wallstreetprompt>_

# Market Positioning Consensus Options Positioning and Implied Volatility, 2026 Year to Date

An Institutional Research Report

**Felipe Sinisterra**

Wall Street Prompt

August 2026

# The Read

Where does US equity positioning sit as of August 28, 2026? The study reads 165 trading sessions of listed-options data and blends three dials: what the crowd is actually buying (SPEC), what protection costs (HEDGE), and what market makers are forced to do (DEALER), weighted 0.45/0.35/0.20 into a Positioning Consensus Index (PCI), scored −100 to +100. Each reading is ranked against every trading day of 2026 so far (the 2026 rank) and against the longest clean history each component has (the multi-year rank).

Bullish consensus, at the overly-bullish boundary, medium confidence. For the first time this year all three dials lean the same way on the 5-day averages: the crowd is buying calls, dealer positioning has flipped supportive, and insurance is cheap. Two mechanics temper the label. The boundary call fires on one leg only: protection cost (HEDGE) clears its gate, the 80th percentile of the year, at the 82nd, a 1.9-point margin (about 3 rank positions in 165 sessions) that dropping any one of three of its seven components erases; the other leg, amplification risk, sits far below its own gate at the 33rd percentile. Confidence is medium: the sign rule on the three 5-day averages counts zero dissent, but the DEALER average is +2.5 on a −100 to +100 scale (0.07 of a daily standard deviation) and the as-of day itself printed two of three dials negative (SPEC −3.1, DEALER −36.0). One ordinary dealer session flips the count.

1. **PCI: 85th percentile of 2026, 61st against the longer view.** The 5-day average is +21.0, higher than about 85 of every 100 trading days this year. On the multi-year convention it ranks 61st, where the same rule would read neutral; that figure is a component-averaging convention rather than a true multi-year rank, since the blend has no pre-2026 history. The state is one week old: the rank swung from the 37th to the 85th percentile in six sessions, first crossed the bullish gate on August 26, and the as-of day itself printed +1.6, the 53rd percentile. The year as a whole shows no trend (Mann-Kendall test, p = 0.50), and the study's own regime table labels July 28 to August 28 "neutral, flat" (mean +7.7). What changed is the last week, not the whole tape.
2. **SPEC: the crowd turned buyer, fast.** Speculative demand faded through mid-August, then jumped from the 21st to the 77th percentile of the year over the final six sessions (5-day average +21.3; 60th multi-year). The year's peak remains +61.1 on August 4.
3. **DEALER: flipped supportive, thinly.** The dealer-direction 5-day average is +2.5, the 51st percentile of 2026 and the newest, thinnest leg of the consensus: the as-of day printed −36.0, the week's most negative. When market makers are net short gamma, their hedging sells as the market falls and buys as it rises, amplifying moves in both directions; the study's gauge of that amplification risk fell from the 73rd to the 33rd percentile of the year in six sessions. Even so, market makers remain net short gamma, as on 116 of 165 sessions this year (the norm, not an elevation), with SPY 0.06% below the 769.6 level where hedging would flip to dampening.
4. **HEDGE: insurance is cheap, and it is the label's weak spot.** The hedging score averages

+31.4 over 5 days, the 82nd percentile of 2026 (85th multi-year). Price levels carry it: VIX option prices sit at 5 on their 0-to-100 one-year range (0 = cheapest), the 1st percentile of 2026, and the other two level components score +74 and +89. Both demand components dissent (scores −55 and −48): the tilt of put prices over call prices is near its most put-heavy reading of the year. Cheap insurance levels, not absent put demand, make this dial bullish.

5. **Validation: days like this leaned the wrong way in-sample.** Only 14 of 165 sessions in 2026 classify bullish; the earlier clusters sat at the late-January top that preceded a 9.1% drawdown and at the early-August peak that faded, and SPY averaged −0.13% over the 10 days after a bullish-class day (positive 18% of the time; n = 11 overlapping windows, descriptive only). The quintile test points the same way: the most crowded-bullish fifth of days preceded SPY returns 2.4 points weaker over 21 days (4.8 over 63) than the most bearish fifth, but the gap is driven by one drawdown episode, beaten by a simple momentum benchmark (−5.3 points at 21 days), and drawn from a single year. Caution, not a forecast.

Three dependencies travel with the verdict. It leans on the insurance dial: drop HEDGE entirely and the composite falls to the 75th percentile, inside the rule's neutral band (0.20 to 0.80), while the other 20 leave-one-out variants keep the bullish class. One insurance component changes definition inside the scored week (the options-versus-realized volatility gap; disclosed in Chapter 6). And every number in Chapters 1 to 12 comes from a single options-data vendor with no in-sample cross-source check; Chapter 13 is the only outside-vendor context. Research, not investment advice.

# Contents


**The Read** 1

**1 Executive Summary** 4
- 1.1 The 2026 positioning dashboard . . . 5
- 1.2 Three dials, one direction . . . 6
- 1.3 Trajectory: a hot week inside a flat year . . . 7
- 1.4 How the report is organized . . . 7

**2 Method and Data Contract** 8
- 2.1 One vendor, one exception . . . 8
- 2.2 Universe . . . 8
- 2.3 Windows and the two ranking frames . . . 9
- 2.4 The three sub-indices and their components . . . 10
- 2.5 Statistical standards and power . . . 11
- 2.6 Endpoint inventory . . . 12
- 2.7 Amendments and wording corrections . . . 12

**3 The Measurement Problem** 13
- 3.1 Why the raw put/call ratio misleads . . . 13
- 3.2 Contracts versus dollars: the S4 gap . . . 14
- 3.3 The structural break audit . . . 15
  - 3.3.1 Actions taken . . . 16

**4 SPEC: Speculative Demand** 18
- 4.1 Construction . . . 18
- 4.2 Component walk . . . 19
- 4.3 Redundancy and known pathologies . . . 20
- 4.4 The 2026 trajectory . . . 20
- 4.5 Where SPEC stands on 2026-08-28 . . . 22

**5 DEALER: What Market Makers Are Forced to Do** 24
- 5.1 Components and construction . . . 24
- 5.2 The direction read: from selling to flat . . . 25
- 5.3 Amplification risk: still the 2026 norm, now at the boundary . . . 26
- 5.4 The strike map on 2026-08-28 . . . 28
- 5.5 Vanna and charm: shown, not scored . . . 28
- 5.6 What enters the verdict . . . 29

**6 HEDGE: The Price of Protection** 30

6.1 Construction and current levels . . . 30
6.2 The year in insurance . . . 31
6.3 The raw inputs: skew, risk premium, term structure . . . 32
6.3.1 H2: the risk reversal stayed positive all year, and is now near its low . . . 32
6.3.2 H3: the volatility risk premium is back to flat, and the series has a seam . . . 33
6.3.3 H4: term structure in firm contango . . . 33
6.4 The complacency read . . . 34
6.5 QQQ and IWM: the reduced-form check . . . 35

**7 The Composite: Positioning Consensus Index 37**
7.1 Construction . . . 37
7.2 Level and range . . . 38
7.3 Monthly attribution . . . 39
7.4 Regimes . . . 39
7.5 Chronology of extremes . . . 40
7.6 Trajectory: trend test . . . 41

**8 Where We Are Versus 2026 43**
8.1 The direct answer: hot for 2026, ordinary for the longer history . . . 43
8.2 The year ranked: notable dates by PCI . . . 43
8.3 The August pattern: a fade, then a six-session surge . . . 44
8.4 Has positioning gotten more bullish? . . . 45

**9 Cross-Section: Names and Sectors 48**
9.1 How each name is scored . . . 48
9.2 The full ranking: four names at the year's highs . . . 49
9.3 A re-risking with breadth . . . 50
9.4 Sector flows: no extremes . . . 51
9.5 Concentration: broad, not narrow . . . 53

**10 Semiconductors and Memory 55**
10.1 How each name is scored . . . 55
10.2 The ranking: crowded supply chain, uncrowded NVDA . . . 56
10.3 Memory versus AI logic versus equipment . . . 57
10.4 What protection costs in semis . . . 58
10.5 Extremes and the takeaway . . . 59

**11 Short-Dated Versus Long-Dated: Positioning by Time to Expiry 60**
11.1 The layered design and delivered coverage . . . 60
11.2 Tenor windows in the delivered series . . . 60
11.3 What agrees and disagrees on 2026-08-28 . . . 62

**12 Validation: Stress-Testing the Composite 64**

12.1 Coherence: three dials, not one . . . 64
12.2 Forward returns by composite quintile . . . 65
12.3 Placebo, benchmark, and the one-episode problem . . . 67
12.4 Break robustness and leave-one-out . . . 69
12.5 Power and single-vendor exposure . . . 70

**13 Long-History Context: A Non-UW Cross-Check 71**
13.1 What a long history can and cannot check . . . 71
13.2 Today against 36 years of insurance pricing . . . 71
13.3 The closest analogues since 2007 . . . 73
13.4 The same test on 26 years of data, 2000 to 2026 . . . 74

**14 The Verdict 77**
14.1 The rule and its inputs . . . 77
14.2 The class is real, and it is young . . . 78
14.3 Why "overly" is only a boundary call . . . 78
14.4 Attribution: what makes the read . . . 79
14.5 What followed readings like this in 2026 . . . 80
14.6 Why the printed confidence is medium . . . 80
14.7 What this read is, and is not . . . 81

**15 What Would Change the Read 82**
15.1 Falsification triggers . . . 82
15.2 The single series to watch: HEDGE . . . 83
15.3 Monitoring checklist . . . 84

**A Data Tables 86**
A.1 Monthly composite summary . . . 86
A.2 Latest component readings . . . 86
A.2.1 SPEC components (M1) . . . 87
A.2.2 DEALER components (M2) . . . 87
A.2.3 HEDGE components (M3) . . . 87
A.2.4 Composite (M4) . . . 87
A.3 Regime segmentation . . . 87

**B Endpoint Inventory 91**
B.1 Probe results . . . 91
B.2 Gated, failed, and excluded paths . . . 93
B.3 Call budget . . . 93

**C Reproducibility and Verification 95**
C.1 Run configuration and modules . . . 95
C.2 Blind QC re-derivation . . . 96

C.3 Sign-convention audit . . . . . 96
C.4 Adversarial review, round 1 (the retired August 20 verdict) . . . . . 97
C.5 Adversarial review, round 2 (the refreshed verdict) . . . . . 97
C.6 Amendment log . . . . . 99

# 1. Executive Summary

This study asks one question: is the US equity market overly bullish, neutral, or overly bearish right now, and how did positioning get here across 2026. The answer matters because positioning tells you who is already leaning which way, and how hard a surprise would push the tape. This chapter gives the answer in one page and sets up the dashboard the rest of the report drills into.

The instrument is the Positioning Consensus Index (PCI), a daily score built from three dials. SPEC (weight 0.45) tracks what the crowd is actually buying in the options market. HEDGE (weight 0.35) tracks what protection costs; it is scored so that cheap, unbid insurance counts as bullish, because a market that will not pay for crash protection is a complacent one. DEALER (weight 0.20) tracks what market makers are forced to do. Each dial and the composite run from $-100$ (extreme bearish consensus) to +100 (extreme bullish consensus). Every reading is ranked two ways: against every trading day of 2026 so far (the 2026 rank; 165 sessions, 2026-01-02 to 2026-08-28) and against the longest clean history each series has (the multi-year rank). All positioning inputs come from a single vendor, Unusual Whales, as of the 2026-08-28 close; Chapter 2 fixes that boundary and what it implies.

The answer, developed in Chapters 8 and 14: bullish consensus, and by the study's own pre-registered rule the reading sits at the overly-bullish boundary. The 5-day average PCI is +21.0, the 85th percentile of 2026, meaning higher than roughly 8 of every 10 trading days this year. The crowd is buying calls, dealer positioning has turned supportive, and insurance is cheap and unbid. Four qualifiers travel with that read, and they are load-bearing.

1. **Hot against 2026, ordinary against the longer view.** On the multi-year rank the same 5-day average sits at the 61st percentile, and the same rule applied there reads neutral. The multi-year composite rank is also an averaging convention across components, not a true rank on a multi-year composite path (Chapter 2), so the 2026 rank is the sharper of the two claims.
2. **It is one week old.** The rule's input, the 5-day average's 2026 rank, jumped from the 37th to the 85th percentile in the final six sessions of August, mostly because speculative demand jumped from the 21st to the 77th. The classification itself is three sessions old, and the as-of day closed soft: the 2026-08-28 daily PCI is +1.6, the 53rd percentile.
3. **Bullish days are rare this year, and their record is poor.** Only 14 of 165 sessions classify bullish under the rule, and the earlier bullish clusters came just before the April drawdown and the mid-August fade. Across the 11 prior bullish days with ten forward sessions of data, SPY's forward 10-day return averaged $-0.13\%$ with gains 18% of the time. That is descriptive, not a forecast.
4. **The year as a whole is still trendless.** A standard statistical trend test puts 2026 at $p = 0.496$, no trend, and the study's own regime segmentation labels the whole July 28 to August 28 stretch "neutral, flat" with a mean PCI of +7.7. What changed is the last week, not the tape.

The "overly bullish" escalation gets its own scrutiny because it fires on one leg by a small margin. The rule escalates when either of two gates trips: insurance complacency (the HEDGE 2026 rank) above 0.80, or

amplification risk above 0.70. The amplification-risk leg is dead at 0.33. The insurance leg fires at 0.8193, a margin of 1.9 percentage points, about three rank positions. Confidence in the direction call is medium. The mechanical sign rule prints high, because all three dials share a positive 5-day average sign, but the DEALER 5-day average clears zero by only +2.5, about 0.07 of a daily standard deviation, and the as-of day's own prints split negative (SPEC −3.1, HEDGE +29.1, DEALER −36.0). One ordinary dealer session flips the sign count, so the daily-sign reading is already medium.

## 1.1 The 2026 positioning dashboard

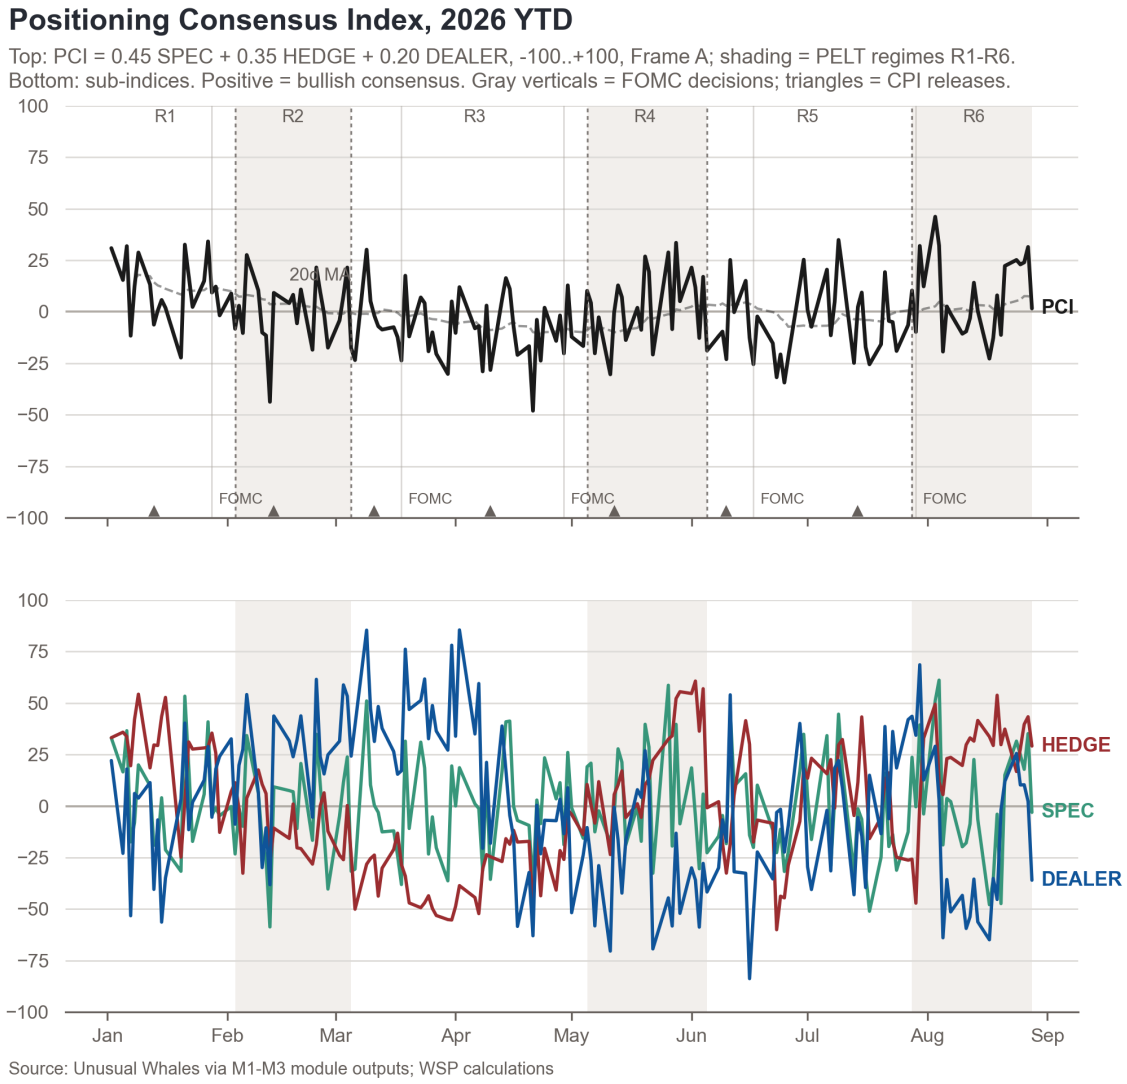


Exhibit 1.1 shows the daily PCI path, the regime shading, the three dials, and the latest attribution; Table 1.1 compresses the state into numbers.

| Series | Weight | 5-day avg | 2026 rank | Multi-year rank | YTD peak (date) | YTD trough (date) |
|---|---|---|---|---|---|---|
| SPEC | 0.45 | +21.3 | 77% | 60% | +61.1 (Aug 4) | −58.7 (Feb 12) |
| HEDGE | 0.35 | +31.4 | 82% | 85% | +60.7 (Jun 2) | −60.0 (Jun 23) |
| DEALER (dir) | 0.20 | +2.5 | 51% | 55% | +85.5 (Apr 2) | −83.8 (Jun 16) |
| PCI | 1.00 | +21.0 | 85% | 61% | +46.2 (Aug 3) | −48.1 (Apr 21) |

Table 1.1. Scoreboard as of 2026-08-28; scores run −100 (extreme bearish consensus) to +100 (extreme bullish consensus). Peaks and troughs are daily 2026 values; HEDGE is the SPY-based market series. Multi-year rank conventions differ by dial, and DEALER's multi-year rank mixes full-depth delta data with 2026-only flow inputs; treat cross-frame comparisons as indicative (Chapter 2). (Source: m1.json, m2.json, m3.json, m4.json headline_numbers and series)

## 1.2 Three dials, one direction

For most of 2026 the dials disagreed with each other. On the latest 5-day averages they line up positive, and each has its own story.

**What the crowd is buying.** Speculative demand printed a −47 day on August 20, then ran five straight positive sessions, peaking at +35 on August 27. The 5-day average of +21.3 ranks in the 77th percentile of 2026, up from the 21st just six sessions earlier. That swing, not a slow build, is what moved the composite.

**What protection costs.** Insurance is cheap and nobody is bidding for it. The clearest single reading: options on the VIX, the market's crash gauge, are priced at 5.4 on a 0-to-100 scale of their own one-year range (0 = cheapest), and that reading sits at the 1st percentile of 2026's daily prints. A positive HEDGE score means the market is not paying up for crash protection; at +31.4 it is the strongest of the three dials.

**What market makers are forced to do.** Net gamma, a measure of whether market makers' hedging dampens market moves (positive) or amplifies them (negative), is still negative on SPY, as it has been for 116 of 165 sessions this year; negative is the 2026 norm, not an alarm. But the tape now sits essentially at the switch line: SPY closed 0.06% below the price level where dealer hedging flips from dampening moves to amplifying them (769.62). Amplification risk, the gauge of how hard that hedging would push a move in either direction once it starts, collapsed in the final week, from the 73rd percentile of the year on August 20 to the 33rd on August 28. Dealer positioning went from a headwind to roughly neutral, with the direction dial at +2.5, the 51st percentile.

