Monthly Consensus Watch #3

November 5, 2026 · MarketMania Research · Monthly series · window Oct 1-31, 2026 UTC

RESEARCH SNAPSHOT			
19,271 LOO-scoreable calls	**7** models (lineage spliced: opus, gpt, grok)	**3,308** thin/tied excluded	window **Oct 1-31, 2026 UTC**
method = **leave-one-out** majority, 3+ directional peers	FH gate **1h / 4h / 1d**	base field hit **48.6%**	cutoff **Sun Nov 1, 2026, 16:00 UTC**
market: BTC net **-7.22%** (prior +6.42%)	ann vol **28.3%** (was 41.2%)	TOP5 volume **$68.94B**, -14% m/m	pairwise corr **0.87** (was 0.78)

KEY FINDING · OBSERVATION (one monthly window)

Herding ran at 98.7% of leave-one-out scoreable calls across the month and the pooled lift came in at +2.7pp: agree-hit 48.9% against 46.1% on disagreement, on a disagree bucket of n=245.

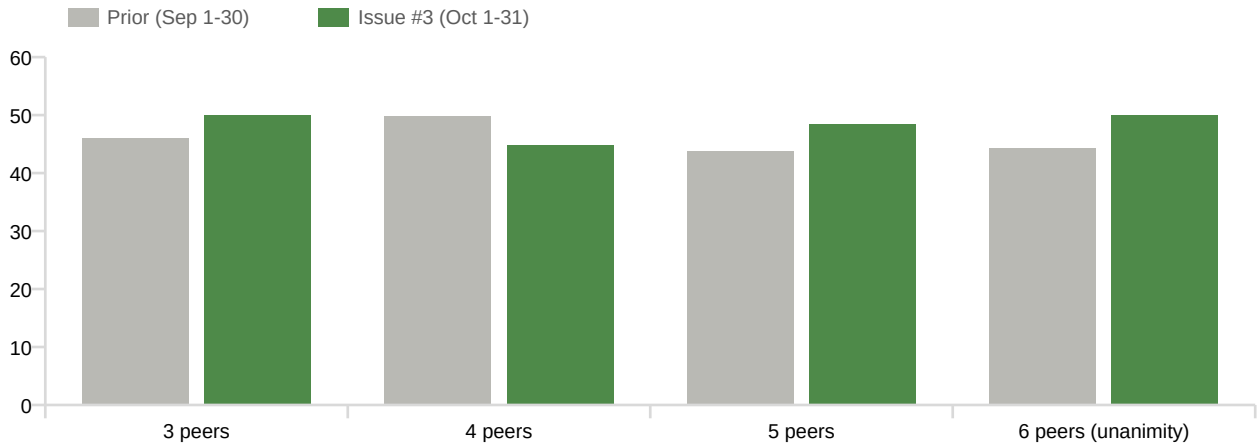
WEEKLY EVIDENCE -> MONTHLY VERDICT (rule below)

Metric	W1 Oct 5-11	W2 Oct 12-18	W3 Oct 19-25	W4 Oct 26-Nov 1	Month	Verdict
LOO lift, pp	+2.2	+1.8	+7.0	-8.5	+2.7	Confirmed (3/4)
Herd share >= 95%	98.1%	98.8%	98.9%	98.7%	98.7%	Confirmed (4/4)
Unanimity hit vs field hit	-0.7	-1.8	+4.7	-3.2	+1.4	Rejected (1/4)

Verdict rule: each row states one reading of the month; a week agrees when it shows the same reading (same sign, or the same name). Confirmed = 3 or 4 of the 4 weeks agree with the month; Inconclusive = 2; Rejected = 0 or 1. Weekly cells come from the frozen weekly packs (7-line roster); the month is the pooled Oct 1-31 pull. W4 straddles the month edge (Nov 1); Oct 1-4 belong to the month only.
The model cutover of Oct 17, 2026 09:01 UTC (gpt-5.6-sol -> gpt-6.1-sol, claude-opus-5 -> claude-opus-5.5, grok-4.6 -> grok-4.7) falls inside W2; W1 runs on the earlier builds, W3 and W4 on the new ones, the month column joins both -- all per-line rows are model lines, not single builds.

