

Consensus Watch #7

September 22, 2026 · MarketMania Research · Weekly series · window Sep 14-20, 2026 UTC


RESEARCH SNAPSHOT			
4,877 LOO-scoreable calls	7 models (stable lineup)	721 thin/tied excluded	window Sep 14-20, 2026 UTC
method = leave-one-out majority, 3+ directional peers	FH gate 1h / 4h / 1d	base field hit 50.7%	cutoff Mon Sep 21, 2026, 16:00 UTC
market: BTC net +5.64% (prior -4.36%)	ann vol 51.0% (was 19.7%)	TOP5 volume \$18.74B , +17% w/w	pairwise corr 0.93 (was 0.74)

KEY FINDING · OBSERVATION (one weekly window)

Herding ran at 98.9% of leave-one-out scoreable calls and the pooled lift came in at +12.1pp: agree-hit 50.6% against 38.5% on disagreement, on a disagree bucket of n=52.

TL;DR

- **OBSERVATION -- 98.9% of the field's leave-one-out calls ran with the herd.** Of 4,877 LOO-scoreable directional calls, 4,825 (98.9%) sided with peers' majority (issue #6: 99.3%); agree-hit was 50.6% [49.2%, 52.0%] vs. 38.5% [26.5%, 52.0%] for the 52 disagreeers, a pooled lift of +12.1pp. Descriptive, one week, not a claim of a durable edge.
- **The margin curve took a seventh reading.** 3 peers 52.1% (was 42.8%), 4 peers 55.5% (was 46.9%), 5 peers 44.2% (was 40.4%), full unanimity 51.6% (was 34.2%). 3 of the 4 tiers sit above the 50.7% field base, and 2 of the 4 kept the side of the base they took in issue #6.
- **Daily-horizon unanimity split around the field base.** 1d/1d margin-6 hit 38.4% (n=112) and 1d/4h margin-6 hit 53.8% (n=91); issue #6 printed 0.0% and 14.3% on the same two cells.
- **Herd failures did not disappear.** 539 cells had 6+ directional models on one side (issue #6: 417); the 5 worst sat at 68.4-70.0 mean confidence and resolved 0.0%-0.0%.



Hit rate (%) by number of agreeing LOO peers, week over week. Issue #7 n: 424 / 685 / 990 / 2,681; field base this week 50.7% (issue #6: 39.6%). A 2-peer bucket exists this week too (n=45, 42.2%) -- footnote only, far too small to plot. "Was" values are the numbers published in issue #6.

Why it matters

MarketMania publishes forecasts from 7 frontier model lines into the same slots every week; 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. Consensus Watch answers that question empirically, one week at a time, using a leave-one-out design built specifically to avoid the obvious trap of 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. The grok rows are one line, so grok never counts as two peers in a cell. The grok 4.5 -> 4.6 flip (Aug 24, 2026 09:22 UTC) sits three windows back (issue #4): both weeks of this issue are pure grok-4.6, with no 4.5-era rows in either window, so the grok week-over-week row is the second pure-4.6 against pure-4.6 comparison of the series.

Pooled result: agree vs. disagree with the LOO consensus

Group	n	Hit rate	Wilson 95% CI	Lift, pp
Agree with LOO peers	4,825	50.6%	[49.2%, 52.0%]	+12.1 (vs. disagree)
Disagree with LOO peers	52	38.5%	[26.5%, 52.0%]	ref.

Read with context. The disagree bucket is n=52 this week against 26 in issue #6 and 26 in issue #5; the two Wilson intervals overlap. This is a single week of dependent observations: not evidence about herding in general, and not a claim that the sign will hold next week.

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	2,952	50.8%	40	42.5%	+8.3	468
4h / 4h	785	52.6%	3	33.3%	+19.3	110
4h / 1h	787	54.1%	8	25.0%	+29.1	113
1d / 1d	146	35.6%	1	0.0%	+35.6	19
1d / 4h	155	31.6%	0	n/a	n/a	11

"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 we also count how many peers were on the winning (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 week-over-week chart is on page 1; the full per-tier numbers: 3 peers 52.1% (n=424), 4 peers 55.5% (n=685), 5 peers 44.2% (n=990), 6 peers 51.6% (n=2,681). Against the 50.7% field base, 3 tier(s) sit above and 1 below, and 2 of the 4 tiers kept the side of the base they took in issue #6. Seven issues in, the margin curve has taken seven readings.

Daily-horizon herds: 1d/4h unanimity above the base, 1d/1d below

The two daily cells stay the report's thinnest. 1d/1d unanimity hit 38.4% (n=112) and 1d/4h unanimity 53.8% (n=91), split around the 50.7% field base; the whole 1d/1d agree bucket ran 35.6% (n=146) -- not the highest of the five FH/TF cells -- while 1d/4h ran 31.6% (n=155), the second-thinnest agree bucket in the table above. Issue #6 printed 0.0% and 14.3% on the same two unanimity cells. Both daily buckets are an order of magnitude thinner than the hourly one (n=2,952), so they are reported, not ranked.