The overly-bullish call is a boundary call, and the fine print matters. The escalation gate is an either-or; today only the insurance leg fires, at 0.8193 against a 0.80 gate, a 1.9-point margin worth about three rank positions or 1.7 score points, while the amplification-risk leg sits at 0.33, nowhere near its 0.70 gate. The firing leg is carried by components that measure where volatility prices sit (scores +74.0, +92.5, +88.7), while both components that track actual paid-for protection demand point the other way (−55.2, −47.7). Drop any one of three of the seven components and the gate un-fires (reranked inputs 0.795, 0.735, 0.753 against 0.80). On the as-of day alone the insurance dial ranks at the 77th percentile, below the gate, and readings at this level have kept the gate lit into the next session about

42% of the time this year. Chapter 15 tracks the same asymmetry going forward.

## 1.3 Trajectory: a hot week inside a flat year

Positioning has not trended across 2026. A standard statistical trend test (Mann-Kendall, with the Sen slope as the size estimate) puts the year at $z = -0.681$, two-sided $p = 0.496$, with a slope of $-0.02$ score points per session, indistinguishable from flat. An algorithm that finds the dates where a series shifts to a new level (PELT) cuts the year into 6 regimes; every one averages inside the neutral band, with regime means spanning $-7.9$ to +11.0, five flat and one rising. The year oscillates between turning points instead: a +34.2 peak on January 27, troughs of $-30.2$ on March 30, $-48.1$ on April 21, and $-34.4$ on June 25, then the +46.2 YTD peak on August 3 and a +31.5 peak on August 27. The latest local signal points up: the rolling 20-day slope is +0.44, the 72nd percentile of the year's slopes, driven entirely by the last week of August. That is the tension Chapter 8 resolves: a flat year, a hot week, and a verdict rule that reads the week.

## 1.4 How the report is organized

Chapter 2 fixes the method and the data contract. Chapter 3 documents why raw put/call measures are broken and the break audit that gates every series. Chapters 4 to 6 walk the three dials: what the crowd buys, what market makers are forced to do, and what protection costs. Chapter 7 assembles the composite and its regimes, and Chapter 8 gives the headline answer against the full 2026 tape. Chapter 9 maps the cross-section of names and sectors, Chapter 10 goes inside the semiconductor and memory complex, and Chapter 11 layers the read by time to expiry. Chapter 12 stress-tests coherence, forward-return behavior, and leave-one-out sensitivity. Chapter 13 adds non-vendor context back to 1990, Chapter 14 states the verdict with attribution and base rates, and Chapter 15 lists the falsification triggers that would change the read.

# 2. Method and Data Contract

This chapter explains how the study was built: what data it uses, how a raw reading becomes a percentile and a score, and which rules were fixed before any answer was computed. When a later chapter cites a percentile, this is where its meaning is pinned down.

The study runs against a pre-registered contract (`RESEARCH_PLAN.md`, run date 2026-08-20, seed 20260820 for every stochastic step) that fixed the universe, windows, component definitions and signs, blend weights, and verdict rule before any composite was computed. Under Amendment A4 the dataset was refreshed through the 2026-08-28 close and every module re-ran under the same contract and seed; every number in this report is the refreshed vintage. Changes forced by the data are logged in Section 2.7.

## 2.1 One vendor, one exception

Every number in Chapters 1 to 12 derives from a single vendor, Unusual Whales (UW), harvested through the REST API into a local parquet lake and processed offline. Chapter 13 is the sole exception: it builds a long-history proxy from non-UW sources (FRED, Cboe-derived series, yfinance closes) to place 2026 against decades rather than years, and carries a banner label to that effect. The two source sets are never mixed.

Single-vendor exposure. Chapters 1 to 12 have no in-sample cross-check from a second source; if the vendor's collection pipeline shifts, the percentiles shift with it. Two known level-shifts sit inside the 2026 window: QQQ put premium stepped to roughly double its prior typical level on 2026-01-28 (a 2.08x jump), and QQQ call delta did the same on 2026-04-09 (2.04x). The multi-year frame cuts history at those breaks; the 2026 ranks for the affected QQQ components knowingly span them. Separately, one insurance-cost input (H3, the gap between what options charge for future volatility and the volatility that actually shows up) splices two definitions: the official series needs the next month of realized volatility, so it ends 2026-07-31, and a trailing 20-day version carries it from 2026-08-03. The most recent sessions ride that extension; the seam is continuity-checked in Chapter 6. These exposures carry into the confidence discussion in Chapter 14. (Source: results/mod0_breaks.json; results/m3.json; results/probe_endpoints.json)

## 2.2 Universe

The market-level read weights the four equity index tickers by capitalization tier: SPY 0.40, QQQ 0.30, IWM 0.15, DIA 0.15. Sectors and mega-caps are reported separately and never enter the headline composite. The dedicated semis universe added under Amendment A5 lives entirely in Chapter 10.

| Group | Members | Role |
|---|---|---|
| Index complex (5) | SPY, QQQ, IWM, DIA, VIX | Headline market read; VIX = options on VIX, VIX spot via the iv-rank `close` field |
| Sector ETFs (11) | XLK, XLF, XLV, XLE, XLI, XLY, XLP, XLU, XLB, XLRE, XLC | Cross-section (Chapter 9) |
| Mega-caps (7) | NVDA, AAPL, MSFT, AMZN, GOOGL, META, TSLA | Cross-section and concentration (Chapter 9) |
| GICS sector tides (11) | Basic Materials, Communication Services, Consumer Cyclical, Consumer Defensive, Energy, Financial Services, Healthcare, Industrials, Real Estate, Technology, Utilities | Intraday sector-tide chronology |

## 2.3 Windows and the two ranking frames

Two ranking frames run in parallel. Frame A, the headline frame, ranks a reading against every trading day of 2026 so far: 2026-01-02 to 2026-08-28, 165 sessions, confirmed by every module's coverage check. The report calls this the 2026 rank. Frame B ranks the same reading against the full multi-year history each series supports after the break audit of Chapter 3: the multi-year rank. It never appears without the 2026 rank beside it, and the 2026 rank leads every citation. Percentiles are literal: the 85th percentile of 2026 means higher than about 8.5 of every 10 trading days this year.

| Family | Frame B start | Rows | Note |
|---|---|---|---|
| Greek exposure: delta, gamma (5Y) | 2022-03-30 | 1,107 | Deepest family; vanna/charm shallower (2023-06-09, 806 rows), break-flagged |
| Options volume, total options volume, short volume | 2024-08-30 | 500 | Vendor cap: rolling 500-session window ending 2026-08-28; requests above 500 are rejected |
| IV rank, realized volatility | 2025-08-29 | 251 | Fixed one-year vendor span |
| Variance risk premium (official) | 2025-08-29 | 231 | Official series ends 2026-07-31; trailing 20-day extension after |
| Risk-reversal splice | 2025-09-29 | | Usable splice start per Module 3 |
| Market tide, NOPE, greek flow, interpolated IV | YTD only | | Per-date harvests; Frame B equals Frame A |

(Source: results/mod0_breaks.json; results/m1.json to m3.json notes; results/probe_endpoints.json)

Every component maps its raw value to a percentile within the frame (rank$/(n+1)$), then to a score $= (\mathrm{pct} - 0.5) \times 200$, from $-100$ to $+100$, positive meaning bullish consensus. Z-scored inputs are capped at 3 standard deviations before averaging, so one wild print cannot dominate a blend.

One caveat matters when reading the verdict. The blended composite has no pre-2026 history: no earlier date carries all of its components. Its multi-year number is therefore a convention, not a true rank: Modules 1, 4, and 5 average the components' multi-year ranks (score$/200 + 0.5$) rather than ranking the composite in a multi-year distribution that does not exist. The DEALER blend's multi-year number mixes one component with history to 2022 (D1) and two that exist only for 2026 (D4, D5), and modules differ in the convention their sub-index headline uses. Cross-module comparisons of multi-year numbers are therefore not made; every cross-module statement runs on the 2026 rank, where the convention is uniform.

## 2.4 The three sub-indices and their components

The Positioning Consensus Index (PCI) blends three sub-indices, computed daily across the 165-session 2026 window:

$$\mathrm{PCI}_t = 0.45 \times \mathrm{SPEC}_t + 0.35 \times \mathrm{HEDGE}_t + 0.20 \times \mathrm{DEALER_dir}_t$$

In human terms: SPEC measures what the crowd is actually buying in the options market. HEDGE measures what protection costs, inverted: cheap, unbid insurance reads as complacency, which counts as bullish consensus. DEALER_dir measures the directional demand customers push onto market makers, inferred from dealer books: what market makers are forced to do. A fourth axis, amplification risk, stays out of the blend. When market makers are net short gamma, their hedging forces them to sell as the market falls and buy as it rises, amplifying moves in both directions. That axis says how hard the tape could move, not which way; it feeds only the "overly" qualifier of the verdict rule in Chapter 14, never the PCI. Within each sub-index, components are equal-weighted.

| Id | What it measures | Sign / role |
|---|---|---|
| **SPEC, weight 0.45: what the crowd is buying. Equal-weight S1, S2, S3, S5, S6, S7, S8, S9** | | |
| S1 | Net premium tilt: how far option dollars lean bullish, (net call − net put premium) / (call + put premium) | + |
| S2 | Premium call/put ratio: call dollars versus put dollars | + |
| S3 | Volume call/put ratio: call contracts versus put contracts | + |
| S4 | Ratio gap $z_A(\mathrm{S3}) - z_A(\mathrm{S2})$: do contracts and dollars disagree | diagnostic, unsigned, not blended |
| S5 | Bullish premium share of all premium | + |
| S6 | Call ask-side share: call volume that hit the seller's price (buyer-initiated) | + |
| S7 | Put ask-side share: buyers paying up for puts (protection lifted at the ask) | − |
| S8 | Out-of-the-money (OTM) net-premium tilt from market tide, EOD cumulative (wording amended; see Section 2.7) | + |
| S9 | Total volume versus its 30-day average, signed by S1's direction | carries S1's sign |
| **DEALER, weight 0.20: what market makers are forced to do. Direction blend of D1, D4, D5** | | |
| D1 | Customer delta demand: directional exposure customers push onto dealers, −(call delta + put delta) | +, in blend |
| D2 | Net gamma: whether dealer hedging dampens market moves (positive) or amplifies them (negative), call gamma + put gamma | amplification-risk axis, not in PCI |
| D3 | Gamma-flip distance: how far spot sits from the level where dealer hedging switches from dampening moves to amplifying them, (spot − flip)/spot | amplification risk |
| D4 | NOPE daily close: how much option trading is pushing around the stock itself | +, in blend |
| D5 | Directional delta flow, daily sum | +, in blend |
| D6 | Vanna/charm: second-order dealer hedging exposures | excluded from scores where break-flagged; shown descriptively |
| **HEDGE, weight 0.35: what protection costs, inverted. Equal-weight H1 to H7** | | |
| H1 | IV rank: where today's option prices sit in their one-year range (0 = cheapest, 100 = most expensive) | − |

| Id | What it measures | Sign / role |
|---|---|---|
| H2 | 25-delta risk reversal: the price of upside calls versus equally out-of-the-money downside puts; when puts cost more, the market pays up for crash protection. 30-45 DTE constant-maturity monthly splice | $+$ |
| H3 | Variance risk premium: the gap between what options charge for future volatility and the volatility that actually shows up (official series to 2026-07-31, trailing 20-day extension after) | $-$ |
| H4 | Term slope: whether 90-day options cost more than 7-day options (calm) or less (stress), interpolated IV(90d) $-$ IV(7d) | $+$ |
| H5 | Short-volume ratio (SPY) | $-$ |
| H6 | VIX IV rank: how bid options on VIX itself are (insurance on volatility) | $-$ |
| H7 | 30-day implied move percentile: how big a move options price in for the next month, versus the past year | $-$ |

Two construction details flagged by the adversarial review are disclosed here. First, S1, S5, and S9 are highly correlated views of the same premium tape (correlations in Chapter 4), so the SPEC blend effectively gives about 3 of its 8 slots to one signal, and S9 flips sign whenever S1 crosses zero, so it is jumpy near flat tape. Second, D1's sign rests on the vendor's assumption about which side of each trade dealers hold, which the delivered fields cannot verify. Both temper how much independent evidence the component count suggests.

## 2.5 Statistical standards and power

Three pieces of statistical machinery run under the report.

**Breaks.** Vendor histories sometimes jump to a new level overnight because collection changed, not the market. Chapter 3 audits every input series with PELT, an algorithm that finds the dates where a data series shifts to a new level, run on the log of absolute values (minimum segment 20 sessions, penalty $3 \ln n$) and confirmed with a second test (CUSUM). A series whose typical level jumps 10x or more at a break is HARD: its multi-year history is cut at the break or the series is dropped. A 2x to 10x jump is SOFT: rescaled where the mechanism is understood, otherwise restricted.

**Trend.** Whether 2026 positioning is drifting is judged with a standard statistical trend test (Mann-Kendall, with the Sen slope for size), read at the 5% level, two-sided.

**Forward returns.** Chapter 12 sorts days into five buckets by PCI and compares forward SPY returns at 5, 10, 21, and 63 trading days. Error bars are built by resampling whole stretches of days (blocks of 10, 5,000 draws, seed 20260820), which respects the fact that markets trend; call these resampled error bars. Significance is read only after correcting for the fact that many tests ran at once (BH-FDR at $q = 0.10$), and results are checked against a shuffled-dates placebo (1,000 shuffles) and a naive momentum benchmark.

Power statement. Frame A holds 165 trading sessions. Ranking claims, where today sits within the year's distribution, are well powered. Forecasting claims, whether a positioning bucket predicts forward returns, are not: a 63-day forward window fits inside the year only twice without overlap. They carry resampled error bars and the multiple-testing correction, and read as suggestive, not established.

## 2.6 Endpoint inventory

The vendor surface was probed on 2026-08-20 with 64 calls across 30 recorded endpoint probes before harvesting; harvests were refreshed through the 2026-08-28 close under Amendment A4. Two plan-gated endpoints (403) are excluded: `/api/volatility/vix-term-structure` and `/api/options-pulse/total`. `/api/market/events` returns 404, so the event chronology routes through the vendor's MCP interface (Amendment A2). Historical per-date greeks return zero rows and are excluded. VIX is an index, not a stock, so its OHLC and stock-state endpoints reject; VIX spot comes from the iv-rank `close` field. The full endpoint table is Appendix B. (Source: results/probe_endpoints.json)

## 2.7 Amendments and wording corrections

Five amendments and one QC-forced wording correction were logged during the run.

- **A1 (2026-08-20).** The vendor's daily price history is capped at one year, so Chapter 12's forward-return tests run on the UW price window (2025-08-28 to 2026-08-28 on this vintage) and the 63-day horizon loses the final quarter's tail (last complete window starts late May 2026). The long-history replication moves to Chapter 13 on non-UW SPY closes, explicitly labelled; Chapter 12 remains 100% UW-derived.
- **A2 (2026-08-20).** The market-events REST path 404s; the regime-naming chronology comes from the vendor's MCP interface instead (684 rows saved to the lake).
- **A3 (2026-08-20).** The Cboe equity put/call archive 403s; Chapter 13's proxies are VIXCLS, VXVCLS, and SKEW plus SPY closes. UW's own put/call ratio covers the current tape in Chapters 1 to 12.
- **A4 (2026-08-30).** Full data refresh through the 2026-08-28 close. One-call families were re-pulled fresh; per-date families gained the six sessions 2026-08-21 to 2026-08-28, taking the 2026 window to 165. The break audit and all modules re-ran in dependency order with the same seed 20260820. The report text was also rewritten in plainer language; no analytical definition changed.
- **A5 (2026-08-30).** A semiconductor and memory chapter was added as Chapter 10, backed by new module M7 (14 semis names, roughly 102 additional vendor calls); later chapters shifted up by one. The author of record changed to Felipe Sinisterra.
- **S8 wording correction (QC).** The contract's original S8 wording ("EOD cum net_call_premium share OTM") does not reproduce any headline read literally. As implemented and QC-verified, S8 is the OTM net-premium tilt from EOD cumulative market tide: (OTM net call − OTM net put premium) / |all-strike net call − net put premium|. The wording is amended; no value, percentile, or verdict changes.

# 3. The Measurement Problem

Every claim in this dossier rests on options data, and options data is easy to misread. This chapter covers the three traps: the headline put/call ratio mixes trades that mean different things, contract counts and dollar amounts often tell different stories, and parts of the vendor's own history were recalculated mid-stream. The first two traps shape how Chapter 4 builds SPEC, the sub-index that measures what the crowd is actually buying. The third is handled by a formal audit that decides how far back each data series can be trusted.

## 3.1 Why the raw put/call ratio misleads

The put/call ratio most people quote simply counts contracts: puts traded divided by calls traded. That count treats an option expiring the same day (0DTE) and a two-year protective put as the same unit. It weighs a $0.05 contract the same as a $50 one. And it cannot tell whether a put was bought by someone who wants protection or sold to a market maker by someone who does not. Three specific problems follow.

- **Expiry mixing.** Same-day options are mostly intraday trading: positions opened and closed before the bell that carry no view about next month. Long-dated options carry durable intent. The study therefore splits all volume and open interest (contracts still outstanding) into buckets by days to expiry: 0DTE (same day), SHORT (1 to 28 days), MID (29 to 91), OTHER (92 to 365, reported but kept out of the headline layers), and LEAP (over a year). The module outputs retain no standalone daily series for the 0DTE share, so the contamination point here is structural; Chapter 11 reads the expiry layers side by side and states each layer's coverage limit.
- **Contracts versus dollars.** One trader paying $5 million for index puts and ten thousand traders paying $20 each for short-dated calls send very different signals, but a contract count can score them the same. A count-based ratio and a dollar-weighted ratio measure different populations of traders. The study computes both every day (S3 counts contracts, S2 counts dollars) and tracks their disagreement as a standing diagnostic, S4, examined next.
- **No side information.** Volume alone cannot separate initiated buying from passive selling. The component set therefore uses the share of trades that hit the seller's price, meaning buyer-initiated (S6 for calls, S7 for puts), the vendor's classification of option dollars as bullish or bearish (S5), and a net dollar tilt (S1) to recover direction in dollar terms.

Each component is then expressed two ways: where today ranks against every trading day of 2026 so far (the 2026 rank), and where it ranks against the longest clean history that series supports (the multi-year rank). Ranks are rescaled to scores from $-100$ to $+100$ so direction and intensity read at a glance. Definitions and signs were fixed in the research contract before any data was scored; Chapter 4 walks through the components.

## 3.2 Contracts versus dollars: the S4 gap

S4 asks one question: would the answer change if trades were weighed by dollars instead of counted? To build it, take the contract-count ratio and the dollar ratio, express each as its distance from its own 2026 average in standard deviations (a z-score), and subtract the dollar reading from the count reading. A wide gap in either direction means the two ratios are telling different stories. S4 never enters the SPEC blend; it exists to show how much the answer depends on which ratio one chooses to read.