TL;DR

- **OBSERVATION -- 98.7% of the field's leave-one-out calls ran with the herd over the month.** Of 19,271 LOO-scoreable directional calls, 19,026 (98.7%) sided with peers' majority; agree-hit was 48.9% [48.1%, 49.6%] vs. 46.1% [40.0%, 52.4%] for the 245 disagreers, a pooled lift of +2.7pp. Descriptive, one month, not a claim of a durable edge.
- **The margin curve read 50.0% / 44.7% / 48.4% / 50.0% from 3 peers to unanimity.** 2 of the 4 tiers sit above the 48.6% field base. The same four tiers read 45.9% / 49.7% / 43.8% / 44.3% in the prior window Sep 1-30, on n=1,712 / 2,400 / 3,972 / 9,317 there.
- **Week by week the lift changed sign 1 time.** The four published weeks ran Oct 5-11 +2.2pp, Oct 12-18 +1.8pp, Oct 19-25 +7.0pp, Oct 26-Nov 1 -8.5pp; the month, which also counts the 4 days outside those weeks, pools at +2.7pp and the prior window read -9.7pp. The month is one pooled pull, not the mean of the weeks: the weekly range is -8.5pp to +7.0pp.
- **2,126 cells had 6+ directional models on one side across the month.** The 5 worst sat at 69.3-71.3 mean confidence and resolved 0.0%-14.3%; the prior window Sep 1-30 had 1,965 such cells, hitting 44.2% on n=13,121.



Hit rate (%) by number of agreeing LOO peers, the prior window against the whole month. Monthly n: 3 peers 1,940 / 4 peers 2,545 / 5 peers 4,266 / 6 peers 10,101. Prior n: 3 peers 1,712 / 4 peers 2,400 / 5 peers 3,972 / 6 peers 9,317. Field base this month 48.6%, prior window 45.5%. The prior window is the full September: window A of the same pull, slots 2026-09-01 to 2026-10-01 (first slot actually seen 2026-09-01 00:01 UTC), matured at the prior cutoff Thu Oct 1, 2026, 16:00 UTC under the same engine, the same FH gate and the same hit rule as October -- 20,682 directional calls against 22,579 in the month. The prior 6-peer tier is the published Monthly #2's six-peer bucket (n=9,317), the same strict leave-one-out definition as this month's; the pack's prior block carries the wider unanimity rule instead (every directional call of a cell with 6 or more models on one side, n=13,121, 44.2%), which is the Unanimity column of the week-by-week table, not a margin tier.

Why it matters

MarketMania publishes forecasts from 7 frontier model lines into the same slots every hour; a natural question for anyone reading the feed is whether a model's call is worth more when it agrees with what everyone else is saying, or whether that agreement is just everyone making the same mistake together. The four weekly issues set against this month answered it with lifts of Oct 5-11 +2.2pp, Oct 12-18 +1.8pp, Oct 19-25 +7.0pp, Oct 26-Nov 1 -8.5pp, and the pooled month reads +2.7pp. This monthly issue asks the same question at monthly n, across four weeks the pack labels as 2 calm and 2 volatile, using the same leave-one-out design built to avoid comparing a model to an index that includes its own vote.

How to read this

For every model M and every forecast it made, we rebuild that slot's consensus using only the OTHER active models in the same symbol / forecast-horizon (FH) / timeframe (TF) cell -- M's own call never counts toward its own consensus. M is scored agree if its side (long or short) matches the strict majority of those peers, and disagree if it sits alone against that majority; a cell needs at least 3 directional peers to count at all, and tied peer splits are excluded rather than forced either way. LOO peers are genuinely external to the model being scored, so the comparison below is not comparing a model to itself. 3 model lines are lineage splices this month (claude-opus-5.5 (incl. opus-5-era, Oct 1-17); gpt-6.1-sol (incl. 5.6-era, Oct 1-17); grok-4.7 (incl. 4.6-era, Oct 1-17)), and a spliced line never counts as two peers in a cell. The month is one pull at one cutoff (Sun Nov 1, 2026, 16:00 UTC); the four weekly blocks quoted alongside it are the published issues at their own frozen cutoffs.

