #5's 4, 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 issues #5, #6 and #7 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 31-Sep 28, 2026) contains WK (Sep 21-28, 2026) -- 7 of the week's 7 days are inside it -- and it shares 21 of its 28 days with issue #7's own M30 window (shared stretch Aug 31-Sep 20), 14 of its 28 days with issue #6's own M30 window (shared stretch Aug 31-Sep 13), 7 of its 28 days with issue #5's own M30 window (shared stretch Aug 31-Sep 6) and 1 of its 28 days with Monthly Config Watch #1's calendar-August window (shared stretch Aug 31). The two columns of this issue are therefore not two independent tests: agreement between them is partly arithmetic, and for the issue-7, issue-6, issue-5 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 2026-09-30T05:59:07Z), M30 the night run (2026-09-30T12:07:43Z -- boards computed at 11:49:01Z, 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 +\$472.6 (fable, gpt, gemini, qwen, grok on SOL+XRP) against the weekly collect's +\$583.4. 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 4 configurations are in the mature net top-40 of both. Issue #7 read the same pair one week earlier and held 16, issue #6 the week before that held 9, issue #5 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 21-28, 2026) lies entirely after that cutover; the M30 window (Aug 31-Sep 28, 2026) lies entirely after 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 101 to 361 entries each, and 15 replay rows carry unfunded entries (up to 18 on cw2-30d1). 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 (cw7-m302: +\$528.4 printed, +\$299.6 today).
- **Research-to-date counter is pinned.** The counter below is pinned from wave 3 onward and is read at the Sep 28, 2026 16:00 UTC cut-off; issue #7 printed 52,971 resolved and 34 published reports against 57,929 and 38 here. The model line keeps the definition issue #4 introduced (7 tracked, current line-up; earlier versions folded into their successors' lineage), so the five 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.** 16,753 simulations across the two collects (5,601 weekly + 11,152 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 3-12 trades and 3 of the 52 replay rows are under 10 closed; on the four weeks 4 replay rows are thin and 0 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)	585
Systematic grid (s2)	4,208	OOS replays (s4cw7+s4cw7o+s4cw6+s4cw6o+s4cw5+s4cw5o+s4cw4+s4cw4o+s4cw1+s4cwm1o+s4cw3+s4cw3o+s4cw2+s4cw2o+s4cw1+s4pre)	88
Refinement: ladder + stake sweep (s2b+s3b)	121	Total	5,601
Preset neighbourhood grid (s2p)	53	Errors / zero-trade	0 / 913
Per-ticker probes (s5t)	70	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)	1,165
Systematic grid (s2)	8,416	OOS replays (s4cw7+s4cw7o+s4cw6+s4cw6o+s4cw5+s4cw5o+s4cw4+s4cw4o+s4cw1+s4cwm1o+s4cw3+s4cw3o+s4cw2+s4cw2o+s4cw1+s4pre)	140
Refinement: ladder + stake sweep (s2b+s3b)	243	Total	11,152
Preset neighbourhood grid (s2p)	106	Errors / zero-trade	0 / 1,698
Per-ticker probes (s5t)	130	Engine	1.1

Scale note (vs issue #7). Issue #7 ran 16,642 sims over two windows in two collects; issue #8 runs 16,753 over two windows in two -- 5,601 in the weekly collect (WK) and 11,152 in the night collect (M30 and the same WK week together). The per-window grid is close to unchanged: 4,208 WK cells and 4,208 M30 cells here against 4,176 per window there. The per-coin branch is 585 and

1,165 against 585 and 1,165. The replay branch grew in rows (52 configs against 47) and in sims (88 and 140 against 78 and 125), 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: 2026-09-30T05:59:07Z (weekly collect) and 2026-09-30T12:07:43Z (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: 16,753 simulations (0 errors) across two frozen collects, of which 228 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): 57,929 directional forecasts resolved since Jul 11 (as of Sep 28, 2026 16:00 UTC cutoff) · 7 models tracked (current line-up; earlier versions folded into their successors' lineage) · 5 assets · 5 horizons · hourly · 38 published reports

Counter pinned from wave 3 onward and read at the Sep 28, 2026 16:00 UTC cut-off: issue #7 printed 52,971 resolved and 34 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 five issues' model counts are comparable. Config Watch #8 is the 38th published report.

Issue #8. The standing walk-forward board reaches 48 cards -- issue #7's 5 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 10 of 48 past cards where the week alone keeps 24 and the month alone 16. The cross-window stability board is measured on this issue's own two windows (WK and M30) and comes back with 4 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 #9 replays the 5 cards published here that carry a short link (both WK councils, the WK solo card, both M30 councils; both M30 solo cards have no link and are not replayed, as issue #7's three unlinked solo cards were not) on its own fresh week, alongside the 48 already on the board. The number to check against: 80.0% of issue #7's 5 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.** 4 configurations are in the mature net top-40 of both windows of this issue's night run, against 16 on issue #7's board, 9 on issue #6's, 0 on issue #5's and 21 on issue #3's, all across the same pair. Config Watch #9's night run computes the same pair one week on, so the question is whether the count moves when the windows do.
- **What the calibrate share reads on a sixth window.** Calibration improves 36.5% of the 1,468 matched WK pairs and 66.8% of the 1,463 M30 pairs this issue, against 25.8% and 65.9% in issue #7. 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 #8 -- <https://marketmania.ai/research/reports/consensus-watch-2026-09-21.pdf>

Weekly Calibration #8 -- <https://marketmania.ai/research/reports/weekly-calibration-2026-09-21.pdf>

Weekly Model Watch #8 -- <https://marketmania.ai/research/reports/model-watch-2026-09-21.pdf>

Config Watch #7 -- <https://marketmania.ai/research/reports/config-watch-2026-09-16.pdf>

Config Watch #6 -- <https://marketmania.ai/research/reports/config-watch-2026-09-09.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_2026w40,
  title      = {Config Watch #8},
  author     = {{MarketMania Research}},
  institution = {MarketMania},
  year      = {2026},
  month     = oct,
  day       = {1},
  type      = {Weekly Research Report},
  series    = {Config Watch},
  number    = {8},
  note      = {Methodology v1.1 + Config Watch amendments, hash e66c7e8c864a2233},
  url       = {https://marketmania.ai/research/reports/config-watch-2026-09-23.pdf}
}
```