**Contracts vs dollars: where volume and premium C/P ratios disagree**

Winsorized YTD z-scores of S3 (volume call/put) and S2 (premium call/put), market-weighted;
S4 diagnostic = z(S3) - z(S2): positive = small-ticket call buying outpacing dollar commitment

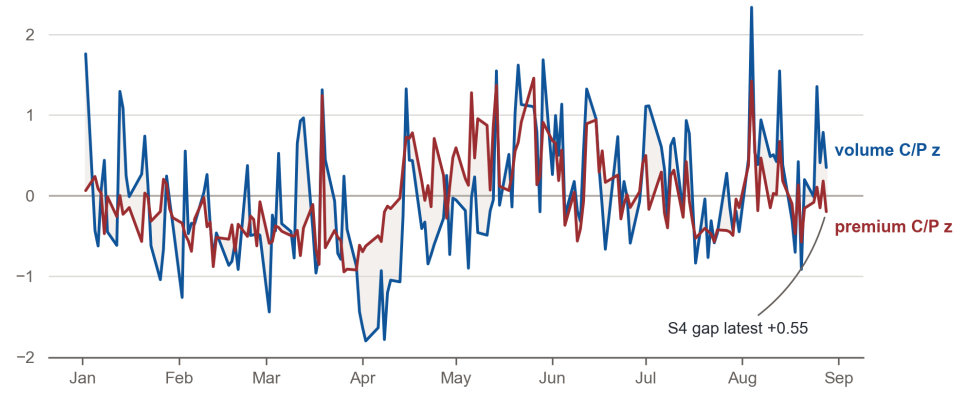


Source: Unusual Whales options-volume; WSP calculations

The gap has mattered for most of 2026, in both directions. On 21 of 165 sessions it exceeded a full standard deviation (11 positive, 10 negative). The widest positive print came on the first trading day of the year, 2026-01-02, at +1.70: contract counts scored +83.6 while dollars scored only +23.4, a call frenzy in contracts that dollar flow did not confirm. The widest negative print, 2026-04-08 at −1.58, was the mirror image: contract counts at −94.3 signaled capitulation while dollars sat near neutral at −7.1. As of 2026-08-28 the gap is +0.55, the 80th percentile of 2026, meaning wider to the bullish side than on 8 of every 10 trading days this year. The late-August wave of call buying is loud in contract counts (+27.8) but nearly silent in dollars (+1.8).

| Month end | S4 gap | S3 volume score | S2 premium score |
|---|---|---|---|
| 2026-01-30 | −0.08 | −15.2 | −21.4 |
| 2026-02-27 | −0.41 | −26.4 | −7.0 |
| 2026-03-31 | −0.83 | −79.2 | −66.1 |
| 2026-04-30 | −0.50 | −2.2 | +43.9 |
| 2026-05-29 | +0.78 | +87.1 | +74.2 |
| 2026-06-30 | −0.04 | +31.4 | +57.4 |
| 2026-07-31 | −0.30 | −23.4 | −12.6 |
| 2026-08-28 | +0.55 | +27.8 | +1.8 |

(Source: results/m1.json series S4_gap, S3_A, S2_A; scores on the $-100$ to $+100$ scale)

The same check runs at the whole-market level, across every optionable name rather than the study's index universe. As of 2026-08-28 the dollar-weighted ratio scores $+12.1$ and the contract-count ratio $+24.1$: both mildly bullish, with contract counts again the louder of the two. That feed reports only these two ratios, with no dollar-direction or buyer-versus-seller fields, so both are carried as diagnostics and neither is blended.

Contracts and dollars disagreed by more than a full standard deviation on 21 trading days of 2026, in both directions. Any single put/call headline silently picks one side of that disagreement. The study blends dollar-weighted and side-tagged components, keeps the contract-count ratio as one voice among eight, and publishes the S4 gap daily. Latest reading: $+0.55$, the 80th percentile of the year. The late-August call-buying wave registers strongly in contract counts and barely in dollars, a fact that travels with the verdict chapters.

## 3.3 The structural break audit

A multi-year rank is only honest if the series behind it was produced the same way for the whole period. Vendors change models and definitions, and when they do, the old and new halves of a series stop being comparable; ranking today against a rebuilt past produces percentiles that look precise and mean nothing. Before any module ran, the study screened all 226 harvested daily series with an algorithm that finds the dates where a series shifts to a new level (PELT change-point detection), confirmed each candidate with a second test (CUSUM), and graded the damage by how far the typical value moved across the break: a jump of $10\times$ or more is a HARD break, $2\times$ to $10\times$ is SOFT, and everything else is CLEAN.

**Results.** Of the 226 series, 52 are HARD-broken, 79 SOFT, and 95 CLEAN. Every HARD break sits in the greek-exposure feed, the vendor's model-derived estimates of dealer hedging pressure. The options-volume feed that supplies SPEC produced no HARD breaks and 3 SOFT ones (all QQQ premium fields). Of the 52 HARD breaks, 41 sit in vanna and charm series, model-derived fields that estimate how dealer hedges drift as volatility and time change, and 39 date their worst break between July and September 2023, every one of them vanna or charm. The same two dates, 2023-09-12 and 2023-09-13, recur across unrelated tickers, and the typical value jumps by $10.3\times$ to $2943.0\times$ outside the VIX and XLC families. Markets do not move dozens of unrelated series by three orders of magnitude on the same day; vendor recalculations do. The pattern is consistent with the vendor rebuilding its vanna and charm fields in late 2023, not with any market event. In the VIX and XLC families the jumps reach the order of $10^{17}$ to $10^{20}$ (VIX call vanna $2.9 \times 10^{20}$; XLC put vanna $7.1 \times 10^{20}$): pre-break values there are numerical garbage, and those history segments are discarded outright.

The table lists all 14 HARD breaks in the index complex (SPY, QQQ, IWM, DIA, VIX). All are vanna or charm series from the greek-exposure feed.

| Series | Worst ratio | Break date | Breaks | Action | Post obs |
|---|---|---|---|---|---|
| SPY call_vanna | 604.3× | 2023-07-25 | 3 | restrict Frame B | 775 |
| SPY put_vanna | 547.8× | 2023-09-12 | 4 | restrict Frame B | 741 |
| SPY call_charm | 18.7× | 2023-09-18 | 2 | restrict Frame B | 737 |
| QQQ call_vanna | 1053.9× | 2023-09-12 | 3 | restrict Frame B | 741 |
| QQQ put_vanna | 667.1× | 2023-09-14 | 3 | restrict Frame B | 739 |
| QQQ call_charm | 26.6× | 2023-09-12 | 2 | restrict Frame B | 741 |
| IWM call_vanna | 483.3× | 2023-07-26 | 4 | restrict Frame B | 773 |
| IWM put_vanna | 304.8× | 2023-09-13 | 3 | restrict Frame B | 739 |
| DIA call_vanna | 1479.8× | 2023-09-13 | 2 | restrict Frame B | 720 |
| DIA put_vanna | 694.4× | 2023-09-13 | 5 | restrict Frame B | 722 |
| VIX call_vanna | $2.9 \times 10^{20}$ | 2023-09-13 | 5 | restrict Frame B | 738 |
| VIX put_vanna | $4.1 \times 10^{19}$ | 2023-09-13 | 1 | restrict Frame B | 738 |
| VIX call_charm | $1.8 \times 10^{20}$ | 2023-09-13 | 1 | restrict Frame B | 738 |
| VIX put_charm | $2.3 \times 10^{19}$ | 2023-09-13 | 1 | restrict Frame B | 738 |

(Source: results/mod0_breaks.json; PELT rbf cost, minimum segment 20 observations, penalty $3 \ln n$, run on $\log_{10} |x|$, CUSUM confirmation; worst ratio = post/pre segment-median ratio at the worst detected break; post obs = observations after the last break)

### 3.3.1 Actions taken

- **Multi-year windows trimmed.** All 52 HARD and all 79 SOFT flags carry the same action: the multi-year (Frame B) window for the affected series starts after its last break. No series was rescaled and none was removed from the 2026 window.
- **Vanna and charm quarantined from scores.** Under the research contract the vanna/charm block (D6) never enters any score regardless of audit outcome; Chapter 5 shows it descriptively with its break status flagged. The validation module (Chapter 12) confirms the consequence: 0 of the 52 HARD-flagged series feeds the composite positioning index. The largest breaks in the dataset touch nothing that is blended.
- **Component-level trims inside SPEC.** The three options-volume SOFT flags map into components: QQQ S1 and S2 use their multi-year rank only from 2026-01-28 (put premium break, 2.1×) and QQQ S5 only from 2025-10-10 (bullish and bearish premium breaks, both 2.1×). A component's multi-year window starts at the latest break date across its raw input fields.
- **Index delta and gamma survive.** The inputs behind Chapter 5's dealer reads, both direction and amplification risk (whether dealer hedging dampens market moves or amplifies them), are CLEAN across the index complex except four SOFT flags (QQQ call delta, IWM call delta, IWM put delta, VIX put gamma); none is HARD-broken.

Two exposures survive the audit and travel to the confidence discussion in Chapter 14. First, the 2026 ranks are computed on the unrestricted year-to-date window, and 15 SOFT worst-break dates fall inside that window; two touch headline index-complex inputs: QQQ put premium (2.1$\times$ on 2026-01-28, feeding S1 and S2) and QQQ call delta (2.0$\times$ on 2026-04-09, feeding D1). Second, every number in Chapters 1 through 12 comes from a single data vendor with no in-sample cross-source check; the long-history context in Chapter 13 is the study's only non-vendor comparison.

# 4. SPEC: Speculative Demand

This chapter answers one question: what is the options crowd actually buying, calls or puts, and how hard? SPEC is the first of the study's three dials. In human terms, SPEC is what the crowd is buying, DEALER (Chapter 5) is what market makers are forced to do, and HEDGE (Chapter 6) is what protection costs. SPEC matters most right now because it is the dial that moved: speculative demand jumped from the 21st percentile of the year to the 77th in the final six sessions of August, and that swing travels straight into the study's verdict (Chapter 14).

SPEC measures speculative directional demand in the options market, in dollars, in contracts, and at the quote side; positive scores mean bullish speculative consensus. It carries the largest weight in the composite Positioning Consensus Index, 0.45 of PCI, and is built from the vendor's `options-volume` family at the ticker level plus one market-wide tide component. Each component's raw reading is ranked two ways: against every trading day of 2026 so far (Frame A, the 2026 rank) and against the full harvested multi-year history (Frame B, the multi-year rank). The rank then becomes a score $(\mathrm{pct} - 0.5) \times 200$, bounded in $[-100, +100]$: zero is the year's typical day, $+100$ its most bullish extreme, $-100$ its most bearish. The YTD panel covers 165 sessions from 2026-01-02 through 2026-08-28. (Source: m1.json; RESEARCH_PLAN sec 6.1)

## 4.1 Construction

Eight components enter the blend at equal weight: S1, S2, S3, S5, S6, S7, S8, S9. S4 is computed and reported but kept out of the blend as a disagreement diagnostic. Ticker-level SPEC is the equal-weight mean of the component scores present for that ticker; market-level SPEC blends the four index ETFs at the contracted cap-tier weights, SPY 0.40, QQQ 0.30, IWM 0.15, DIA 0.15, with the market-wide S8 series injected into each ticker frame before blending.

| Id | Definition | Sign | Role |
|---|---|---|---|
| S1 | Net premium tilt = (net_call_premium − net_put_premium) / (call_premium + put_premium) | + | Blend |
| S2 | Premium C/P ratio = call_premium / put_premium | + | Blend |
| S3 | Volume C/P ratio = call_volume / put_volume | + | Blend |
| S4 | Ratio gap = $z_A$(S3) − $z_A$(S2), contracts-versus-dollars disagreement | unsigned | Diagnostic only (Ch. 3) |
| S5 | Bullish premium share = bullish_premium / (bullish_premium + bearish_premium) | + | Blend |
| S6 | Call ask-side share = call_volume_ask_side / (call_volume_ask_side + call_volume_bid_side) | + | Blend |
| S7 | Put ask-side share, symmetric to S6 on the put side; protection lifted at the ask | − | Blend |
| S8 | OTM net-premium tilt, market-wide: (OTM net_call_premium − OTM net_put_premium) / \|all net_call_premium − all net_put_premium\|, EOD cumulative from market-tide all-versus-OTM [amended wording] | + | Blend, market-wide |
| S9 | Volume vs 30d average = (call_volume + put_volume) / (avg_30_day_call + avg_30_day_put), signed by sign(S1) | sign(S1) | Blend |

(Source: RESEARCH_PLAN sec 6.1; S8 stated as implemented per m1.json notes and the QC pass)

One wording amendment is disclosed here: the plan's original S8 text, "EOD cumulative net_call_premium share OTM," does not describe the implemented quantity and fails to reproduce any headline number under its literal reading. All results use the implemented, QC-verified OTM net-premium tilt given in the table.

## 4.2 Component walk

The table's definitions carry most of the load; four points matter for reading them.

First, S5 folds in trade side and moneyness through the vendor's bullish/bearish classification of each trade, not option type alone; a put sold aggressively counts as bullish premium.

Second, S6 and S7 read the quote side. A trade at the seller's asking price is buyer-initiated; a trade at the buyer's bid is seller-initiated. Calls lifted at the ask are aggressive bullish demand, so S6 enters positively. Puts lifted at the ask are protection being paid up for, so S7 enters negatively.

Third, S4 stays out of the blend and works as a gauge of when the contract tape and the dollar tape disagree. On 2026-08-28 the gap is +0.55, the 80th percentile of 2026: the call-to-put ratio in contract counts runs hotter than the same ratio in dollars, consistent with cheap, short-dated contracts, including options expiring the same day (0DTE), carrying more of the call buying than the dollars do. Chapter 3 treats this measurement problem in full.

Fourth, S8 is the only market-wide component. It tracks net premium in out-of-the-money (OTM) contracts, where convexity is bought. The market-tide feed exists for 2026 only, so S8's multi-year rank equals its 2026 rank by construction.

## 4.3 Redundancy and known pathologies

The adversarial review found no component sign provably backwards, but it established that the blend double-counts one signal. S1, S5, and S9 form a near-duplicate block: score correlations through August 20 ran 0.871 (S1 to S5), 0.762 (S1 to S9), and 0.677 (S5 to S9), and S9 carries sign(S1) by construction. Net-premium direction therefore holds roughly 3 of 8 units of SPEC weight, an effective triple weight on one underlying signal.

S9 is also discontinuous: with magnitude always positive and sign inherited from S1, the series jumps whenever S1 crosses zero. The review counted 57 sign flips on SPY across the year's first 159 sessions, with a median raw jump of 2.0 on a series whose typical scale is about 1, and the whipsaw stayed live into the refresh: SPY's raw S1 flipped sign three times in the final ten sessions alone. The component is retained because the contract fixed it ex ante, but its day-to-day moves should not be over-read.

The round-two review then turned this block against the late-August move itself: if a triple-counted signal drove the jump, the week's bullish reading would be an artifact. The attack failed. The swing is broad-based. Six of the eight component 5-day averages rose against ten sessions earlier (S1 +57.5, S6 +57.2, S5 +39.2, S7 +30.3, S9 +25.9, S8 +24.8 in score points), the other two eased only slightly (S2 −8.2, S3 −8.3), and collapsing S1, S5, and S9 into a single voice raises the 5-day SPEC average from +21.3 to +23.7 rather than lowering it, leaving the composite's verdict class unchanged. The disclosure stands; it does not bite on this reading. (Source: adversarial_log.json, sign-convention attack and round-2 redundancy attack)

## 4.4 The 2026 trajectory

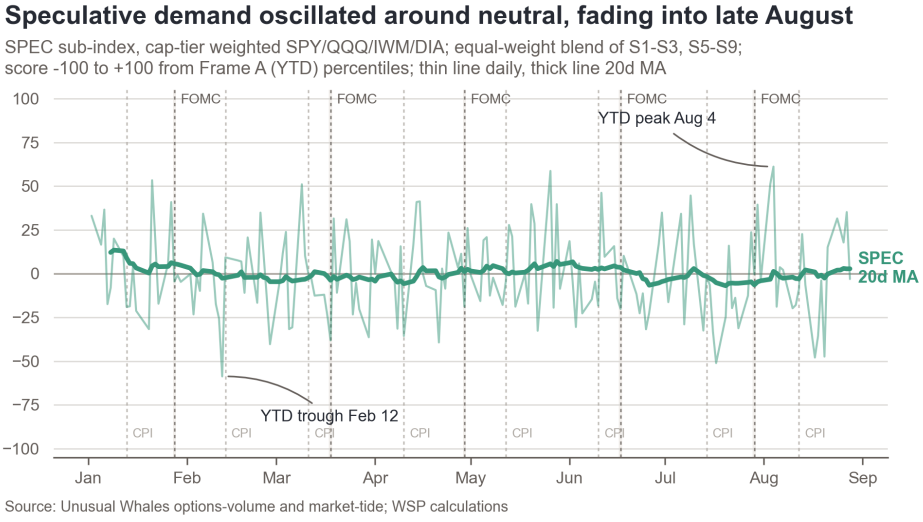


The timeline (daily SPEC with a 20-day moving average, FOMC and CPI sessions shaded) shows an oscillating, mean-reverting year: 92 of 165 sessions printed negative, and no direction held for long. Four episodes structure the path. January brought high-amplitude two-sided chop, including +53.3 on January 21 (daily rank 0.982). SPEC then collapsed to its YTD low of −58.7 on February 12, the single lowest print of the year, and snapped back within a session to +9.3 on February 13. Spring and early summer

oscillated: bullish spikes on March 9 (+51.0) and May 26 (+58.7) alternated with bearish runs into April 21 (−39.3, the composite's YTD trough session, Chapter 8) and a July slide to −51.1 on July 17.

August then produced the year's defining feature: a round trip in each direction inside one month. First up, a two-session burst to +51.2 on August 3 and the YTD peak of +61.1 on August 4, the single highest print of the year. Then down: nine of the next twelve sessions printed negative, including −47.9 on August 17 and −47.4 on August 20, the year's fourth-lowest day (daily rank 0.024). Then up again: five consecutive positive sessions from August 21 through August 27 (+15.1, +31.5, +24.9, +17.8, +35.2) lifted the 5-day average to +21.3 by the close of the window. The final session eased to −3.1, a median day (daily rank 0.494). The strength on the page is one week old, not one day old.

The multi-year rank disciplines the narrative. The February 12 trough, the lowest day within 2026, sits only at the 21st percentile of the multi-year component window, and the August 4 peak reaches the 80th rather than the 99th. 2026's speculative extremes are loud within the year and ordinary against the fuller history.

**SPEC component scores by month**

Monthly mean of Frame A component scores (-100 to +100), market-weighted across SPY/QQQ/IWM/DIA; bottom row = blended SPEC sub-index

| | Jan | Feb | Mar | Apr | May | Jun | Jul | Aug |
|---|---|---|---|---|---|---|---|---|
| S1 | +10 | +19 | +26 | -4 | -4 | -14 | -13 | -19 |
| S2 | +5 | -38 | -58 | +9 | +54 | +24 | -11 | +18 |
| S3 | +4 | -25 | -15 | -36 | +18 | +18 | +8 | +30 |
| S5 | +12 | +14 | +12 | +2 | -6 | -11 | -9 | -12 |
| S6 | +21 | +2 | -5 | +18 | +11 | -14 | -13 | -17 |
| S7 | -13 | -14 | +0 | +15 | -14 | -0 | +15 | +7 |
| S8 | -8 | -4 | -4 | +17 | -4 | -20 | +0 | +24 |
| S9 | +7 | +9 | +25 | -0 | -12 | -10 | -11 | -8 |
| SUB | +5 | -5 | -3 | +2 | +5 | -3 | -4 | +3 |

Source: Unusual Whales options-volume and market-tide; WSP calculations

The heatmap shows the late-August swing is broad rather than the work of one input: as of August 28 all eight component 5-day averages are positive (table below). The final day itself is split, five components positive and three negative, with the deepest negative in S7 (−53.2): buyers paid up for puts into the last close of the window.

## 4.5 Where SPEC stands on 2026-08-28

| Component | Latest score | 5d mean | Implied pct (A) |
|---|---|---|---|
| S1 net premium tilt | +16.33 | +16.62 | 0.582 |
| S2 premium C/P | +1.75 | +21.15 | 0.509 |
| S3 volume C/P | +27.83 | +36.93 | 0.639 |
| S5 bullish premium share | +13.61 | +15.70 | 0.568 |
| S6 call ask-side share | −24.70 | +17.70 | 0.377 |
| S7 put ask-side share | −53.19 | +22.07 | 0.234 |
| S8 OTM net-premium tilt | −22.89 | +30.60 | 0.386 |
| S9 volume vs 30d, signed | +16.69 | +9.34 | 0.583 |
| **Market SPEC (blend)** | −3.07 | +21.26 | 0.494 [daily rank] |

(Source: m1.json series, 2026-08-28; implied pct (A) via the contract identity pct $=$ score$/200 + 0.5$; the blend row's percentile is the daily rank of SPEC within YTD)

The equal-weight mean of the component 5-day means reproduces the sub-index 5-day mean of +21.26 exactly. Two conventions turn that into a percentile. The verdict-rule convention (the latest 5-session mean of the daily SPEC rank within YTD, the same convention Chapter 14 applies to the composite) gives 0.7723: the 77th percentile of 2026, meaning the recent days rank higher than about three of every four trading days this year. The component-average convention is milder, 0.6063 within 2026 and 0.6039 against the multi-year window. SPEC reads bullish under either; how bullish depends on the yardstick, hot by the year's own standard and moderate by the longer one. Six sessions earlier, on August 20, the verdict-rule statistic sat at the 21st percentile; the entire move is the final week. By ticker:

| Ticker sub-index | Frame A | Frame B |
|---|---|---|
| SPY | +15.81 | +14.48 |
| QQQ | +2.26 | +8.30 |
| IWM | −25.90 | −32.93 |
| DIA | −41.27 | −35.52 |
| **Market (.40/.30/.15/.15)** | −3.07 | −1.98 |

(Source: m1.json SUB series, 2026-08-28)

The final-session split is large-cap versus the rest: SPY and QQQ hold positive scores while IWM and DIA print negative. The single-name cross-section behind that split is Chapter 9's subject. Two market-wide diagnostics computed outside the blend, the same call-to-put ratios taken on the vendor's whole-market totals rather than the four ETFs, point the same way as the week: 5-day averages of +8.7 on the dollar ratio and +44.6 on the contract ratio.