Pooled result: agree vs. disagree with the LOO consensus

Group	n	Hit rate	Wilson 95% CI	Lift, pp
Agree with LOO peers	19,026	48.9%	[48.1%, 49.6%]	+2.7 (vs. disagree)
Disagree with LOO peers	245	46.1%	[40.0%, 52.4%]	ref.

The disagree bucket is n=245 over the whole month; the two Wilson intervals overlap. Prior window Sep 1-30: agree 45.1% / disagree 54.7%, lift -9.7pp, on n=17,488 agree and n=137 disagree. These are dependent observations pooled across the whole month: two calm weeks, two volatile weeks and 4 days outside them.

By forecast horizon / timeframe

The pooled result above mixes 5 different FH/TF cells. Broken out:

FH / TF	Agree n	Agree hit	Dis. n	Dis. hit	Lift, pp	Excluded
1h / 1h	11,708	49.9%	176	48.9%	+1.0	2,058
4h / 4h	3,124	47.0%	32	43.8%	+3.3	556
4h / 1h	3,135	46.5%	24	41.7%	+4.9	496
1d / 1d	514	45.5%	2	0.0%	+45.5	105
1d / 4h	545	54.9%	11	27.3%	+27.6	93

"Excluded" = thin (fewer than 3 directional peers) or tied peer splits in that cell; not folded into agree or disagree. "Dis." = disagree. Cells whose disagree bucket is under N=10 are insufficient and are never used to rank anything.

Margin: does more agreement mean a safer call?

For every LOO-scored forecast that sided with its peers' majority (the agree bucket, n=19,026) we also count how many peers were on that majority side -- from the bare minimum of 3 up to the full peer set agreeing (6 of 6, since 7 model lines were active and LOO always drops one); the 245 disagreeing calls sit in no tier, so the tiers sum to the agree bucket, not to the 19,271 scoreable calls. The month-over-month chart is on page 1; the full per-tier numbers: 3 peers 50.0% (n=1,940), 4 peers 44.7% (n=2,545), 5 peers 48.4% (n=4,266), 6 peers 50.0% (n=10,101). Against the 48.6% field base: 2 tiers above, 2 tiers below; in the prior window the same tiers read 45.9% / 49.7% / 43.8% / 44.3%. The week-by-week table further down prints each tier for each of the four weeks. The prior 6-peer tier is the published Monthly #2's six-peer bucket (n=9,317), the same strict leave-one-out definition as this month's; the pack's prior block carries the wider unanimity rule instead (every directional call of a cell with 6 or more models on one side, n=13,121, 44.2%), which is the Unanimity column of the week-by-week table, not a margin tier.

Daily-horizon herds: the thinnest cell, the best rate

The two daily cells stay the report's thinnest even at monthly n. 1d/1d unanimity hit 45.9% (n=301) and 1d/4h unanimity 62.7% (n=308) against the 48.6% field base; the whole 1d/1d agree bucket ran 45.5% (n=514) -- not the highest of the five FH/TF cells -- while 1d/4h ran 54.9% (n=545). Both daily buckets are an order of magnitude thinner than the hourly one (n=11,708), so they are reported, not ranked.

Herd failures: the 5 worst unanimous misses

When the whole field agreed, and was wrong. This month 2,126 slot / symbol / FH / TF cells had 6 or more directional models sitting on the same side.