Herd failures: the 5 worst unanimous misses

When the whole field agreed, and was wrong. This week 539 slot / symbol / FH / TF cells had 6 or more directional models sitting on the same side (issue #6: 417).

Slot (UTC)	Symbol	FH / TF	Side	Models	Mean conf	Hit rate
Sep 18, 20:01	XRP	4h / 4h	long	7	70.0	0.0%
Sep 16, 09:01	BNB	1h / 1h	short	7	69.6	0.0%
Sep 14, 19:01	BTC	1h / 1h	long	7	68.6	0.0%
Sep 19, 04:01	XRP	4h / 4h	long	7	68.4	0.0%
Sep 20, 00:01	BTC	1d / 1d	long	7	68.4	0.0%

The 5 worst cells were long/short swarms at 68.4-70.0 mean confidence and resolved 0.0%-0.0%. 'Sideways' was 39.9% of mature forecasts (3,710 of 9,308) this week.

Per-model: agree vs. disagree

Per-model disagree counts run 0 to 30 this week (issue #6: 1 to 12). 2 of the 7 lines clear the N=10 floor (deepseek-v4-pro 12, gemini-3.1-pro 30); every other per-model contrarian cell below is insufficient.

Model	Agree n	Agree hit	Disagree n	Disagree hit
claude-fable-5	713	51.3%	0	n/a
claude-opus-5	659	51.4%	1	0.0%
deepseek-v4-pro	736	49.0%	12	41.7%
gemini-3.1-pro	695	52.5%	30	43.3%
gpt-5.6-sol	720	50.8%	2	50.0%
grok-4.6	577	49.2%	6	16.7%
qwen-3.8-max	725	49.7%	1	0.0%

Individual per-model contrarian records are anecdotes, not findings: at these counts a handful of flipped calls swings a 'lift' by tens of points, so models are never ranked by contrarian lift here. The grok 4.5 -> 4.6 flip (Aug 24, 2026 09:22 UTC) sits three windows back (issue #4): both weeks of this issue are pure grok-4.6, with no 4.5-era rows in either window, so the grok week-over-week row is the second pure-4.6 against pure-4.6 comparison of the series.

Market check: an up week on a louder tape

OBSERVATION — mean [1d move] 1.35% -> 2.69% (+100% rel), BTC realized vol 31.4% -> 35.1% (ann., hourly); field directional accuracy 39.6% -> 50.7% (+11.2pp), 7 of 7 models improved, sim win-rate up for 7 of 7, field sim PnL -\$1018.93 -> -\$337.72 (adjacent calendar weeks Sep 7-13 vs Sep 14-20; descriptive, one pair of weeks, not a claim).

Robustness: raw price-sign accuracy 38.7% -> 54.2% — the same direction as the trade-based hit rule.

Measure	Sep 7-13 (week A)	Sep 14-20 (week B)	Change
Mean [1d move] (5 assets)	1.35%	2.69%	+100% rel
BTC realized vol (ann., hourly)	31.4%	35.1%	+3.7pp
Field directional accuracy	39.6%	50.7%	+11.2pp
Models improving hit-rate	--	7 of 7	--
Raw price-sign accuracy	38.7%	54.2%	+15.4pp
Field sim win-rate	32.1%	43.4%	+11.2pp
Field sim net PnL	-\$1018.93	-\$337.72	--

Week A is exactly the issue-#6 window, so this slice and every week-over-week column in this report compare the same two back-to-back weeks. Descriptive, one pair of weeks; 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 splice ACTIVE: grok-4.5 rows are aggregated into grok-4.6 (flip 2026-08-24T09:22:00+00:00, three windows back, inside issue #4's window). Week A (2026-09-07..2026-09-13) is pure grok-4.6 (it reproduces the published issue #6); week B is pure grok-4.6.

Practical implications

- Seven issues, seven margin-curve readings. The tier ordering moved again -- 4 peers still on top, but unanimity came off the bottom and the 5-peer tier took it -- and all four tiers printed a higher rate than in issue #6 (4 peers 46.9% -> 55.5%, full unanimity 34.2% -> 51.6%), so a consensus filter still should not carry weight in a config.
- The pooled lift is +12.1pp this week against -30.4pp in issue #6 and -27.6pp in issue #5. The sign flipped after three issues below zero: treat the number as this week's weather, not an edge.
- Read the consensus number together with the market check on the same page. The field base moved +11.2pp week over week; 4 of the 4 consensus tiers moved up and 0 down against issue #6 (3 peers +9.4pp, 4 peers +8.6pp, 5 peers +3.9pp, unanimity +17.3pp).