SPEC enters the verdict window hot but already cooling. The defining feature is August's double round trip: the YTD peak of $+61.1$ on August 4, a fade to $-47.4$ by August 20 (daily rank 0.024), then five straight positive sessions that lifted the 5-day average to $+21.3$, the 77th percentile of 2026 on the verdict-rule convention. Three qualifiers travel with the number. First, it is one week old: the final session printed $-3.1$, a median day, so the bullish reading describes August 21 to 27, not the as-of close. Second, the longer context is milder: the same week sits near the 60th percentile on the multi-year component yardstick. Third, the swing is real rather than an artifact of the blend's known double-count: all eight component 5-day averages are positive, and removing the S1/S5/S9 duplication makes the week read slightly hotter, not cooler.

Frame and data caveats, in plain terms. (1) All SPEC inputs are single-vendor (Unusual Whales) with no in-sample cross-source validation. (2) One flagged discontinuity sits inside the 2026 ranking window with no restriction: QQQ put premium's median roughly doubles (2.1x) at 2026-01-28, so 2026 ranks for QQQ dollar-based components knowingly span that jump. The market-blend effect is small, about 3.8% of SPEC (S2 is one eighth of QQQ's blend at 0.30 market weight). The refreshed break audit adds no new flag on any series inside the market blend; the new SOFT flags are confined to sector and megacap series. (3) Multi-year ranks are restricted where the audit found breaks: QQQ S1 and S2 use post-2026-01-28 history only, QQQ S5 post-2025-10-10. The Frame B window for options-volume components is the rolling 500-session span 2024-08-30 to 2026-08-28. (4) S8's Frame B equals Frame A because market-tide exists YTD only. (5) Frame B percentiles here use M1's component-average convention (score$/200 + 0.5$) because the blended sub-index has no pre-2026 path to rank against; cross-module Frame B comparisons are therefore not exact. (Source: m1.json notes; mod0_breaks.json; adversarial_log.json)

# 5. DEALER: What Market Makers Are Forced to Do

The DEALER sub-index tracks what market makers are forced to do. (Its two siblings: SPEC tracks what the crowd is actually buying, and HEDGE tracks what protection costs.) It answers two questions that sound similar and are not: which way customer option flow is leaning, and whether dealer hedging will make the next market move bigger or smaller. This week the answers are that flow went from clearly selling to roughly flat in six sessions, and that the hedging machinery sits almost exactly on the line between dampening moves and amplifying them.

## 5.1 Components and construction

DEALER carries a 0.20 weight in the composite. The research contract (section 6.2) splits it into two axes that are never blended. The direction blend (D1, D4, D5) estimates which way customer option flow leans; its percentile is what enters the composite. The amplification flag F, built from net gamma (D2), describes how violently the market transmits whatever flow arrives; it stays out of the composite entirely and touches the verdict only through the "overly" qualifier of the study's pre-registered verdict rule (contract section 8). This chapter reports the two axes separately throughout. Table 5.1 lists the six contracted components and the 2026-08-28 readings. The direction blend is the equal-weight percentile score of D1 (all four index tickers), D4 (SPY only), and D5 (SPY and QQQ); the market-level series is the cap-tier weighted mean over per-ticker sub-indices at SPY 0.40, QQQ 0.30, IWM 0.15, DIA 0.15. D3 is a snapshot metric with no YTD distribution. D6 failed the break audit and is excluded from all scores, shown descriptively in section 5.5.

| Id | Definition | Role | Latest | 2026 pct | Multi-yr pct |
|---|---|---|---|---|---|
| D1 | Customer delta demand, −(call_delta + put_delta), SPY shown | Direction, in blend | −88,059,059.7 | 31% | 37% |
| D2 | Net gamma, call_gamma + put_gamma, SPY | Amplification axis only | −12,649.7 | 70% | 72% |
| D3 | Gamma-flip distance, (spot − flip)/spot | Amplification, snapshot | −0.06% | n/a | n/a |
| D4 | NOPE daily close, SPY | Direction, in blend | −5.47 | 13% | n/a |
| D5 | dir_delta_flow daily sum, SPY shown | Direction, in blend | 496,162.3 | 48% | n/a |
| D6 | Vanna / charm exposure | Excluded, descriptive | see 5.5 | n/a | n/a |
| F | Amplification risk = percentile(−net gamma), market | Sec-8 qualifier only | 0.3301 | 33% | 29% |

Table 5.1. DEALER components as of 2026-08-28. Raw levels for D1, D2 and D5 are in the greek-exposure and greek-flow endpoints' native exposure units. D4 and D5 harvests exist only from January 2026, so they carry no multi-year depth and their multi-year percentiles are null by construction. F is itself a percentile; its "latest" is the 2026 value and its multi-year column is the re-rank against history back to 2022-03-30. (Source: m2.json headline_numbers; contract section 6.2)

One caveat attaches to the direction axis before any number is read. D1, the largest-capacity input, is customer delta demand: the negative of dealers' net delta as the vendor reports it. Its economic sign rests

on the vendor's assumption about which side of each trade the dealer holds, and nothing inside this study's data boundary can verify that assumption independently. The adversarial pass (Appendix C) found no evidence the sign is wrong, but also no independent confirmation that it is right. The direction read below is conditional on that assumption.

## 5.2 The direction read: from selling to flat

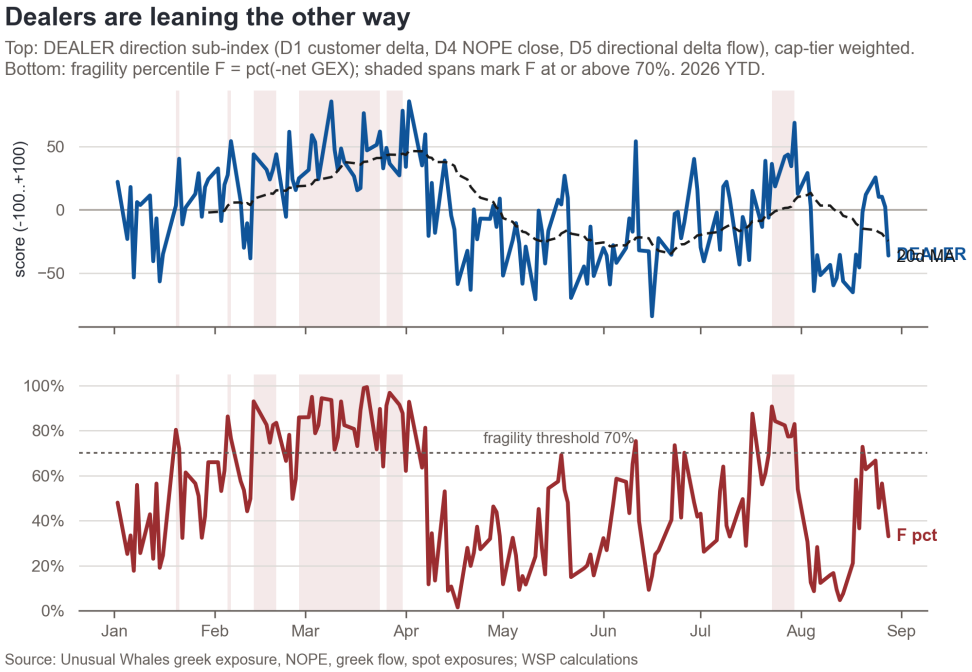


Exhibit 5.1 plots the market-level direction blend across 2026 with the amplification-risk band beneath it. The latest 5-day average is +2.5, which ranks at the 51st percentile of every trading day of 2026 so far (the 2026 rank), meaning higher than about half of this year's sessions: the middle of the range. Ranked against the full multi-year history (the multi-year rank) it sits at the 55th percentile, with the caveat in the data note below: two of the blend's three inputs exist only from January 2026, so the multi-year rank leans on one input's longer history.

The move to flat is recent and sharp. Six sessions earlier, on August 20, the same 5-day average stood at −40.7. Customer flow, on this measure, went from clearly selling to neutral in a week. Flat is the right word, not bullish. The +2.5 average sits on a scale that runs −100 to +100, and the series swings about 37 points on a typical day, so the margin above zero is roughly 7% of one ordinary daily move. The five days inside the average printed +25.6, +10.2, +10.3, +2.3 and −36.0: the newest print is the most negative of the scored week. The average stays positive through the next session only if that session prints above +13.2, which 39% of 2026 sessions have done. Chapter 14 weighs what this thin margin means for the verdict's confidence label.

The flip is not broad either. Table 5.2 splits the blend by ticker. On the 5-day averages, QQQ and IWM lean positive, SPY is flat, and DIA is firmly negative. On the as-of day itself, three of the four tickers printed

negative.

| Ticker | Cap weight | 2026-08-28 | Latest 5d avg |
|---|---|---|---|
| SPY | 0.40 | −39.0 | +2.4 |
| QQQ | 0.30 | −38.6 | +14.8 |
| IWM | 0.15 | +18.1 | +32.0 |
| DIA | 0.15 | −77.1 | −51.6 |
| Market | 1.00 | −36.0 | +2.5 |

Table 5.2. Per-ticker DEALER direction scores (−100 to +100) as of 2026-08-28. The market row is the cap-tier weighted blend. 5-day averages computed from the per-ticker series in m2.json. (Source: m2.json series DEALER_dir_*_A)

Across the year the blend has run nearly its full range: a peak of +85.5 on April 2, a trough of −83.8 on June 16, and 81 of 165 sessions closing negative, a genuinely two-sided year. August supplied the whole recent story. The blend was mildly positive on August 3 (+29.0) and August 4 (+2.7), the day speculative buying peaked (Chapter 4). It dropped to −64.0 on August 5 and held between roughly −35 and −65 through August 19. It then printed positive for five straight sessions, August 21 through August 27, before closing the month at −36.0 on August 28.

Within the blend, the inputs disagree. D4 is the bearish extreme: NOPE, a measure of how much option trading is pushing around the stock itself, closed at −5.5, the 13th percentile of 2026, against a YTD range of −40.9 (January 15) to +30.0 (February 6). D1 sits at the 31st percentile and D5 at the 48th. The as-of day's −36.0 blend is led by that NOPE close; the 5-day average's positive sign is led by the QQQ and IWM legs. Neither side of the disagreement is at an extreme, which is the point: this sub-index is currently a weak vote, not a strong one.

## 5.3 Amplification risk: still the 2026 norm, now at the boundary

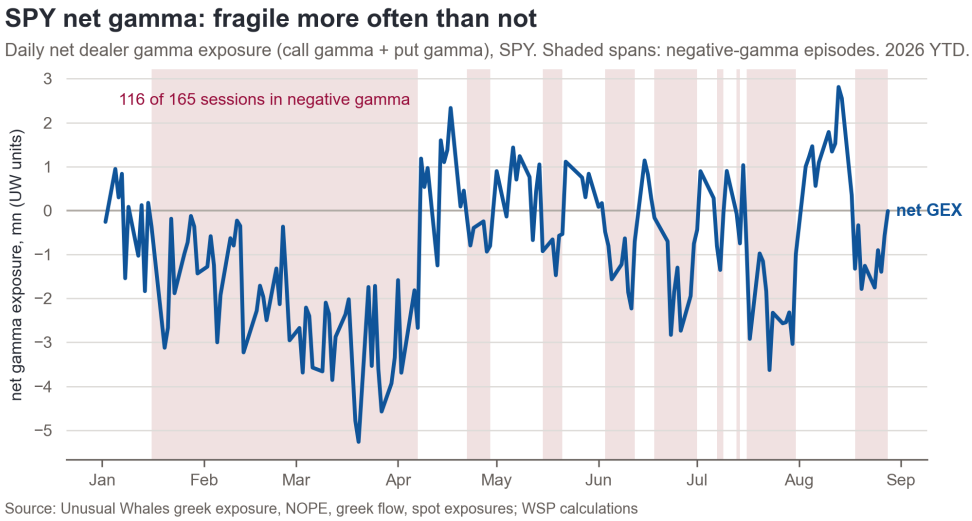


The second axis asks how hard the market transmits whatever flow arrives. The input is net gamma, a measure of whether market makers' hedging dampens market moves (positive) or amplifies them (negative).

When dealers are net short gamma, hedging forces them to sell as the market falls and buy as it rises, which makes moves bigger in both directions. This report calls that condition amplification risk. It says nothing about direction; it is about transmission.

Exhibit 5.2 shades the negative-gamma stretches across the year, and the correct summary is that they are not episodes so much as the resting state. SPY net gamma has been negative in 116 of 165 sessions, 70% of the year, with a YTD median of about −0.8 million in the vendor's exposure units. Dealer hedging positioned to amplify moves rather than dampen them is simply how 2026 has traded. The latest reading extends that streak in sign but not in substance: net gamma printed −0.01 million on August 28, the year's closest print to zero, effectively on the flip line itself. The path there ran through both extremes inside one month. Net gamma was positive from August 3 through August 17, including the year's most positive print, +2.81 million on August 13. It turned negative on August 18 and has stayed negative for nine sessions, but the depth faded from −1.79 million on August 20 to essentially zero on August 28. The year's opposite extreme, −5.26 million, printed on March 20.

The study compresses this axis into one flag, F: the percentile of negative net gamma across the book, so a higher F means hedging is positioned to amplify moves more than usual. F closed at 0.3301, the 33rd percentile of 2026, and the 29th percentile against the multi-year history back to March 2022. Both frames now agree amplification risk is below normal. The flag has been violent lately: across the last ten sessions it traversed 0.21 (August 17) to 0.73 (August 20) to 0.33 (August 28). Its year's extremes are 0.9931 on March 20 and 0.0145 on April 17, and 45 of 165 sessions, 27% of the year, have printed F at or above 0.70. The flag visits its high zone often, and it was there six sessions ago.

F matters to the verdict only through the "overly bullish" qualifier of the section 8 rule, which arms when either F reaches 0.70 or the HEDGE percentile reaches 0.80. The F leg is not armed: 0.3301 sits far below its 0.70 gate. The qualifier currently fires on the other leg alone, the HEDGE 5-day percentile at 0.8193 against its 0.80 gate, a margin of 1.9 percentage points. Chapter 14 prints that margin and its fragility next to the verdict, and Chapter 15 tracks both legs going forward.

Direction and mechanics are different axes, and this week both moved fast. The direction blend went from clearly bearish (5-day average −40.7 on August 20) to flat (+2.5, the 51st percentile of 2026) in six sessions; that is the number that enters the composite at a 0.20 weight. The amplification flag collapsed over the same stretch, from 0.73 to 0.33, and net gamma sits at the year's closest print to zero, right on the flip line. Neither axis is at an extreme. The caution is freshness: the +2.5 average is about 7% of one typical daily swing, the newest daily print (−36.0) is the most negative of the scored week, and negative net gamma remains the 2026 norm (116 of 165 sessions).

## 5.4 The strike map on 2026-08-28

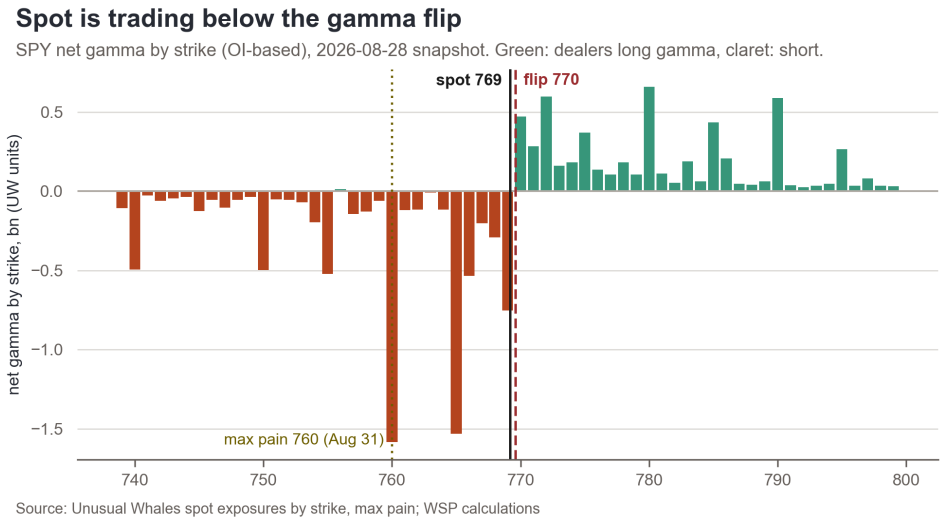


Exhibit 5.3 gives the strike-level snapshot behind the mechanics read. The gamma flip is the price level where dealer hedging switches from dampening moves to amplifying them. The interpolated flip for SPY sits at 769.62, and spot closed 0.06% below it, so the market ended the week just on the amplifying side of the line. That matches the barely negative aggregate net gamma in the prior section. Max pain, the strike where the most option value would expire worthless, is 760 for the 2026-08-31 expiry, the nearest expiry after the as-of date; it sits below spot. The full strike range shows multiple zero crossings of net gamma; the flip quoted is the interpolated zero at the sign-change pair nearest to spot, and for SPY that crossing is unique near spot. These are one-day snapshot quantities with no YTD distribution, so no percentile attaches; the historical proxy for which side of the flip the market sat on remains the sign of D2.

## 5.5 Vanna and charm: shown, not scored

Vanna and charm are the second-order dealer pressures that options commentary often cites: vanna tracks how dealer hedges must shift when volatility changes, and charm tracks how they must shift as expiration approaches. The contract lists them as component D6, and D6 is the one contracted component the break audit removed from scoring. All four SPY series break inside the mid-2023 to mid-2024 window Chapter 3 documents across the greek-exposure family, three at 10x-plus magnitude (Table 5.3).

| Series | Verdict | Break date | Ratio | Post obs | Action |
|---|---|---|---|---|---|
| SPY call_vanna | HARD | 2023-07-25 | 604.3x | 775 | restrict multi-yr |
| SPY put_vanna | HARD | 2023-09-12 | 547.8x | 741 | restrict multi-yr |
| SPY call_charm | HARD | 2023-09-18 | 18.7x | 737 | restrict multi-yr |
| SPY put_charm | SOFT | 2024-08-28 | 4.7x | 501 | restrict multi-yr |

Table 5.3. Break-audit entries for the SPY vanna and charm series, from an algorithm that finds the dates where a data series shifts to a new level (PELT, on log10 absolute levels) with a second test (CUSUM) confirming each break; ratio is the worst post/pre segment-median ratio. Per contract 6.2, broken D6 series are excluded from all scores and shown descriptively only. (Source: mod0_breaks.json hard/soft entries)

Descriptively, on post-break segments only: the latest SPY vanna exposure is −631.7 million, the 5th percentile of 2026 and the 37th percentile of its own segment since 2023-09-12; the latest charm exposure is +245.2 million, the 67th percentile of 2026 and the 65th of its segment since 2024-08-28. Neither number feeds the direction blend, the amplification flag, or the composite. They locate today's surface for readers who track vanna and charm commentary, with the explicit warning that the pre-break history behind popular multi-year comparisons of these series is not usable.

## 5.6 What enters the verdict

DEALER reaches the verdict through two separate doors, and this week both matter. The direction blend enters the composite at its 0.20 weight, and its 5-day average of +2.5 is the reading the section 8 rule scores. It is the thinnest of the three sub-index inputs: on the as-of day alone, DEALER subtracted 7.2 points from the daily composite (0.20 times the day's −36.0; Chapter 7 carries the full attribution). The composite crossed its bullish gate this week (P = 0.8494 on the 2026 frame; the multi-year frame reads 0.6126, below the gate), and the rule's confidence stamp counts DEALER's +2.5 as sign-agreement with SPEC and HEDGE. Chapter 14 prints that stamp with its margin attached: measured on the day's prints rather than 5-day averages, DEALER disagrees at −36.0, and one ordinary dealer session flips the sign.

The amplification flag stays outside the index. F at 0.3301 leaves its leg of the "overly bullish" qualifier unarmed; the qualifier rides on the HEDGE leg alone, 0.8193 against a 0.80 gate, a 1.9 percentage point margin. Chapter 14 carries the qualifier's full statement and Chapter 15 the thresholds that would change the read on both legs.