Slot (UTC)	Symbol	FH / TF	Side	Models	Mean conf	Hit rate
Oct 2, 12:01	BTC	4h / 4h	long	7	71.3	0.0%
Oct 5, 00:01	BTC	1d / 4h	long	7	70.9	0.0%
Oct 5, 00:01	BTC	4h / 4h	long	7	70.4	0.0%
Oct 21, 16:01	SOL	4h / 1h	long	7	70.1	0.0%
Oct 4, 16:01	XRP	1h / 1h	long	7	69.3	14.3%

The 5 worst cells were long swarms at 69.3-71.3 mean confidence and resolved 0.0%-14.3%. 'Sideways' was 44.8% of mature forecasts (18,310 of 40,889) this month.

Per-model: agree vs. disagree

Per-model disagree counts run 11 to 96 over the month. 7 of the 7 lines clear the N=10 floor (claude-fable-5 11, claude-opus-5.5 (spliced) 17, deepseek-v4-pro 41, gemini-3.1-pro 96, gpt-6.1-sol (spliced) 29, grok-4.7 (spliced) 31, qwen-3.8-max 20); every other per-model contrarian cell below is insufficient.

Model	Agree n	Agree hit	Disagree n	Disagree hit
claude-fable-5	2,630	48.6%	11	27.3%
claude-opus-5.5 (spliced)	2,443	49.3%	17	76.5%
deepseek-v4-pro	2,834	48.8%	41	46.3%
gemini-3.1-pro	2,671	48.8%	96	43.8%
gpt-6.1-sol (spliced)	2,970	49.3%	29	24.1%
grok-4.7 (spliced)	2,655	48.7%	31	58.1%
qwen-3.8-max	2,823	48.7%	20	55.0%

Prior window Sep 1-30: per-model agree-side hit ran 43.5% to 47.2% on n=2,097 to 2,660 per line. At these counts a handful of flipped calls moves a per-model lift by tens of points, so models are not ranked by contrarian lift here. claude-opus-5.5 (spliced) = claude-opus-5.5 (incl. opus-5-era, Oct 1-17). gpt-6.1-sol (spliced) = gpt-6.1-sol (incl. 5.6-era, Oct 1-17). grok-4.7 (spliced) = grok-4.7 (incl. 4.6-era, Oct 1-17). In the prior window the lineage splice folds 6,300 claude-opus-5 rows into claude-opus-5.5, 6,300 gpt-5.6-sol rows into gpt-6.1-sol and 6,300 grok-4.6 rows into grok-4.7 (18,900 rows in all).

Week by week: how often the field agreed, and what it paid

The same leave-one-out scoring run on each of the four Mon-Sun weeks from Oct 5 to Nov 1 gives four weekly lifts; the pooled monthly number also counts the 4 days outside those weeks. Herd share ran 98.1% to 98.9% and the lift ran -5.1pp to +7.0pp across the four weeks against +2.7pp for the month as a whole.

Week	Scoreable	Herd share	Agree hit	Disagree hit	Lift pp	Unanimity n	Unanimity hit
W1 Oct 5-11	3,735	98.1%	43.6%	41.4%	+2.2	2,476	44.1%
W2 Oct 12-18	4,330	98.8%	46.3%	44.4%	+1.8	3,166	45.5%
W3 Oct 19-25	5,156	98.9%	57.9%	50.8%	+7.0	3,931	59.1%
W4 Oct 26-31	3,874	98.7%	40.9%	46.0%	-5.1	2,879	41.1%
Month	19,271	98.7%	48.9%	46.1%	+2.7	14,199	49.4%

Source: the month pull split by slot week at the single monthly cutoff (Sun Nov 1, 2026, 16:00 UTC), so the four week rows and the Month row are one consistent pool. The frozen-cutoff cross-check is the published weekly issues at their own cutoffs: the recomputed lift reproduces the published lift for 3 of 4 weeks (W1 +2.2pp vs +2.2pp, W2 +1.8pp vs +1.8pp, W3 +7.0pp vs +7.0pp, W4 -5.1pp vs -8.5pp). W4 holds Oct 26-31 here: the month pull ends on Oct 31. * = N<10 (insufficient) -- such cells are never ranked.