Limitations

- Single week (Sep 14-20, 2026 UTC). Every finding is descriptive for this window only.
- Observations are not independent: the same models watch overlapping symbol / FH / TF cells hour after hour. Wilson intervals here are descriptive, not inferential.
- Correlational, not causal: all models see the same market data, so 'agreement' and 'hit' can rise together simply because a slot was easy to read. LOO removes self-match bias, not this confound.
- The pooled lift rests on 52 disagreeing calls; per-model counts run 0-30. 721 calls (12.9% of the 5,598-call mature directional universe) were thin or tied and are excluded from every consensus table.
- This window ran 3 trend days of 7 (issue #6: 2 of 7); every comparison with issue #6 is a comparison across regimes as well as across weeks. Series density and lineage notes (wave 7): (1) rolling-1w forecasts emit daily since Aug 22, 2026 and rolling-1M daily since Aug 23, 2026, so inside this window (Sep 14-20) daily coverage is FULL, as it was in issues #4, #5 and #6: the 1w series has slots on 7 of 7 days and the 1M series on 7 of 7 days; both sit outside this report's FH gate (1h/4h/1d) and appear only in the exclusions counter (1w 488, 1M 485). (2) The grok line flipped 4.5 -> 4.6 at Aug 24, 2026 09:22 UTC, three windows back (inside the issue-#4 window): this window carries 1,470 grok rows and none from the 4.5 era, and neither does week A (the issue-#6 window), so every grok number in this issue and in issue #6 is a pure grok-4.6 line -- the grok week-over-week row is the second pure-4.6 against pure-4.6 comparison of the series. (3) qwen-3.8-max succeeded qwen-3.7-max on Aug 5, 2026 (fully before this window).
- Market-state row is computed on a single-exchange (binance) daily candle series, as in issue #6 after the audit that found the raw table mixes two exchanges; issue #6 values are as published.
- Research-to-date counter: directional forecasts resolved inside the FH gate (1h/4h/1d) since Jul 11, counted at the issue cutoff (Sep 21 16:00 UTC) -- the definition pinned in issue #3 and carried by issues #4, #5 and #6, which printed 47,373 at the Sep 14 16:00 cutoff. The pack reproduces that pin at the previous cutoff (control OK), so this issue's 52,971 is an additive step under one definition; counter deltas against issue #5 and earlier remain definitional.

Counters & lineage

Counter	Value	Counter	Value
Forecasts total	10,290	Mature (scored pool)	9,308
OK in gate	9,308	-- of them directional	5,598
Out of gate (1w / 1M)	488 / 485	-- of them sideways	3,710
Invalid	9	Pending (next issue)	0
Late closes	0	Uptime, all grids	266 / 266 slots
Source file	weekly_metrics_2026-09-14.json	Report cutoff	Mon Sep 21, 2026, 16:00 UTC
Generated at (pipeline)	2026-09-22 06:56 UTC	Methodology	v1.1 (2026-08-10) · hash e66c7e8c864a2233

THIS REPORT: 4,877 LOO-scored observations. MARKETMANIA RESEARCH TO DATE (as of Sep 21, 2026 cutoff): 52,971 directional forecasts resolved since Jul 11 · 7 models tracked (current line-up; earlier versions folded into their successors' lineage) · 5 assets · 5 horizons · hourly · 31 published reports.

Issue #7. Consensus Watch is a living weekly comparison; each issue appends one more week of leave-one-out agreement-vs-hit data. Issue #6 asked whether the pooled agree-vs-disagree sign keeps the same direction for a fourth issue running: it reads +12.1pp this week after -30.4pp in issue #6 (agree 50.6% vs disagree 38.5% on n=52), and whether the same margin tiers hold their side of the field base for a fourth issue: 2 of the 4 tiers did, and the 4-peer tier came in above the field base. The grok row is pure grok-4.6 this issue and was pure grok-4.6 in issue #6; the 4.5 -> 4.6 flip (Aug 24, 2026 09:22 UTC) sits three windows back. Engine 1.1 has powered the sandbox since Aug 18, i.e. before this window; no cross-engine PnL comparisons are claimed.

What we're testing next

- Next issue: the pooled agree-vs-disagree lift flipped to +12.1pp after three issues below zero -- does the positive sign hold for a second issue?
- The 4-peer tier has sat above the field base two issues running -- does it hold for a third?
- Monthly series: the agreement curve across regimes at monthly n.

Related research

Report	Direct PDF link
Weekly Calibration #7	https://marketmania.ai/research/reports/weekly-calibration-2026-09-14.pdf
Weekly Model Watch #7	https://marketmania.ai/research/reports/model-watch-2026-09-14.pdf
Consensus Watch #6	https://marketmania.ai/research/reports/consensus-watch-2026-09-07.pdf

The three weekly reports publish together as one issue each week; each links straight to the others' PDF and to its own previous issue. Direct links are the posting rule from wave 2 on.

CITE THIS REPORT

```
@misc{mm_consensus_2026w38,
  title = {Consensus Watch #7: does agreeing with the crowd make an LLM's market call safer?},
  author = {{MarketMania Research}},
  year = {2026}, month = {September}, day = {22},
  url = {https://marketmania.ai/research/reports/consensus-watch-2026-09-14.pdf},
  note = {Methodology v1.1; window Sep 14-20, 2026 UTC; source weekly_metrics_2026-09-14.json}
}
```