Integrity items for this chapter. (1) The D1 hand-recompute for SPY matches the stored series exactly (−88,059,059.7); D4 is verified as the session close after sorting the reverse-chronological NOPE parquets, where a naive daily-last aggregation would return the 13:30 opening bar (checked against the 2026-08-28 file); and the flip-strike sign change is verified for all four index tickers (m2.json checks). (2) Multi-year restrictions from the break audit: the QQQ D1 leg is restricted to post 2026-04-09 (SOFT call_delta break, 2.0x, 98 observations) and the IWM leg to post 2025-06-24 (SOFT put_delta, 2.3x, 297 observations); the QQQ break date falls inside the 2026 window, where the contract applies no restriction, so 2026 percentiles for that leg span a flagged discontinuity. (3) D4 and D5 are YTD-only harvests, so the blend's multi-year rank mixes full-depth D1 with 2026-only D4 and D5; multi-year DEALER percentiles are indicative, not a clean multi-year rank. (4) All DEALER inputs are single-vendor Unusual Whales series, and the D1 sign convention rests on the vendor's dealer-inventory assumption, which has no independent confirmation inside the study window.

# 6. HEDGE: The Price of Protection

This chapter answers two questions. What does downside protection cost right now, and is anyone actually buying it. As of the August 28 close the answers point in opposite directions: insurance is priced near its cheapest levels of 2026, while the two gauges that track live hedging activity show investors still paying up for it. That gap between a cheap price and live demand is the single most important nuance behind the study's verdict (Chapter 14).

HEDGE is one of the study's three dials: SPEC tracks what the crowd is buying, DEALER tracks what market makers are forced to do, and HEDGE, this chapter, tracks what protection costs. It carries a 0.35 weight in the composite and blends seven equal-weighted inputs, H1 through H7, all measured on SPY at the market level. The sign convention is fixed by the contract: positive means insurance is cheap or unbid, which reads as complacency and bullish consensus; negative means insurance is bid, which reads as fear.

As of 2026-08-28 the HEDGE 5-day average stands at +31.4, the 82nd percentile of 2026 so far, meaning higher than roughly 8 of every 10 trading days this year. Ranked against the fuller multi-year history it sits at the 85th percentile (a component-average convention; see the note at the end of this chapter). The daily close is +29.1 on the 2026 ranking, its 77th percentile, and +28.1 on the multi-year ranking. HEDGE remains the highest of the three dials: SPEC's 5-day average ranks at the 77th percentile of the year and DEALER's at the 51st.

## 6.1 Construction and current levels

Each input is converted to a percentile, sign-oriented so that positive reads complacent, then mapped to a $-100$ to +100 score; the blend is the equal-weighted average. Two rankings appear throughout: the 2026 rank (the study's Frame A), against every trading day of 2026 so far, and the multi-year rank (Frame B), against the full harvested history of each input. Raw values are in volatility points, the units implied volatility is quoted in, where applicable. 5-day scores are the average of the last five sessions.

| Component | Sign | Raw | 2026 pct | M-yr pct | Score (day) | Score (5d) |
|---|---|---|---|---|---|---|
| H1 SPY 1-year IV rank: where SPY option prices sit in their own one-year range, 0 = cheapest, 100 = most expensive (high = vol bid = fear) | – | n/a | n/a | n/a | +97.6 | +74.0 |
| H2 25-delta risk reversal: upside call IV minus equally out-of-the-money put IV (high = calls bid over puts) | + | +2.6 | 0.1761 | 0.1577 | −64.8 | −55.2 |
| H3 Volatility risk premium: what options charge for future volatility minus the volatility that shows up, spliced per Section 6.3.2 (high = insurance rich = fear) | – | +1.4 | 0.4639 | 0.4405 | +7.2 | +34.2 |
| H4 Term slope: 90-day IV minus 7-day IV; longer-dated options costing more reads calm | + | +4.1 | 0.7771 | n/a | +55.4 | +33.0 |
| H5 SPY short-volume ratio: share of SPY volume sold short (high = hedging pressure) | – | n/a | n/a | n/a | −88.0 | −47.7 |
| H6 VIX 1-year IV rank: the cost of insuring against a volatility spike (high = crash protection bid) | – | 5.4 | 0.0120 | 0.0397 | +97.6 | +92.5 |
| H7 30-day implied-move percentile: how large a move option prices are braced for, vendor-provided (high = large moves priced) | – | n/a | n/a | n/a | +98.8 | +88.7 |
| **HEDGE blend (equal weight)** | | | 0.8193 | 0.8494 | **+29.1** | **+31.4** |

(Source: Definitions from RESEARCH_PLAN.md section 6.3; raw values and percentiles from m3.json headline_numbers; day scores from the m3.json score series; 5-day scores are means of the last five sessions of those series and match the round-2 adversarial recomputation. Blend-row percentiles rank the 5-day average. n/a: raw level or multi-year rank not published in m3.json.)

Five of seven components sit at complacent scores. The two that do not, H2 and H5, are the two that track what hedgers are actually doing rather than what insurance costs. That split shapes the whole reading of this chapter.

## 6.2 The year in insurance

Insurance started the year unbid, was aggressively bid through the February to April drawdown, normalized in May, and whipsawed in June, which contains both the YTD maximum (+60.7 on 2026-06-02) and minimum (−60.0 on 2026-06-23) of the daily series. Late July brought one more fear flash (−47.1 on 2026-07-29). The turn after it was abrupt: three sessions later August opened at +49.5, and all twenty August sessions printed positive, the only month of 2026 with that property (April is its mirror image, the only all-negative month). That August run is what holds the 5-day average at the 82nd percentile.

| 2026 | Jan | Feb | Mar | Apr | May | Jun | Jul | Aug |
|---|---|---|---|---|---|---|---|---|
| Sessions | 20 | 19 | 22 | 21 | 20 | 21 | 22 | 20 |
| Mean | +26.7 | −6.5 | −35.1 | −28.1 | +10.8 | +0.9 | +4.7 | +30.5 |
| Min | −24.8 | −32.6 | −55.3 | −52.2 | −23.6 | −60.0 | −47.1 | +5.5 |
| Max | +54.3 | +17.5 | +0.2 | −1.8 | +55.5 | +60.7 | +43.3 | +53.8 |

(Source: Monthly aggregation of HEDGE_SPY_frameA dailies through 2026-08-28, m3.json.)

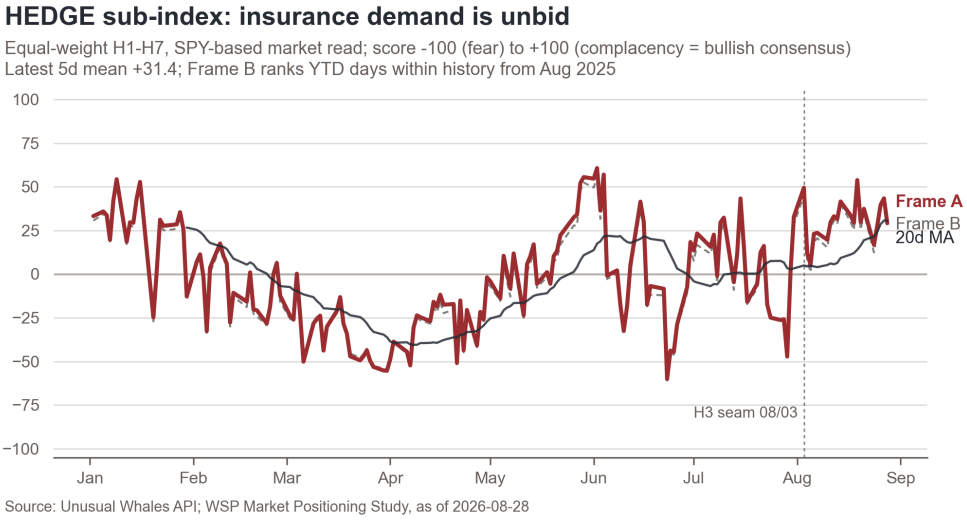


## 6.3 The raw inputs: skew, risk premium, term structure

### 6.3.1 H2: the risk reversal stayed positive all year, and is now near its low

The risk reversal compares the price of upside call options with equally out-of-the-money downside puts; when puts cost more, the market is paying up for crash protection. This year the measure was positive on all 158 sessions with a valid observation: calls carried higher implied volatility than puts every observed day. The YTD band runs from +1.8 volatility points (2026-08-04) to +10.8 (2026-03-06). The latest print, +2.6, sits at the 18th percentile of 2026 and the 16th of the multi-year window, close to the bottom of an unusually call-tilted year.

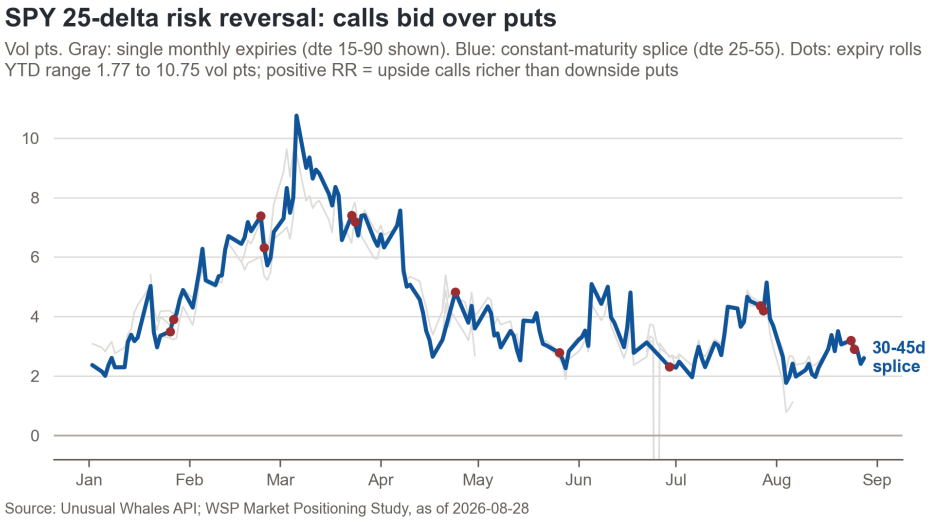


The position within the band matters more than the level. The least call-tilted reading of the year printed on 2026-08-04, and the latest value sits only 0.8 volatility points above it. Relative to this year's own range, puts are getting bid at the margin even as volatility itself stays cheap. That is the demand-side dissent behind H2's −55.2 5-day score.

### 6.3.2 H3: the volatility risk premium is back to flat, and the series has a seam

The volatility risk premium is the gap between what options charge for future volatility and the volatility that actually shows up. It stands at +1.4 volatility points, the 46th percentile of 2026 and the 44th multi-year: options are priced roughly in line with delivered movement, neither rich nor cheap. The measure hugged zero through August (monthly average −0.6) and flipped positive in the final two sessions; 51 of 165 sessions this year printed negative. The YTD range runs from −4.9 (2026-05-29) to +11.4 (2026-03-30, the fear peak). The sign convention behaved as contracted in-sample: the premium averaged +6.1 volatility points across April, when insurance was rich, and H3's April scores averaged −78 (fear); the near-zero August average scores +59 (complacent).

H3 carries the one construction seam in this sub-index. The vendor `risk_premium` series defines the premium against volatility realized after each date, which goes null once that forward window passes the data edge; its last value is 2026-07-31. From 2026-08-03 the series is extended as implied volatility minus trailing 20-day realized volatility, indexed to the date each trailing value was known. The two definitions differ by construction, so pre- and post-seam values are not strictly comparable. The continuity check passed: the first extension day moved 0.030 against a typical daily move of 0.014, inside the pre-set three-sigma limit of 0.042. Two facts keep this honest. All five sessions in the current scoring window are extension values, so the latest H3 reading rests entirely on the post-seam definition. And the headline does not depend on it: dropping H3 outright moves the blend's 5-day percentile from 0.8193 to 0.8434 and changes no conclusion (Section 6.4).

### 6.3.3 H4: term structure in firm contango

The term slope asks whether longer-dated options cost more than short-dated ones, the calm configuration, or less, the stress configuration. The 90-day minus 7-day slope is +4.1 volatility points, the 78th percentile of

2026, firm calm-side territory. The year's range runs from −5.8 (2026-03-09, inverted during the drawdown) to +6.1, a ceiling the series touched twice, on 2026-01-16 and 2026-08-17. August has been the steepest sustained stretch of the year, averaging +4.3. Interpolated implied volatility exists YTD only, so H4 and H7 have no distinct multi-year rank.

**SPY implied volatility term structure, YTD**

Interpolated IV at 7d, 30d and 90d tenors; 90d over 7d = contango = calm (H4 raw = iv90 - iv7)
Latest H4 slope +4.1 vol pts

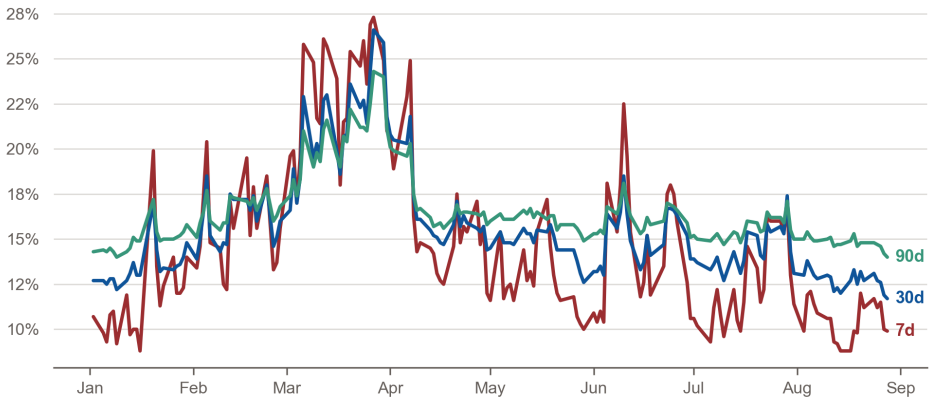


Source: Unusual Whales API; WSP Market Positioning Study, as of 2026-08-28

## 6.4 The complacency read

The 82nd-percentile reading matters because the study's pre-registered rule treats it as an overheating trigger. A bullish composite is upgraded to overly bullish when either of two legs fires: HEDGE's 5-day percentile at or above 0.80, or the amplification-risk gauge (the DEALER-side measure of how strongly dealer hedging would magnify a selloff) at or above 0.70. As of August 28 the HEDGE leg is the one firing, and it is the leg that puts the verdict at the overly-bullish boundary rather than comfortably past it: 0.8193 clears the gate by 1.9 percentile points, about 3 rank positions among 165 trading days, or 1.7 points on the score scale. The other leg is dead at 0.33, down from 0.73 six sessions earlier.

Cheap insurance is real, but the reading is carried entirely by price-of-volatility components while both demand-side components dissent. The three carriers are H1 (+74.0), H6 (+92.5; VIX options, the direct price of insuring against a volatility spike, have been cheaper on only one session all year), and H7 (+88.7). H3 and H4 sit mid-pack (+34.2, +33.0). Against them, H2 (−55.2) says the put side of the options market is near its most-bid of the year, and H5 (−47.7) says short-sale hedging pressure is elevated. Drop any one of the three carriers and the overheating label un-fires: without H1 the blend percentile is 0.7952, without H6 it is 0.7349, without H7 it is 0.7530, each below the 0.80 gate (the other four drops keep it, 0.8313 to 0.8614). At the prior cut, when HEDGE stood at the 90th percentile, the reading survived every drop tested; at the 82nd it depends on keeping all three carriers in the blend. The daily percentile on August 28 itself is 0.7711, already below the gate, and staying at or above 0.80 requires the next daily print to reach +8.8 or better, which happened on 42% of 2026

sessions. The reading is not a ranking-convention artifact (all three admissible conventions clear 0.80: 0.8193, 0.8060, 0.8642); it is sensitive to which components you keep. Complacency here means the price of insurance is low. It does not mean nobody is buying puts, and with the margin this thin, that demand-side dissent is decisive rather than color.

(Source: Leave-one-out, margin, persistence and convention figures: results/ADVERSARIAL_LOG2.md, finding 2; component scores from m3.json.)

Against the composite, HEDGE's role has changed. For most of 2026 it was the lone complacent dial while the composite sat neutral to bearish. In late August the other two dials moved toward it on the 5-day averages (SPEC +21.3, DEALER +2.5), the first such alignment of the year, and that convergence is what tips the study's verdict to bullish consensus (Chapter 14). Two facts keep the convergence in proportion. DEALER's +2.5 is barely distinguishable from zero, about 0.07 of that series' typical daily swing. And on the as-of day itself the daily prints split again: SPEC −3.1, DEALER −36.0, HEDGE +29.1, leaving HEDGE once more the only bullish print on the day. Chapters 7 and 8 carry the composite read; Chapter 15 lists the specific levels that would change it.

## 6.5 QQQ and IWM: the reduced-form check

Ticker-level HEDGE runs in reduced form for QQQ and IWM using H1 to H4 plus H7 (H5 and H6 are SPY-only; DIA lacks a harvested risk-reversal family).

| Ticker | Form | Latest 5d | 2026 pct |
|---|---|---|---|
| SPY | Full H1 to H7 | +31.4 | 0.8193 |
| QQQ | Reduced (H1 to H4, H7) | +34.8 | 0.8976 |
| IWM | Reduced (H1 to H4, H7) | +46.2 | 0.9157 |

(Source: m3.json headline_numbers; multi-year rank not published for the reduced forms.)

The complacency read is not an SPY quirk, and it is broadest in small caps: the IWM dailies printed their YTD maximum (+62.5) on 2026-08-19 (YTD minimum −66.0, 2026-06-23), and IWM's 5-day average ranks at the 92nd percentile of its year against SPY's 82nd.

**Coverage and conventions.** (1) Break audit: no HARD or SOFT breaks in any HEDGE input family; all 129 flagged series sit in the greek-exposure and options-volume families, and the refreshed audit added no flag inside the market blend. (2) Multi-year (Frame B) depth by family: IV rank, realized volatility and the risk premium from 2025-08-21 (251 rows), risk reversal from 2025-09-29, short-volume from 2024-08-22; interpolated IV is YTD-only, so H4/H7 have no distinct multi-year rank. (3) The blend's multi-year percentile (0.8494) is a component-average convention, not a rank of the blend in a real multi-year history; the blend has no pre-2026 path, and other modules' conventions differ, so cross-module multi-year comparisons are indicative only. (4) H2 splice: 158 of 165 sessions have

a monthly expiry inside the 25 to 55 day keep-window (7 gap days between the 13 roll dates; the blend skips H2 on gap days); the 2026-08-28 splice value recomputes to machine precision; one corrupt risk-reversal print on 2026-06-25 falls outside the keep-window (off-scale in the exhibit; H2 unaffected). (5) The vendor `vol_stats` family is absent from the data lake; the latest IV-rank row supplies the equivalent snapshot; no deliverable was dropped. (6) Every number in this chapter is single-vendor (Unusual Whales) with no in-sample cross-source validation, and the H3 definitional seam (2026-07-31 to 2026-08-03) sits inside the current 5-day scoring window: all five H3 inputs behind the latest reading use the extension definition.

# 7. The Composite: Positioning Consensus Index

This chapter blends the three sub-indices of Chapters 4 to 6 into one number, the Positioning Consensus Index (PCI). In human terms the three inputs are: SPEC, what the crowd is actually buying; DEALER, what market makers are forced to do; and HEDGE, what protection costs. The chapter answers two questions: what pushed the composite to where it now sits, and whether 2026 shows any build-up toward a positioning extreme. All figures use the 165-session YTD window ending 2026-08-28.

## 7.1 Construction

The composite follows the contract's fixed aggregation rule:

$$\mathrm{PCI}_t \;=\; 0.45 \times \mathrm{SPEC}_t \;+\; 0.35 \times \mathrm{HEDGE}_t \;+\; 0.20 \times \mathrm{DEALER}_t$$

where SPEC is the M1 market blend, HEDGE the M3 SPY-based market score, and DEALER the M2 cap-tier direction score (`DEALER_dir`). The weights were fixed in the pre-registered plan before any data was scored. One related quantity never enters the blend: the amplification-risk flag. When market makers are net short gamma, their hedging forces them to sell as the market falls and buy as it rises, which amplifies moves in both directions; the flag (F in the verdict rule of Chapter 14) tracks how strong that condition is, separately from direction. Attribution is exact: the audit check confirms the three weight-times-score contributions sum to the PCI with zero residual on all 165 sessions. Table 7.1 decomposes the latest session.

| Component | Weight | Contribution (2026-08-28) | Frame A pct (YTD) |
|---|---|---|---|
| SPEC | 0.45 | −1.38 | 0.4940 |
| HEDGE | 0.35 | +10.19 | 0.7711 |
| DEALER (dir.) | 0.20 | −7.20 | 0.1747 |
| PCI (sum) | 1.00 | +1.61 | 0.5301 |

Table 7.1: Latest-session attribution waterfall. Contribution = weight × sub-index score; percentiles rank each daily contribution series within the YTD window.
(Source: results/m4.json, headline_numbers waterfall_*_contrib_latest and pci_latest_day)

On the day itself, cheap insurance is the largest positive push: the HEDGE contribution of +10.19 ranks at the 77th percentile of its own YTD contribution series. Dealer direction subtracts 7.20 (17th percentile) and the crowd's buying is flat (−1.38, 49th percentile). The day nets to +1.61, near the middle of the year's range.