Margin	W1	W2	W3	W4	Month
3 peers	44.0% (452)	51.4% (432)	55.0% (480)	44.7% (356)	50.0% (1,940)
4 peers	38.5% (605)	46.0% (615)	53.5% (620)	36.0% (520)	44.7% (2,545)
5 peers	48.4% (858)	47.1% (1,068)	52.8% (1,128)	46.4% (858)	48.4% (4,266)
6 peers (unanimity)	43.1% (1,708)	44.8% (2,128)	61.7% (2,821)	39.1% (2,051)	50.0% (10,101)

Cells are hit rate with n in parentheses; * = N<10 (insufficient), "n/a (0)" = no call landed in that tier that week. Weeks are W1 Oct 5-11, W2 Oct 12-18, W3 Oct 19-25, W4 Oct 26-31. Same source and cutoff as the table above; the Month column is the pooled month, not the mean of the four weeks. Month-over-month columns compare back-to-back windows: Sep 1-30 vs Oct 1-31 (this issue). "Was" values are the pack's own prior-window block -- window A of the same pull, same engine, same FH gate and same hit rule, matured at the prior cutoff Thu Oct 1, 2026, 16:00 UTC. The prior month is the full September (30 days against 31 in October), and it holds 20,682 directional calls against 22,579 in October. Deltas are computed on unrounded rates.

Market check: October ran calmer than September

OBSERVATION — mean |1d move| 1.91% -> 1.57% (-17% rel), BTC realized vol 34.9% -> 32.4% (ann., hourly); field directional accuracy 45.5% -> 48.6% (+3.1pp), 6 of 7 models improved, sim win-rate up for 6 of 7, field sim PnL -$2554.15 -> -$1648.92 (Sep 1-30 vs Oct 1-31; descriptive, one pair of windows, not a claim).
Robustness: raw price-sign accuracy 48.4% -> 48.1% — the opposite direction to the trade-based hit rule.

Measure	Sep 1-30	Oct 1-31	Change
Mean |1d move| (5 assets)	1.91%	1.57%	-17% rel
BTC realized vol (ann., hourly)	34.9%	32.4%	-2.5pp
Field directional accuracy	45.5%	48.6%	+3.1pp
Models improving hit-rate	--	6 of 7	--
Raw price-sign accuracy	48.4%	48.1%	-0.3pp
Field sim win-rate	38.4%	39.6%	+1.2pp
Field sim net PnL	-$2554.15	-$1648.92	--

Window A is the pack's prior window -- the full September -- so this slice and every month-over-month column in this report compare October against the full September. Descriptive, one pair of windows; the hit rule is methodology v1.1 (trade-based) and the raw price-sign row is the robustness check (hourly closes at :00 against slots at :01). BTC realized vol here is the hourly-return series from the alive slice; the snapshot row's figure is the audited daily-candle estimate -- different estimators, both reported as measured. Lineage splices in this window: claude-opus-5.5 (incl. opus-5-era, Oct 1-17); gpt-6.1-sol (incl. 5.6-era, Oct 1-17); grok-4.7 (incl. 4.6-era, Oct 1-17). 10,395 rows were remapped inside the month against 18,900 in the prior window.

Practical implications

- 2 of the 4 margin tiers sit above the 48.6% field base this month; the four tier rates span 44.7% to 50.0%, a spread of +5.4pp, and the four weeks set against the month ran lifts of -8.5pp to +7.0pp against a pooled monthly lift of +2.7pp.
- The pooled monthly lift is +2.7pp while the four weeks ran -8.5pp to +7.0pp and the prior window read -9.7pp.
- The field base moved +3.1pp from September to October (45.5% -> 48.6%), and every consensus tier is measured against that base.

Limitations

- The prior window is the full September: every "prior" column, the gray series of the page-1 chart and the market-check A column describe Sep 1-30. It is window A of the same pull, slots 2026-09-01 to 2026-10-01 (first slot actually seen 2026-09-01 00:01 UTC), matured at the prior cutoff Thu Oct 1, 2026, 16:00 UTC under the same engine, the same FH gate and the same hit rule as October -- 20,682 directional calls against 22,579 in the month. Month-over-month deltas are a step between two windows of different length (30 and 31 days).
- Two cutoffs live in this issue. The month is frozen at Sun Nov 1, 2026, 16:00 UTC; the four weekly blocks are frozen at their own published cutoffs (W1 Oct 5-11, W2 Oct 12-18, W3 Oct 19-25, W4 Oct 26-Nov 1). A forecast still pending at its week's cutoff is immature in the weekly pack and mature in the month pull, so the month does not equal the sum of the weeks (22,579 vs 20,635 directional, plus 2,505 on the 4 remainder days); W4 also runs to Nov 1, a day the month pull does not hold, so its pack counts 5,094 directional calls against 4,533 for Oct 26-31 inside the month pull -- that gap is arithmetic, not a data problem.
- Observations are not independent: the same models watch overlapping symbol / FH / TF cells hour after hour for 31 days. Wilson intervals here are descriptive, not inferential, and pooling a month does not make them inferential.
- The pooled lift rests on 245 disagreeing calls; per-model counts run 11-96. 3,308 calls (14.7% of the 22,579-call mature directional universe) were thin or tied and are excluded from every consensus table. Correlational, not causal: LOO removes self-match bias, not the confound that an easy slot lifts agreement and hit together.
- This window ran 9 trend days of 31, and the pack's own volatility rule splits the four Mon-Sun weeks from Oct 5 to Nov 1 as W1 volatile, W2 calm, W3 volatile, W4 calm; every month-over-month comparison is a comparison across regimes as well as across windows. Series density and lineage notes (monthly #3): (1) rolling-1w forecasts emit daily since Aug 22, 2026 and rolling-1M daily since Aug 23, 2026, so across this 31-day window daily coverage is FULL as measured -- the 1w series has slots on 31 of 31 days and the 1M series on 31 of 31

days; both sit outside this report's FH gate (1h/4h/1d) and appear only in the exclusions counter (1w 2,160, 1M 2,135). (2) Model lineage is a splice inside this window. Three lines flipped at Oct 17, 2026 09:01 UTC, day 17 of 31 of the month: "gpt-6.1-sol (incl. 5.6-era, Oct 1-17)" (3,465 rows from the gpt-5.6-sol era, 3,045 from gpt-6.1-sol); "claude-opus-5.5 (incl. opus-5-era, Oct 1-17)" (3,465 rows from the claude-opus-5 era, 3,045 from claude-opus-5.5); "grok-4.7 (incl. 4.6-era, Oct 1-17)" (3,465 rows from the grok-4.6 era, 3,045 from grok-4.7). Every per-model aggregate sees the merged head id (the raw id is kept in the pack); 10,395 rows were remapped inside the month, and in the prior window 6,300 rows of each of the three lines (18,900 in all) are earlier-build rows, so each month-over-month row for those lines joins a spliced month to a single-build prior month -- a remapped comparison is not the same model build. (3) The two hourly slots before the flip (Oct 17, 07:01 and 08:01 UTC) hold 60 invalid rows (claude-opus-5 20, gpt-5.6-sol 20, grok-4.6 20), counted in the invalid counter as they are; in those two hours the other 4 lines (claude-fable-5, deepseek-v4-pro, gemini-3.1-pro, qwen-3.8-max) carry ok rows. (4) Past lineage flips: the grok 4.5 -> 4.6 flip (Aug 24, 2026 09:22 UTC) and the qwen-3.7-max -> qwen-3.8-max succession (Aug 5, 2026) both sit before the prior window.