The 5-day averages tell a hotter story than the single day. All three sub-index 5-day averages are positive at the window's end: SPEC +21.3, HEDGE +31.3, DEALER +2.5. Because all three point the same way, the composite's sign no longer depends on how the weights split; any weighting of three positive inputs reads positive. That alignment is thinner than it looks, though. On the as-of day itself two of the three sub-indices

printed negative (SPEC −3.1 and DEALER −36.0, against HEDGE +29.1), and the dealer 5-day average clears zero by only 2.5 points, about 0.07 of that sub-index's typical daily move. One ordinary dealer session flips its sign. Chapter 14 weighs what this margin means for the verdict rule's confidence grade. (Source: results/m4.json contribution series; results/adversarial_log.json, round-2 verdict audit)

## 7.2 Level and range

Two ranking frames appear throughout. Frame A ranks a reading against every trading day of 2026 so far (the 2026 rank). Frame B ranks it against the full multi-year history (the multi-year rank). Table 7.2 places the latest prints inside the year's range in both frames.

| Measure | Value | Frame A pct | Frame B pct | Window |
|---|---|---|---|---|
| Latest daily PCI | +1.61 | 0.5301 | 0.5227 | 2026-08-28 |
| Latest 5-day average | +21.04 | 0.8494 | 0.6126 | to 2026-08-28 |
| YTD maximum | +46.15 | 0.9940 | 0.7264 | 2026-08-03 |
| YTD minimum | −48.06 | 0.0060 | 0.2661 | 2026-04-21 |

Table 7.2: Composite level and YTD range in both percentile frames.
(Source: results/m4.json, headline_numbers pci_latest_day / pci_latest_5d / pci_ytd_max / pci_ytd_min)

The 5-day average is the quantity the pre-registered verdict rule of Chapter 14 consumes; that chapter calls its 2026 rank P. As of August 28, P is 0.8494, meaning the current 5-day average ranks above roughly 85% of the year's trading days, past the rule's 0.80 bullish gate. Four pieces of context travel with that number. First, the frames disagree: the multi-year rank of the same reading is 0.6126, and at the 61st percentile the same rule would read neutral. The composite's multi-year rank is also built by the component-average convention described in the note below, not by a true ranking against pre-2026 history. Second, the reading is young. P stood at 0.3735 on August 20 and reached 0.8494 six sessions later; the rule first crossed the bullish gate on August 26, three sessions before the window's end, and the newest single day printed +1.6, the 53rd percentile. Third, the state is rare and has not aged well in 2026. Only 14 of 165 days this year classify bullish under the rule, and the earlier bullish clusters, January 27 to 28 and August 3 to 5, came just before the year's −9.1% drawdown and the mid-August fade; across the 11 bullish-class days with ten forward sessions available, SPY's mean 10-day forward return was −0.13% with an 18% hit rate (overlapping windows, descriptive, not a forecast). Fourth, the year around it is trendless: the regime table below labels the segment containing these sessions neutral and flat (mean +7.7), and the trend test in the final section finds no year-long trend ($p = 0.496$). The bullish reading is real, and it is one week old. (Source: results/m4.json, headline_numbers verdict_P; results/adversarial_log.json, round-2 verdict audit)

Smoothing shapes how the rule input behaves at extremes, in both directions. The current +21.0 average is carried by four hot sessions: August 24 to 27 printed +25.2, +23.0, +24.0 and +31.5 before the newest day cooled to +1.6. The same lag once worked in reverse: on April 21 the daily PCI hit its YTD low (0.6th percentile), yet the 5-day average stayed just above the rule's 0.20 bearish gate, so the rule read neutral on the single most bearish day of the year. (Source: results/m4.json, series PCI_A; results/adversarial_log.json, verdict-rule audit)

Frame B conventions are not uniform across modules. The blended composite has no pre-2026 path to rank against, so the M4 Frame B percentile uses the weighted component-average convention (score/200 + 0.5), consistent with M1 and M5 but not directly comparable with the M2 and M3 Frame B headline conventions. Frame A leads throughout. Separately, regime and extreme naming uses the FOMC/CPI/NFP/PCE/PPI calendar (events within ±3 calendar days of a breakpoint); the news file holds only 100 headlines timestamped 2026-08-28 and 2026-08-21, so headlines could not be attached to earlier breakpoints.

## 7.3 Monthly attribution

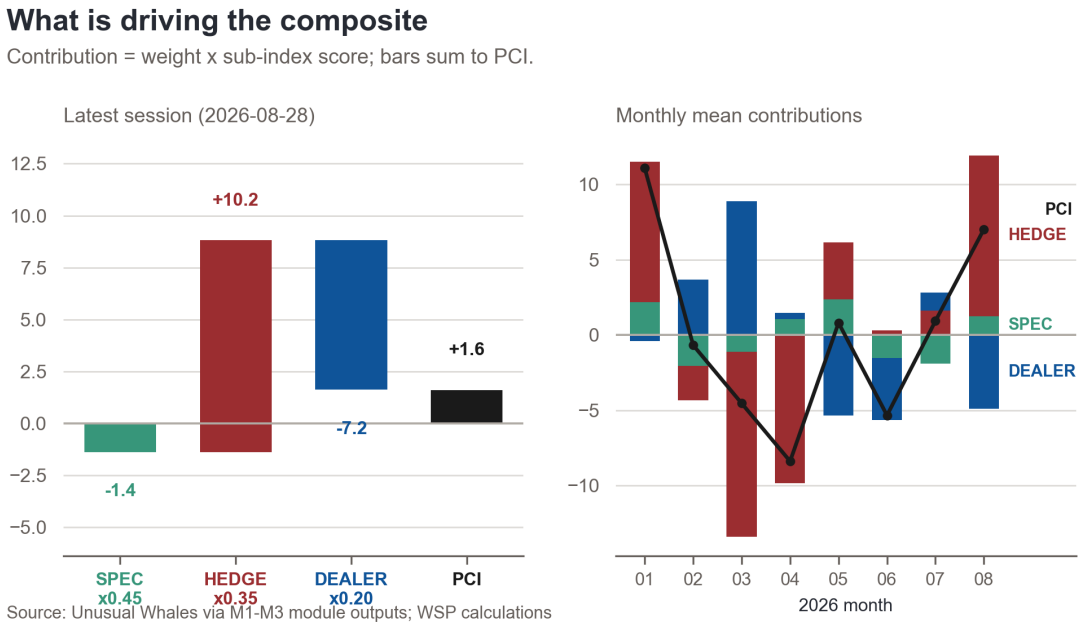


Table 7.3 averages the daily contributions by month. January remains the only month whose mean composite leaves single digits (+11.1). HEDGE is the most volatile contributor, swinging from −12.3 (March) to +10.7 (August) across monthly means. August is the strongest month since January at +7.0, and its mix is distinctive: a large positive HEDGE contribution (insurance cheap and unbid, which the inverted sign convention reads as complacency-bullish) plus mildly positive crowd buying, against a −4.9 drag from dealer direction.

## 7.4 Regimes

Regime segmentation follows the section-9 standard. It uses an algorithm that finds the dates where a data series shifts to a new level (PELT change-point detection; contract settings: rbf cost, minimum segment 20 sessions, jump 1, penalty scanned over a fixed log grid from $10^{-2}$ to $10^{3}$ with 151 points, segment count chosen as the count in the 3-to-6 range stable over the widest span of the grid). Six-segment and three-segment solutions tie for the widest stable stretch (10 grid points each, against 6 for five segments);

| Month | SPEC | HEDGE | DEALER | PCI |
|---|---|---|---|---|
| 2026-01 | +2.167 | +9.331 | −0.411 | +11.086 |
| 2026-02 | −2.066 | −2.289 | +3.676 | −0.680 |
| 2026-03 | −1.133 | −12.276 | +8.870 | −4.538 |
| 2026-04 | +1.070 | −9.846 | +0.391 | −8.385 |
| 2026-05 | +2.366 | +3.785 | −5.347 | +0.804 |
| 2026-06 | −1.536 | +0.319 | −4.118 | −5.334 |
| 2026-07 | −1.901 | +1.628 | +1.206 | +0.933 |
| 2026-08 | +1.256 | +10.672 | −4.905 | +7.024 |

Table 7.3: Monthly means of daily contributions (weight × sub-score). Rows sum to the monthly mean PCI up to rounding.

(Source: results/m4.json, series.waterfall_monthly)

the scan's selected penalty, 0.355, lands on six. Names combine level (mean PCI at or above +20 reads bullish tilt, at or below −20 bearish tilt, otherwise neutral) with the within-segment drift (Sen slope, the median pairwise slope, in points per session).

| R | Start | End | Days | Mean PCI | Pct (A) | Sen slope | Events at start |
|---|---|---|---|---|---|---|---|
| 1 | 2026-01-02 | 2026-02-02 | 21 | +10.98 | 0.7169 | −0.370 | YTD start |
| 2 | 2026-02-03 | 2026-03-04 | 21 | +0.16 | 0.5060 | +0.320 | none |
| 3 | 2026-03-05 | 2026-05-04 | 42 | −7.85 | 0.3735 | −0.111 | NFP 03-06 |
| 4 | 2026-05-05 | 2026-06-04 | 22 | +3.74 | 0.5783 | +0.901 | NFP 05-08 |
| 5 | 2026-06-05 | 2026-07-27 | 35 | −4.97 | 0.4277 | +0.164 | NFP 06-05 |
| 6 | 2026-07-28 | 2026-08-28 | 24 | +7.72 | 0.6446 | +0.454 | FOMC 07-29, PCE 07-30 |

Table 7.4: PELT regimes R1-R6 on the daily YTD PCI. Five of six classify as neutral, flat under the naming rule; R4 classifies as neutral, rising. Pct (A) ranks the regime mean within the YTD daily distribution.

(Source: results/m4.json, series.regime_table and headline_numbers regime_*_mean_pci)

All six regime means sit far inside the ±20 naming thresholds, so every regime classifies neutral on level; the consensus never established a directional regime. Five of six also classify flat, and the one exception is last spring: the May 5 to June 4 segment classifies rising, with drift of +0.90 points per session, the strongest of the year. Macro releases line up with most of the boundaries: four of the five interior regime starts fall within three calendar days of a scheduled release (an NFP report for R3, R4 and R5; the FOMC meeting and the PCE release for R6). The newest segment, July 28 to August 28, contains the late-August jump and still averages just +7.7, neutral and flat. Five hot sessions cannot form their own regime under the 20-session minimum, and inside a 24-session segment they lift the mean only modestly.

## 7.5 Chronology of extremes

Table 7.5 lists the six local extremes of the YTD path with the sub-index that moved most over the prior 10 sessions and any macro event within ±3 calendar days.

| Date | Type | PCI | Pct (A) | Driver (prior 10 sessions) | Event near |
|---|---|---|---|---|---|
| 2026-01-27 | peak | +34.15 | 0.9759 | SPEC +13.9 | FOMC 01-28 |
| 2026-03-30 | trough | −30.20 | 0.0361 | HEDGE −11.9 | none |
| 2026-04-21 | trough | −48.06 | 0.0060 | DEALER −24.5 | none |
| 2026-06-25 | trough | −34.37 | 0.0181 | SPEC −6.1 | PCE 06-25 |
| 2026-08-03 | peak | +46.15 | 0.9880 | SPEC +34.2 | none |
| 2026-08-27 | peak | +31.46 | 0.9458 | DEALER +7.6 | none |

Table 7.5: Chronology of YTD PCI extremes with dominant driver and event alignment.
(Source: results/m4.json, series.chronology_table and headline_numbers chron_1..6)

Ownership of the extremes rotates. SPEC drove three of the six, including the year's high on August 3; DEALER drove both the deepest trough of the year (April 21) and the newest peak (August 27); HEDGE drove the March 30 trough. That rotation is consistent with the low cross-correlation among the sub-indices documented in Chapter 12. Two of the six extremes fall within ±3 days of a scheduled macro release; the two largest (April 21 and August 3) do not. The newest entry, the August 27 peak (+31.46, the 95th percentile of the year's daily prints), lands one session before the window's end and feeds directly into the hot 5-day average discussed above.

## 7.6 Trajectory: trend test

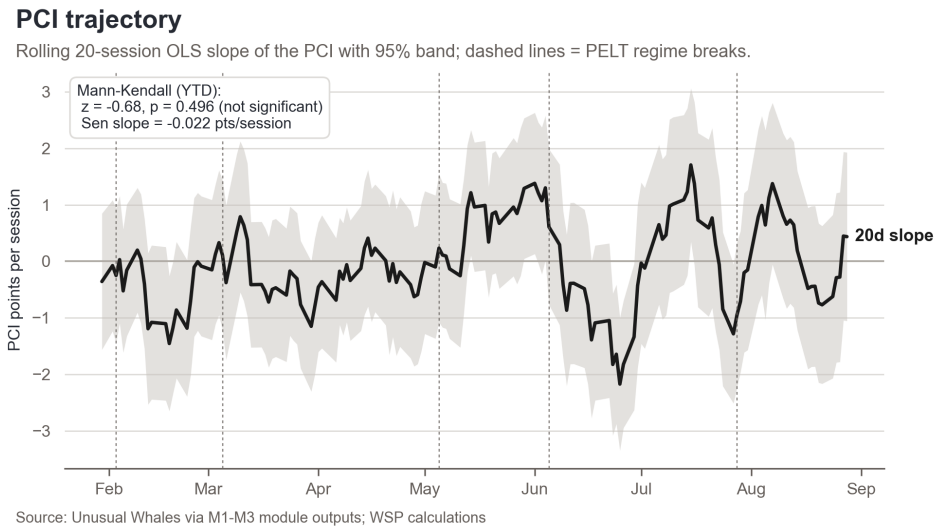


Under the contract's trend standard, a standard statistical trend test (Mann-Kendall; two-sided, alpha 0.05, tie-corrected variance, all 165 daily observations) gives $S = -484$, $z = -0.681$ and $p = 0.496$: no year-long trend. The Sen slope across the full year is $-0.022$ score points per session, tiny against daily swings that routinely exceed 20 points. The rolling 20-day slope ends the window at $+0.437$ points per session, the 72nd percentile of the year's slope readings: the latest month leans up, but within a range the year has visited repeatedly.

The composite's YTD path is oscillation, not accumulation. The trend test finds no significant trend across 2026 ($p = 0.496$), all six regimes classify neutral on level with means inside $\pm 11$ points, and the year's six extremes were driven by three different sub-indices at different times. The late-August jump that lifted the 5-day average to the 85th percentile of the year is one week old: it registers as the year's third peak, not as a new regime, and the 20-day slope it produced sits at the 72nd percentile, inside the year's visited range. What changed is the last week, not the whole tape.

Chapter 8 places the 2026-08-28 reading against every other session of the year and identifies which sub-index sits at a year-to-date extreme. Chapter 14 runs the pre-registered verdict rule on the inputs assembled here.

# 8. Where We Are Versus 2026

This chapter answers the question the study was built to answer: how stretched is options positioning right now, measured against everything 2026 has printed? The yardstick is the Positioning Consensus Index (PCI) from Chapter 7, a single score from −100 to +100, positive meaning bullish. It blends three dials: SPEC (what the option-buying crowd is actually buying), DEALER (what market makers are forced to do), and HEDGE (what protection against a fall costs). All numbers run through the August 28 close, 165 trading sessions into the year.

## 8.1 The direct answer: hot for 2026, ordinary for the longer history

The five-day average of the PCI ends August at +21.0, the 85th percentile of 2026, meaning higher than about 8 of every 10 trading days this year (the 2026 rank; the study's Frame A). Ranked against the full multi-year history of its components instead (the multi-year rank; Frame B), the same reading sits at the 61st percentile: hot within the year, unremarkable against the longer record. The August 28 print itself is +1.6, the 53rd percentile, because the last day of the late-August surge had already cooled.

| Observation | Date | PCI | 2026 pct | Multi-yr pct |
|---|---|---|---|---|
| YTD peak | 2026-08-03 | +46.15 | 0.99 | 0.73 |
| YTD trough | 2026-04-21 | −48.06 | 0.01 | 0.27 |
| Latest 5-day average | to 2026-08-28 | +21.04 | 0.85 | 0.61 |
| Latest day | 2026-08-28 | +1.61 | 0.53 | 0.52 |

(Source: m4.json headline_numbers: pci_ytd_max, pci_ytd_min, pci_latest_5d, pci_latest_day)

Positioning ended August hot by 2026's own standard (85th percentile of the year) and ordinary by the longer one (61st against the fuller component history). The heat is one week old: speculative demand jumped from the 21st to the 77th percentile of the year over August's final six sessions, and the as-of day itself had already cooled to the 53rd. The year as a whole still shows no statistical trend. What changed is the last week, not the whole tape.

## 8.2 The year ranked: notable dates by PCI

The table ranks the six inflection points from the Chapter 7 chronology by PCI level, with the latest daily print inserted for reference. Two of the year's three peaks now sit inside August.

| Rank | Date | Type | PCI | 2026 pct | Driver (prior 10 sessions) | Event near |
|---|---|---|---|---|---|---|
| 1 | 2026-08-03 | peak | +46.15 | 0.99 | SPEC +34.2 | none |
| 2 | 2026-01-27 | peak | +34.15 | 0.98 | SPEC +13.9 | FOMC 01-28 |
| 3 | 2026-08-27 | peak | +31.46 | 0.95 | DEALER +7.6 | none |
| | 2026-08-28 | latest | +1.61 | 0.53 | surge cooling (Sec. 8.3) | none |
| 4 | 2026-03-30 | trough | −30.20 | 0.04 | HEDGE −11.9 | none |
| 5 | 2026-06-25 | trough | −34.37 | 0.02 | SPEC −6.1 | PCE 06-25 |
| 6 | 2026-04-21 | trough | −48.06 | 0.01 | DEALER −24.5 | none |

(Source: m4.json series.chronology_table and headline_numbers chron_1..chron_6; latest day from pci_latest_day)

The new entry is August 27, the year's third-highest close and the first 2026 peak driven by the dealer dial rather than by crowd demand or the cost of protection: dealer positioning added +7.6 points over the prior ten sessions as market makers' hedging flows turned supportive. One caveat attaches to any ranking built on peaks: the classification rule works on five-day averages, which lag daily extremes, so peak days and rule firings do not line up one for one (Section 8.4).

## 8.3 The August pattern: a fade, then a six-session surge

August ran in two acts. Through August 20 it looked like a speculative washout: the crowd-demand dial printed −47.9 on August 17 and −47.4 on August 20, both in the bottom 3% of the year's days, and its five-day average sat at the 21st percentile of 2026. The final six sessions reversed nearly all of it. Daily SPEC swung to +35.2 by August 27, a 91st-percentile day, and its five-day average finished at the 77th percentile. The composite followed: four of the last five sessions printed daily PCI between +23.0 and +31.5, the 88th-to-95th percentile band, including the new August 27 peak. (Source: m1.json series SPEC_A, SPEC_A_pct; m4.json series PCI_A, PCI_A_pct)

The other two dials shaped the month rather than the week. HEDGE, the cost-of-protection dial, was elevated all month (+11.7 on August 4, +53.8 on August 19, +29.1 at the close); a positive HEDGE score means protection is cheap and unbid, which reads as complacency and scores as bullish consensus. DEALER swung twice: deeply negative mid-month (−64.9 on August 17), positive through the surge week (+25.6 on August 24), then back to −36.0 on August 28, the most negative dealer print of the final five sessions. (Source: m3.json series HEDGE_SPY_frameA; m2.json series DEALER_dir_market_A)

| Contribution | July mean | August mean | Aug 28 daily | Aug 28 2026 pct |
|---|---|---|---|---|
| SPEC (×0.45) | −1.90 | +1.26 | −1.38 | 0.49 |
| HEDGE (×0.35) | +1.63 | +10.67 | +10.19 | 0.77 |
| DEALER dir (×0.20) | +1.21 | −4.91 | −7.20 | 0.17 |
| PCI | +0.93 | +7.02 | +1.61 | 0.53 |

(Source: m4.json series.waterfall_monthly (2026-07, 2026-08); headline_numbers waterfall_*_contrib_latest, pci_latest_day)

August's mean HEDGE contribution of +10.67 is the largest positive monthly pull from any dial all year, and it carried the month's mean PCI to +7.02. The as-of day is less one-sided: HEDGE still in the top quarter of the year, SPEC mid-pack, DEALER in the bottom fifth, two of three dials negative on the day. The surge is real, and it was already cooling at the close.