- Market-state row is computed on a single-exchange (binance) daily candle series, as in the weekly series after the audit that found the raw table mixes two exchanges. The prior-month block is the full September: it covers 30 days (Sep 1-30), and the TOP5 volume figure compares a 31-day sum with a 30-day sum -- a level difference, not a like-for-like change.
- Research-to-date counter: directional forecasts resolved inside the FH gate (1h/4h/1d) since Jul 11, counted at the issue cutoff (Sun Nov 1, 2026, 16:00 UTC) -- the D2 definition pinned in the weekly series; this issue prints 82,334. The pack carries no prior-cutoff pin for a monthly window (the prior-cutoff pin is null by construction), so the control this issue is the weekly reproduction: the W4 block of this pack (Oct 26-Nov 1) reproduces the published weekly #13 on 7 of 7 checks, and the prior window of this pack (Sep 1-30, prior cutoff Thu Oct 1, 2026, 16:00 UTC) reproduces the published Monthly #2 on 7 of 7 checks. Counter deltas against pre-wave-3 issues remain definitional.

Counters & lineage

Counter	Value	Counter	Value
Forecasts total	45,540	Mature (scored pool)	40,889
OK in gate	40,889	-- of them directional	22,579
Out of gate (1w / 1M)	2,160 / 2,135	-- of them sideways	18,310
Invalid	356	Pending (next issue)	0
Late closes	0	Uptime, all grids	1,178 / 1,178 slots
Source file	monthly_metrics_2026-10.json	Report cutoff	Sun Nov 1, 2026, 16:00 UTC
Generated at (pipeline)	2026-11-03 12:39 UTC	Methodology	v1.1 (2026-08-10) · hash e66c7e8c864a2233

THIS REPORT: 19,271 LOO-scored observations. MARKETMANIA RESEARCH TO DATE (as of cutoff Nov 1, 16:00 UTC): 82,334 resolved since Jul 11 · 7 models tracked (current line-up; earlier versions folded into their successors' lineage) · 5 assets · 5 horizons · hourly · 67 published reports.

Issue #3. Monthly Consensus Watch is a living monthly comparison; each issue appends one more calendar month of leave-one-out agreement-vs-hit data. This third issue reads October against the full September: herd share 98.7%, pooled lift +2.7pp (agree 48.9% vs disagree 46.1% on n=245), four weekly lifts set against the month running -8.5pp to +7.0pp, and the prior window Sep 1-30 at herd share 99.2% and lift -9.7pp. 3 model lines are lineage splices this month (claude-opus-5.5 (incl. opus-5-era, Oct 1-17); gpt-6.1-sol (incl. 5.6-era, Oct 1-17); grok-4.7 (incl. 4.6-era, Oct 1-17)). Engine 1.1 has powered the sandbox since Aug 18, i.e. before this window and before the prior window; no cross-engine PnL comparisons are claimed.

What we're testing next

- Monthly #4 (November): the same three series over Nov 1-30, cut on Dec 1 plus the maturity lag, so its month-over-month comparison runs against this October -- the first month whose prior window carries the cutover splice.
- Monthly Config Watch #3: the same October window through the config search engine, published as the fourth monthly report of this issue.

- Quarterly Stability Index: the next repeat-prompt stability run is the next cross-quarter point on model stability; its date is announced with the run.

Related research

Report	Direct PDF link
Monthly Calibration #3	https://marketmania.ai/research/reports/calibration-monthly-2026-10.pdf
Monthly Model Watch #3	https://marketmania.ai/research/reports/model-watch-monthly-2026-10.pdf
Consensus Watch #13	https://marketmania.ai/research/reports/consensus-watch-2026-10-26.pdf

The three monthly reports publish together as one issue each month; each links straight to the others' PDF and to the last weekly issue of the same series. Direct links are the posting rule from wave 2 on.

CITE THIS REPORT

```
@misc{mm_consensus_monthly_2026m10,
  title  = {Monthly Consensus Watch #3: does agreeing with the crowd make an LLM's market call safer?},
  author = {{MarketMania Research}},
  year   = {2026}, month = {November}, day = {5},
  url    = {https://marketmania.ai/research/reports/consensus-watch-monthly-2026-10.pdf},
  note   = {Methodology v1.1; window Oct 1-31, 2026 UTC; source monthly_metrics_2026-10.json}
}
```