## 8.4 Has positioning gotten more bullish?

For the year: no. For the last week: sharply, yes. The two answers are not in tension, and telling them apart is this chapter's main finding.

Across all 165 sessions, a standard statistical trend test (Mann-Kendall) finds nothing: $z = -0.681$, $p = 0.496$, with a median pairwise slope (Sen) of $-0.02$ score points per session, indistinguishable from flat. An algorithm that finds the dates where a series shifts to a new level (PELT) cuts the year into six regimes and names all six neutral; the current one, July 28 through August 28 (24 sessions), has a mean PCI of +7.7 and is labeled "neutral, flat." The latest 20-session rolling slope, +0.44, ranks at the 72nd percentile of the year's slope readings: a mild local uptilt, not a trend. The year remains oscillation, not accumulation. (Source: m4.json: mann_kendall_z, mann_kendall_p, sen_slope_per_session, rolling20_slope_latest, n_regimes, series.regime_table)

The week is another matter. The rule input $P$, the 2026 rank of the five-day PCI average, jumped from the 37th percentile on August 20 to the 85th on August 28. That move tripped the study's pre-registered classification rule (Chapter 14), which read bullish consensus on August 26, 27, and 28, its first firings since early August. Four pieces of context travel with that label everywhere it appears in this report.

- **It is a within-2026 statement.** Under the multi-year rank the same reading is the 61st percentile, and the same rule would read neutral. The multi-year figure for the composite is itself a convention (a weighted average of component ranks) rather than a rank in any lived multi-year history, because the blend has no pre-2026 path.
- **It is three sessions old and moving fast.** $P$ swung from the 37th to the 85th percentile in six sessions, and the as-of day printed +1.6, the 53rd percentile, with two of the three dials negative on the day (SPEC $-3.1$, DEALER $-36.0$).
- **It is a rare state whose in-sample history is unflattering.** Only 14 of 165 days this year classify bullish: January 2 to 8, January 27 and 28, June 1, August 3 to 5, and August 26 to 28. The late-January firings came just before the January 27 to March 30 slide of 9.1% in SPY, and the August 3 to 5 firings came just before the mid-August fade. SPY's forward 10-day return after bullish-class days averages $-0.13$% with an 18% hit rate (descriptive, 11 overlapping windows, not a forecast).
- **It coexists with a flat regime.** The same module that fires the rule labels July 28 through August 28 "neutral, flat" with a mean PCI of +7.7, and the trend test finds no year-long trend ($p = 0.496$). A one-week state inside a flat month is exactly what the numbers describe.

(Source: results/ADVERSARIAL_LOG2.md finding 1; m4.json series.regime_table)

At the dial level, no sub-index ends the month at a year-to-date extreme on the five-day view.

| Series (5-day avg) | Weight | Score | 2026 pct | Multi-yr pct |
|---|---|---|---|---|
| SPEC | 0.45 | +21.26 | 0.77 | 0.60 |
| HEDGE (SPY) | 0.35 | +31.35 | 0.82 | 0.85 |
| DEALER dir | 0.20 | +2.48 | 0.51 | 0.55 |
| PCI | | +21.04 | 0.85 | 0.61 |

(Source: m1.json spec_latest_5d; m3.json HEDGE_latest_5d; m2.json dealer_dir_latest5d; m4.json pci_latest_5d. See the frame note below.)

For the first time this year, all three five-day averages are positive at once. Two qualifiers attach. DEALER's +2.5 sits at the 51st percentile, about 0.07 of the series' typical daily swing away from zero, and the as-of day's dealer print was $-36.0$, so one ordinary dealer session flips its sign. And the composite's 85th percentile outranks every one of its own dials: three moderate positives stacked, rather than one dial at an extreme. (Source: m2.json dealer_dir_latest5d; results/ADVERSARIAL_LOG2.md finding 3)

The extreme in the system sits one level down, inside HEDGE. The VIX's own option prices sit at the 1st percentile of their one-year range (an IV rank of 5.4 on a 0-to-100 scale where 0 is the year's cheapest): insurance on volatility itself is priced near its cheapest in a year. That cheapness is also what puts the reading at the overly-bullish boundary of the pre-registered rule rather than cleanly over it. The qualifier fires on the HEDGE leg, which clears its 0.80 gate by 1.9 points (0.8193, about three rank positions out of 165 days). The rule's other trigger, amplification risk, sits far below its own gate at 0.33. That trigger tracks how strongly dealer hedging would amplify a selloff: when market makers are net short gamma, hedging forces them to sell as the market falls and buy as it rises. The condition has been the norm this year (116 of 165 sessions) but is currently milder than on two thirds of 2026's days. The firing margin is fragile: dropping any one of the three option-price-level components inside HEDGE un-fires the gate (the leg's percentile falls to 0.795, 0.735, or 0.753), and the two demand-side components lean the other way. The price gap between downside puts and upside calls scores $-55$, meaning investors are still paying up for crash protection, and the short-sale volume component scores $-48$. Cheap volatility, not absent put demand, is what the complacency label means here. (Source: m3.json H6_VIX_iv_rank_latest; m2.json fragility_F_latest, days_negative_gamma_ytd; results/ADVERSARIAL_LOG2.md finding 2)

The bullish read also leans on HEDGE structurally. In the validation module's 21 removal variants (drop one component or one dial and re-score; Chapter 12), the classification stays bullish in 20 of 21, with $P$ ranging from 0.747 to 0.8735. The single exception is removing the HEDGE dial entirely, which drops $P$ to 0.747, under the 0.80 gate, back to neutral: crowd demand and dealer positioning alone place the week high but not over the line. The weights were fixed before any data were scored, so this is disclosure rather than re-fitting, but it is why Chapter 15 names HEDGE the single series to watch, and Chapter 14 turns the full state into the study's verdict. (Source: m6.json series.breakrob_table, headline_numbers breakrob_loo_flips, breakrob_P_min, breakrob_P_max)

The multi-year columns are not one convention across dials. For the composite PCI and for SPEC, the multi-year figure is a weighted average of component ranks mapped to a 0-to-1 scale (score/200 + 0.5), because the blends have no pre-2026 path to rank against. For HEDGE it is a true rank within the multi-year daily history. For DEALER it mixes full-history components (customer delta) with 2026-only ones (NOPE, a gauge of how much option trading pushes the underlying market around, and directional delta flow). The 2026 columns (Frame A) are the comparable ones and lead throughout. Separately, every number in this chapter derives from a single vendor (Unusual Whales) with no in-sample cross-source validation; Chapter 14 prices this into the study's confidence grade.

# 9. Cross-Section: Names and Sectors

This chapter asks where the positioning actually sits, name by name, across the 23 tickers in the study universe (4 index products, VIX options, 11 sector ETFs, 7 mega-caps), and whether any single name has reached a year-to-date extreme. As of August 28, four have, and all four sit on the bullish side: the small-cap index IWM, GOOGL, MSFT, and the financials ETF XLF. Nothing in the universe sits at a bearish extreme.

## 9.1 How each name is scored

Each ticker gets a compressed version of the market positioning score from Chapter 7, called the mini-PCI. The recipe uses the same three ingredients at the same weights: what the crowd is buying in that name's options (weight 0.45), what protection costs there (weight 0.35), and what market makers are positioned to absorb (weight 0.20). Single-name data availability forces the compression. The crowd leg uses seven of the market version's nine components (S1, S2, S3, S5, S6, S7 and S9; S8 is market-wide and S4 is diagnostic). The dealer leg reduces to one series, customer delta demand (D1). The protection leg reduces to IV rank, which is where the name's option prices sit in their own one-year range (0 = cheapest, 100 = most expensive). Scores run −100 to +100; positive means the options crowd in that name is leaning bullish.

Rankings use the latest 5-day average of the daily score, the same convention as the study's verdict rule (Chapter 14), and the one-month change compares against the 5-day average ending 21 sessions earlier. Each score carries two percentiles. The 2026 rank places it against all 165 trading days of 2026 through August 28. The multi-year rank places it against the longest clean history each input allows. A name flags as a year-to-date extreme when its 5-day average lands in the top or bottom tenth of the year: a 2026 rank at or above 0.90, meaning higher than on at least 9 of every 10 trading days this year, or at or below 0.10.

Reduced-form caveats carry over from the component chapters. The S1/S5/S9 overlap and the S9 sign-flip construction disclosed in Chapter 4, and the D1 dealer-inventory assumption from Chapter 5, are inherited by every mini-PCI here. Multi-year (Frame B) depths differ by input: options volume to 2024-08-22, dealer delta to 2022-03-30 subject to per-ticker break restrictions, IV rank to 2025-08-21. Composite multi-year percentiles use M5's component-average convention (B-score/200 + 0.5) while raw-series headlines use rank/(n+1); they are not directly comparable to the Frame B conventions behind the Chapter 6 HEDGE headline. One name-specific item: QQQ's put-premium field carries a SOFT structural break dated 2026-01-28 (post/pre median ratio 2.1x) inside the YTD window, and because this module's input map includes premium fields, that break restricts QQQ's multi-year window here in a way Chapter 4's market-level read does not. All break actions restrict the multi-year window only; no series is dropped from scores. All inputs are single-vendor (Unusual Whales). (Source: results/m5.json notes; results/mod0_breaks.json, options_volume.QQQ.put_premium)

## 9.2 The full ranking: four names at the year's highs

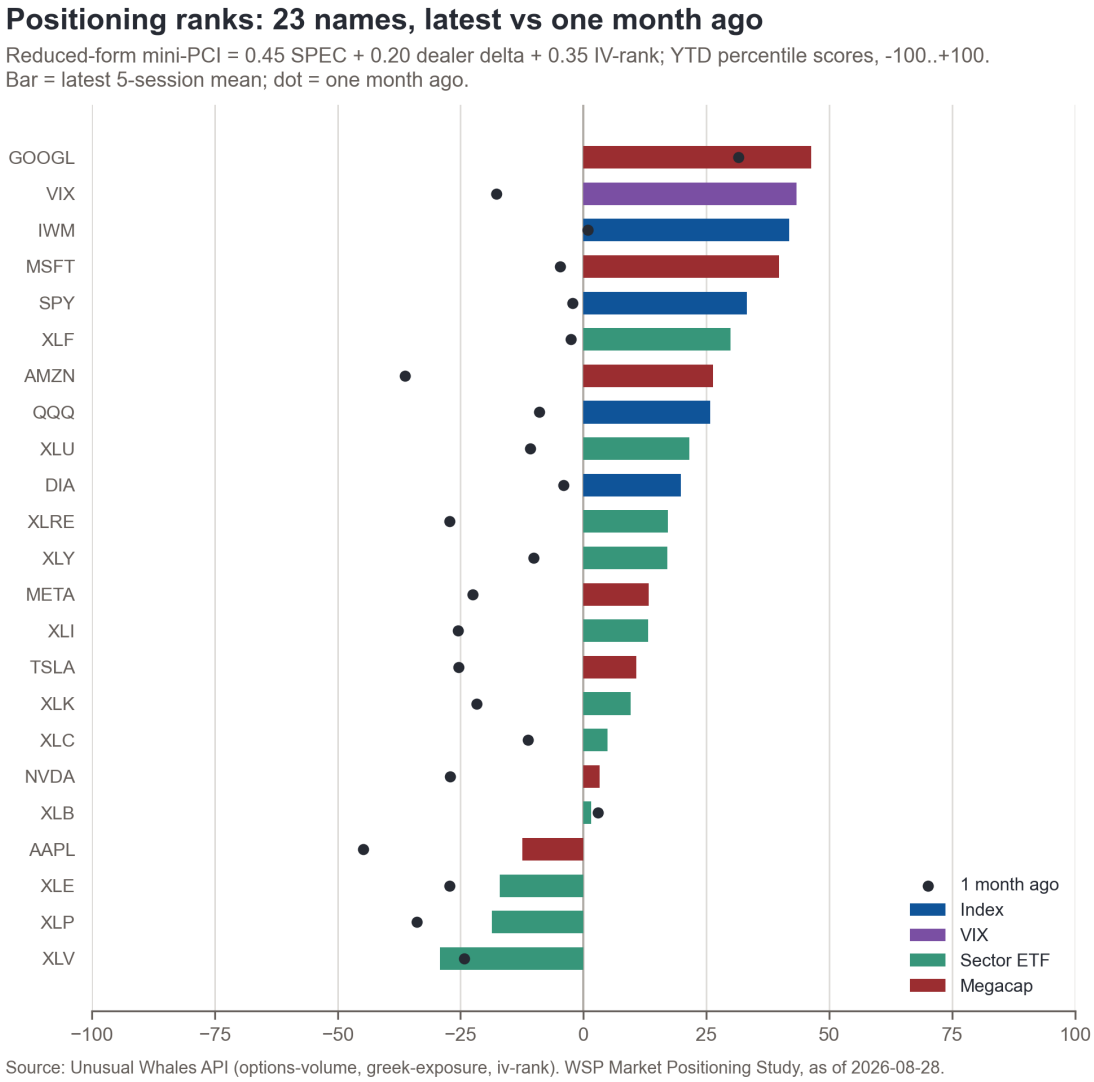


Table 9.1 gives the complete ranking as of August 28. Four readings stand out. First, GOOGL leads the universe at +46.3, its 94th percentile of 2026, and it is the only name that also led a month ago. Second, the single most stretched reading belongs not to a mega-cap but to the small-cap index IWM (95th percentile), and the index products as a group sit near the top: IWM 3rd, SPY 5th, QQQ 8th, DIA 10th. The market-level turn of late August is visible name by name, not carried by one corner of the tape. Third, four names clear the extreme bar: IWM, GOOGL, MSFT and XLF (marked * in the table). Fourth, the bottom of the table is defensives and energy: XLV is the most bearish name at −29.2 (12th percentile), with XLP and XLE just above it, and none of the three reaches the bearish extreme bar.

Rank 2 needs its own explanation. The VIX row scores positioning in VIX options themselves. The VIX gauge rises when stocks fall, so paying up for bullish VIX exposure is in effect buying exposure to turbulence, and the row reads inversely for stock-market direction. That demand jumped 61 points in a month, from 12th place to 2nd, and now sits just under the extreme bar at the 89th percentile of 2026. It is the loudest

reading in the universe that leans against the equity consensus.

Table 9.1: Mini-PCI ranking, all 23 names, as of 2026-08-28 (5-day averages)

| Rank | Ticker | Group | Score 5d | 2026 pct | Multi-yr pct | 5d, 1M ago | Rank 1M | 1M chg |
|---|---|---|---|---|---|---|---|---|
| 1 | GOOGL* | Megacap | +46.31 | 0.94 | 0.67 | +31.59 | 1 | +14.72 |
| 2 | VIX | VIX | +43.30 | 0.89 | 0.70 | −17.71 | 12 | +61.00 |
| 3 | IWM* | Index | +41.90 | 0.95 | 0.68 | +0.84 | 3 | +41.06 |
| 4 | MSFT* | Megacap | +39.75 | 0.94 | 0.76 | −4.71 | 7 | +44.46 |
| 5 | SPY | Index | +33.24 | 0.89 | 0.65 | −2.23 | 4 | +35.47 |
| 6 | XLF* | Sector ETF | +29.97 | 0.91 | 0.69 | −2.55 | 5 | +32.52 |
| 7 | AMZN | Megacap | +26.41 | 0.77 | 0.57 | −36.29 | 22 | +62.70 |
| 8 | QQQ | Index | +25.85 | 0.73 | 0.62 | −9.00 | 8 | +34.86 |
| 9 | XLU | Sector ETF | +21.51 | 0.75 | 0.56 | −10.84 | 10 | +32.35 |
| 10 | DIA | Index | +19.85 | 0.80 | 0.59 | −4.05 | 6 | +23.89 |
| 11 | XLRE | Sector ETF | +17.23 | 0.74 | 0.59 | −27.23 | 20 | +44.45 |
| 12 | XLY | Sector ETF | +17.07 | 0.77 | 0.56 | −10.18 | 9 | +27.25 |
| 13 | META | Megacap | +13.31 | 0.66 | 0.59 | −22.55 | 14 | +35.86 |
| 14 | XLI | Sector ETF | +13.16 | 0.63 | 0.48 | −25.54 | 17 | +38.71 |
| 15 | TSLA | Megacap | +10.71 | 0.61 | 0.55 | −25.39 | 16 | +36.10 |
| 16 | XLK | Sector ETF | +9.61 | 0.66 | 0.44 | −21.72 | 13 | +31.33 |
| 17 | XLC | Sector ETF | +4.87 | 0.58 | 0.44 | −11.23 | 11 | +16.10 |
| 18 | NVDA | Megacap | +3.30 | 0.56 | 0.54 | −27.09 | 18 | +30.39 |
| 19 | XLB | Sector ETF | +1.57 | 0.51 | 0.46 | +2.99 | 2 | −1.42 |
| 20 | AAPL | Megacap | −12.43 | 0.33 | 0.42 | −44.73 | 23 | +32.30 |
| 21 | XLE | Sector ETF | −17.03 | 0.25 | 0.39 | −27.18 | 19 | +10.15 |
| 22 | XLP | Sector ETF | −18.62 | 0.25 | 0.36 | −33.89 | 21 | +15.27 |
| 23 | XLV | Sector ETF | −29.15 | 0.12 | 0.35 | −24.27 | 15 | −4.88 |

(Source: results/m5.json, rank_table; * = YTD extreme under the top/bottom-tenth rule; 2026 pct = rank within 2026-01-02..2026-08-28, multi-yr pct = max break-clean depth per input; 1M chg = 5d average vs 5d average 21 sessions earlier)

## 9.3 A re-risking with breadth

The one-month column carries the breadth story: 21 of 23 names score higher than a month ago. The only two fallers are XLB, off −1.4 and down 17 places to 19th, and XLV, off −4.9 and down to last place.

The largest riser in the universe is AMZN at +62.7, which re-ranks from 22nd to 7th; the VIX row's +61.0 is second.

All seven mega-caps participate (Table 9.2). A month ago six of the seven carried negative scores; today only AAPL does, and even AAPL rose 32 points. MSFT's +44.5 takes it into the extreme band at the 94th percentile of 2026. NVDA rose 30 points yet holds rank 18, because nearly everything around it rose as much or more. The equal-weight average of the seven now stands at +18.2, the 81st percentile of 2026 and the 59th on the multi-year rank: clearly re-risked, close to but not through the extreme bar. Chapter 10 takes the same scoring into the semiconductor complex in detail.

Table 9.2: Megacap mini-PCI, one-month change (as of 2026-08-28)

| Ticker | Score 5d | 5d, 1M ago | 1M chg | Rank 1M → now |
|---|---|---|---|---|
| GOOGL | +46.31 | +31.59 | +14.72 | 1 → 1 |
| MSFT | +39.75 | −4.71 | +44.46 | 7 → 4 |
| AMZN | +26.41 | −36.29 | +62.70 | 22 → 7 |
| META | +13.31 | −22.55 | +35.86 | 14 → 13 |
| TSLA | +10.71 | −25.39 | +36.10 | 16 → 15 |
| NVDA | +3.30 | −27.09 | +30.39 | 18 → 18 |
| AAPL | −12.43 | −44.73 | +32.30 | 23 → 20 |

(Source: results/m5.json, rank_table, Megacap rows)

## 9.4 Sector flows: no extremes

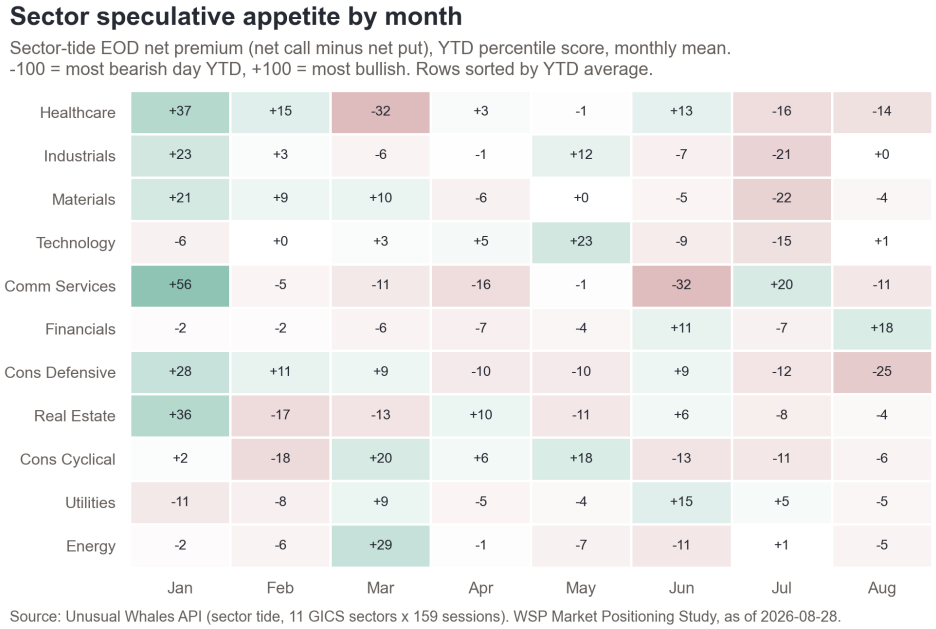


A second, independent read comes from sector tide data: the net option dollars flowing into each of the 11

GICS sectors each day, ranked against every trading day of 2026. This feed starts in January 2026, so it has no multi-year rank. Under the study's extreme rule, no sector flags in either direction.

Table 9.3 shows the latest one-day snapshot. Real estate and communication services lead; financials and utilities print at the bottom. That bottom line deserves a caution rather than a conclusion, because single-day dollar flows whipsaw. The five-day options read on the financials ETF XLF is the strongest of its year (91st percentile, an extreme flag), while the one-day flow percentile for the financial services sector on August 28 alone is near zero. One number is a week of positioning; the other is one day of dollars. The defensives are weak on both reads: healthcare and consumer defensive sit in the bottom third of the flow table, and their ETFs (XLV, XLP) hold the last two places in the positioning ranking.

Table 9.3: Sector flow percentile (2026 rank, single day, 2026-08-28)

| Sector | 2026 pct | Sector | 2026 pct |
|---|---|---|---|
| Real Estate | 0.8012 | Consumer Defensive | 0.1807 |
| Communication Services | 0.747 | Industrials | 0.1506 |
| Consumer Cyclical | 0.5301 | Technology | 0.1446 |
| Basic Materials | 0.247 | Utilities | 0.0241 |
| Energy | 0.2169 | Financial Services | 0.0181 |
| Healthcare | 0.2169 | | |

(Source: results/m5.json, sector_spec_pct_frameA, 2026-08-28 values; UW sector-tide EOD = last intraday cumulative row of each session)

## 9.5 Concentration: broad, not narrow

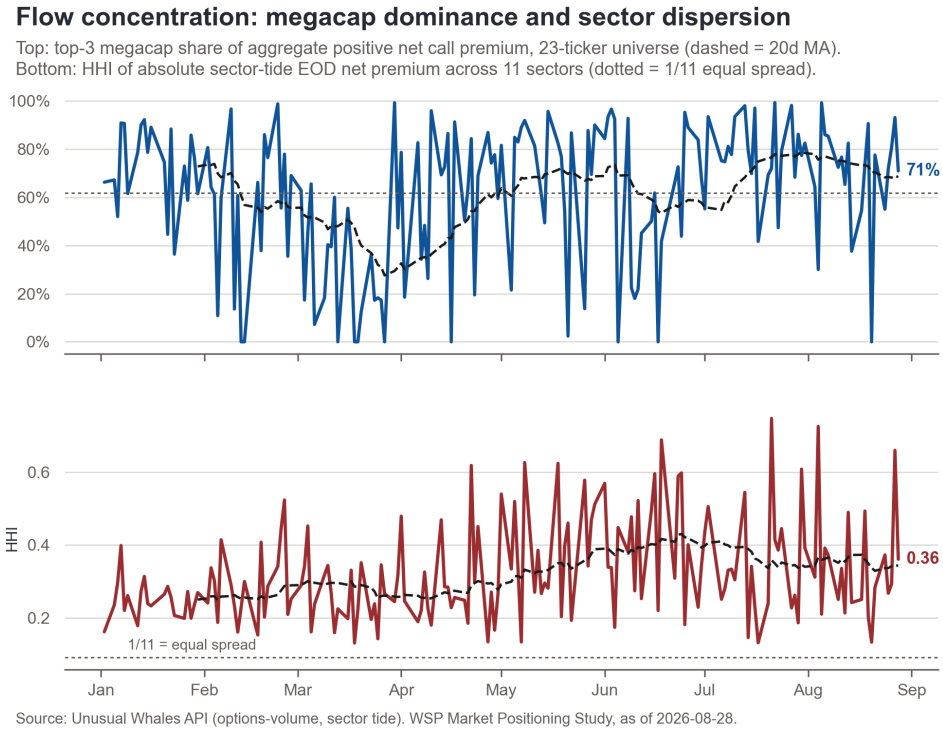


Source: Unusual Whales API (options-volume, sector tide). WSP Market Positioning Study, as of 2026-08-28.

Two dispersion checks complete the picture, and both ask the same plain question: is this a narrow squeeze into a handful of names, or a broad lean? First, the share of bullish call dollars going to the biggest three mega-caps. Measured against the full 23-ticker base, the top-3 share of positive net call premium runs 74.4% on the latest 5-day average. Three names taking roughly three-quarters of the call dollars sounds narrow, but for this universe it is normal: the 57th percentile of 2026 and the 52nd on the multi-year rank (options-volume depth to 2024-08-22). August's re-risking is not a narrowing into three names.

Second, the sector version of the same check: a standard concentration score (a Herfindahl index) on each day's sector flow shares, where higher means the dollars pile into fewer sectors. It reads 0.391 on the latest 5-day average, the 74th percentile of 2026: somewhat more piled-up than the year's norm, short of the extreme bar, and with no multi-year rank because the sector-tide feed is YTD-only.

Two construction notes, briefly. Shares clip negative net call premium to zero so they stay between 0 and 1, and use the 23-ticker base; a mega-cap-only base is stored as a secondary series but is degenerate (median 1.0), because on most days three or fewer mega-caps show positive net call premium. The concentration score uses raw flow dollars with no size weighting.

The cross-section corroborates the market-level shift and shows its breadth. 21 of 23 names score higher than a month ago; four sit at year-to-date bullish extremes (IWM at the 95th percentile of 2026, GOOGL and MSFT at the 94th, XLF at the 91st) and none at a bearish one; the index

products themselves rank near the top, while the bottom is defensives (XLV at the 12th percentile). Concentration stays ordinary, with the top-3 call-dollar share at the 57th percentile of the year, so this is a broad lean rather than a crowded squeeze into a few names. Two moderators travel with the read. Every multi-year rank in the table sits at or below the 76th percentile: stretched by 2026's standard, ordinary by the longer one. And the number-two name in the universe is VIX options, where a 61-point one-month jump in demand for volatility exposure is the loudest positioning against the equity consensus.

# 10. Semiconductors and Memory

Semiconductors carry more of the market's direction in 2026 than any other industry group, and memory is the year's supply story. This chapter maps options positioning inside the complex, name by name and cluster by cluster: which names the crowd is leaning on, whether memory really is more crowded than AI logic, and what protection costs in the group. The reason to care is concentration: if positioning anywhere is stretched enough to matter for the broad tape, it should show up here first.

## 10.1 How each name is scored

Every name gets the reduced-form positioning score introduced for the broad cross-section in Chapter 9, the mini-PCI. The ingredients are the same three dials as the market composite, in human terms: SPEC is what the crowd is actually buying, DEALER is what market makers are forced to do, and HEDGE is what protection costs. At the single-name level the last two collapse to one component each. The blend is the seven single-name speculative-demand components of Chapter 4 (weight 0.45); D1, the net market direction of customers' option positions, which market makers take the other side of (weight 0.20, Chapter 5); and H1, built on IV rank, which is where today's option prices sit in their one-year range, 0 = cheapest and 100 = most expensive (weight 0.35, Chapter 6). H1 is signed so that cheap options raise the score: when nobody pays up for protection, positioning reads as bullish consensus. Scores run −100 to +100, and every percentile in this chapter ranks the latest 5-day average within that series' own 2026 trading days, 165 sessions through August 28.

Because two of the three dials are single components here, the mini-PCI is not comparable one-for-one with the full market PCI. Where the two appear on the same chart, the market line shows direction, not level.

The reduced-form caveats of Chapter 9 carry over: the SPEC redundancy block disclosed in Chapter 4 and the D1 dealer-inventory assumption of Chapter 5 are inherited by every score here, and all inputs are single-vendor (Unusual Whales). Rankings are within-2026 only; this module stores no multi-year Frame B. The 14-name universe is the SMH ETF plus seven AI-logic names (NVDA, AMD, AVGO, TSM, MRVL, ARM, INTC), four memory names (MU, WDC, SNDK, STX) and two equipment names (AMAT, LRCX). NVDA also sits in the megacap set of Chapter 9; the clusters here are semis-internal aggregates. SNDK relisted in February 2025, so its options history is shorter than the rest of the universe (376 sessions), but it covers all of 2026 and the within-2026 rankings are unaffected. The SMH and MU risk reversals shown later are supplementary evidence, not mini-PCI inputs. One-month changes compare the current 5-day average with the 5-day average ending 21 sessions earlier. (Source: results/m7.json notes; module M7 added under amendment A5)

## 10.2 The ranking: crowded supply chain, uncrowded NVDA

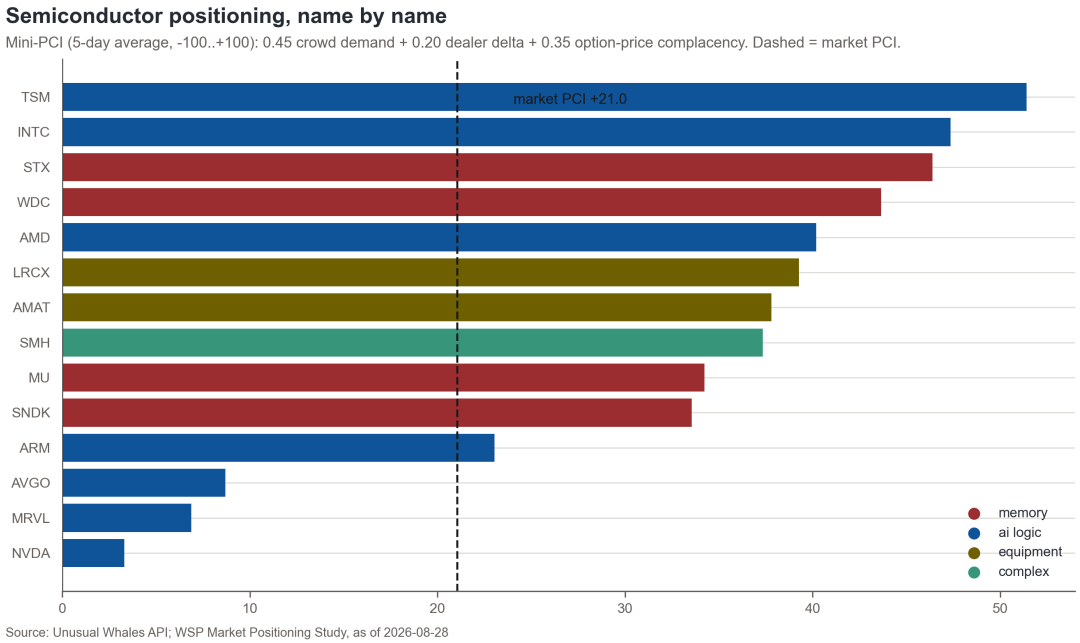


TSM tops the complex at +51.4, the 94th percentile of its own 2026 trading days, meaning a stronger reading than on roughly 9 of every 10 days this year. INTC follows at +47.3, then the hard-drive memory pair, STX at +46.4 and WDC at +43.6. At the other end sits the most-watched stock in the market: NVDA is the least crowded name in its own complex at +3.3, barely above its 2026 midpoint, with MRVL and AVGO just above it. The whole group re-risked over the past month. All 14 one-month changes are positive, from AVGO's +12.1 to TSM's +58.2 (Table 10.1).

(Source: results/m7.json, series.ranking_table. 2026 pct ranks each name's 5-day mean within that name's own 2026 dailies (2026-01-02..2026-08-28, 165 sessions), which is why the column is not monotone in rank; 1M chg = 5-day mean vs the 5-day mean ending 21 sessions earlier)

The component columns say where the heat comes from. Nine of the 14 SPEC readings are negative; a SPEC score below zero means that name's speculative buying mix runs below its own 2026 midpoint, lighter call-chasing than the name's typical day this year. What lifts the complex instead is the other two columns: dealer-absorbed direction, with D1 above +42 for seven names, and above all option prices, with H1 above +87 for nine names, meaning options near their cheapest levels of the past year. The crowd is not chasing semis with calls; semis options have simply stopped charging for fear.

Table 10.1: Semis mini-PCI ranking, all 14 names, as of 2026-08-28 (5-session means)

| Rank | Ticker | Group | Mini-PCI 5d | 2026 pct | SPEC | D1 | H1 | 1M chg |
|---|---|---|---|---|---|---|---|---|
| 1 | TSM | AI logic | +51.41 | 0.94 | +27.09 | +42.17 | +87.95 | +58.16 |
| 2 | INTC | AI logic | +47.34 | 0.92 | −2.27 | +77.35 | +93.98 | +42.91 |
| 3 | STX | Memory | +46.39 | 0.93 | −6.09 | +77.83 | +95.90 | +43.77 |
| 4 | WDC | Memory | +43.64 | 0.93 | −15.25 | +85.54 | +95.42 | +46.89 |
| 5 | AMD | AI logic | +40.18 | 0.89 | +7.23 | +28.19 | +89.40 | +52.63 |
| 6 | LRCX | Equipment | +39.27 | 0.95 | −12.74 | +57.59 | +95.66 | +49.65 |
| 7 | AMAT | Equipment | +37.79 | 0.95 | −30.05 | +94.22 | +92.77 | +46.69 |
| 8 | SMH | Complex ETF | +37.35 | 0.89 | +18.24 | +25.54 | +68.67 | +48.62 |
| 9 | MU | Memory | +34.23 | 0.84 | −2.41 | +9.16 | +95.66 | +43.22 |
| 10 | SNDK | Memory | +33.55 | 0.83 | −3.89 | +8.67 | +95.90 | +26.10 |
| 11 | ARM | AI logic | +23.04 | 0.72 | +10.36 | +19.76 | +41.20 | +46.79 |
| 12 | AVGO | AI logic | +8.70 | 0.61 | −13.39 | +77.83 | −2.41 | +12.08 |
| 13 | MRVL | AI logic | +6.88 | 0.55 | +5.92 | −8.43 | +16.87 | +35.62 |
| 14 | NVDA | AI logic | +3.30 | 0.56 | −10.50 | +7.23 | +18.80 | +30.39 |

## 10.3 Memory versus AI logic versus equipment

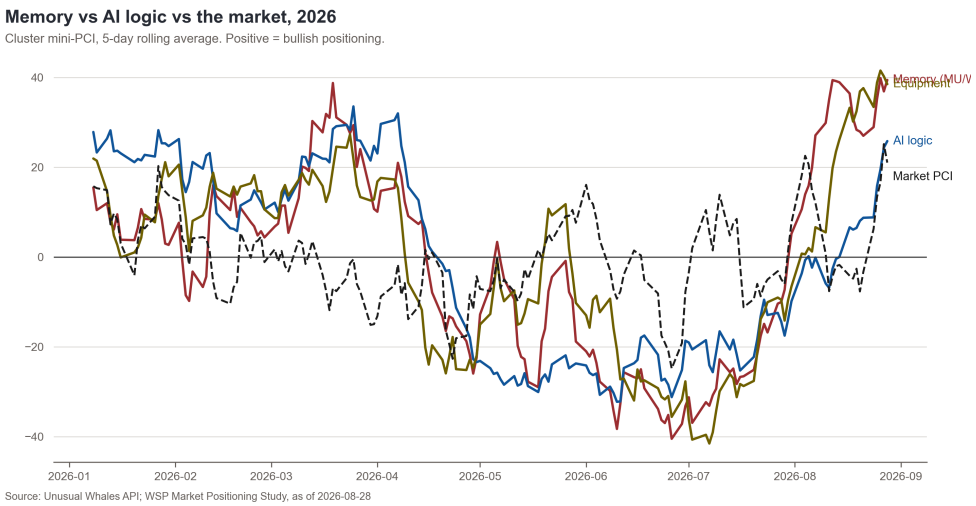


All four clusters are positive and rank in the 85th percentile of their own year or higher (Table 10.2). Memory holds the highest score, +39.5 (92nd percentile of its year). Equipment sits a hair behind at +38.5 but ranks highest against its own range (96th percentile), and it is the only cluster whose YTD daily peak came in August (+51.4 on August 20) rather than January. AI logic, home of the biggest names, is the laggard at +25.8 (85th). The all-semis aggregate stands at +32.4, its 91st percentile of 2026.

(Source: results/m7.json, headline_numbers and series.clusters_5d_rolling; peaks and troughs are daily values with dates)

Table 10.2: Cluster mini-PCI summary, as of 2026-08-28

| Cluster | Names | Mini-PCI 5d | 2026 pct | YTD daily peak | YTD daily trough |
|---|---|---|---|---|---|
| Memory | 4 | +39.45 | 0.92 | +60.99 (Jan 2) | −51.62 (Jun 5) |
| Equipment | 2 | +38.53 | 0.96 | +51.40 (Aug 20) | −54.02 (Jul 1) |
| Complex (SMH) | 1 | +37.35 | 0.89 | +66.45 (Jan 21) | −66.40 (Jun 10) |
| AI logic | 7 | +25.83 | 0.85 | +52.14 (Jan 2) | −48.38 (Jun 10) |
| All semis | 14 | +32.36 | 0.91 | +52.56 (Jan 2) | −46.92 (Jun 5) |

The year's path explains the figure. Every cluster started 2026 positive, slid to a deep trough between June 5 and July 1, and has climbed almost without interruption since late July. Memory turned positive first, on July 31; SMH and the equipment pair followed on the next session, August 3; AI logic did not hold above zero until August 14. The gap was widest in mid-August: on August 12 the memory cluster printed +39.4 while AI logic still sat at −2.6. The summary spread, memory minus AI logic on the 5-day window, now reads +13.6, the 80th percentile of its 2026 readings. Memory really is the more crowded of the two, though the gap narrowed in the final week of August as AI logic caught up, from +8.8 on August 24 to +25.8 on August 28.

The figure's market line is the full three-dial PCI of Chapter 7, a different construction, so compare direction, not level. On the same 5-day window the market composite stands at +21.0, the 85th percentile of 2026 and the 61st under the study's longer-history convention (Chapter 14 unpacks that distinction), and the all-semis line runs 11.3 points above it.

## 10.4 What protection costs in semis

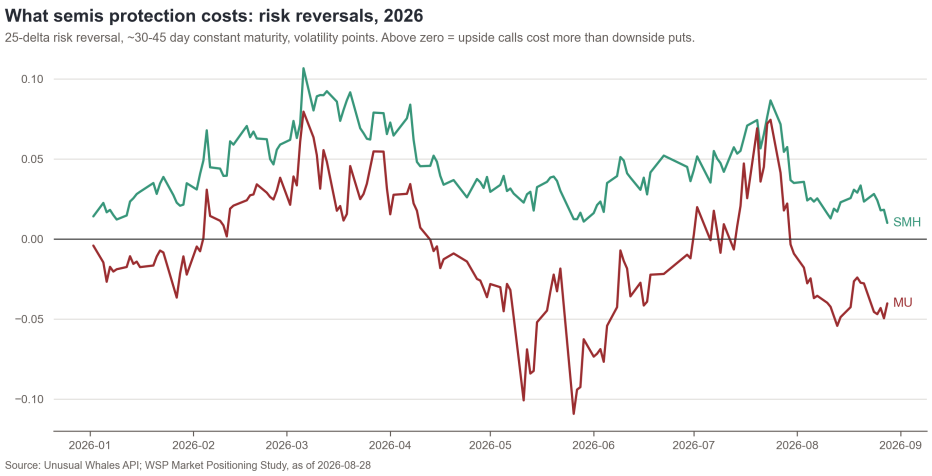


A risk reversal compares the price of upside call options with equally out-of-the-money downside puts: when puts cost more, the market is paying up for crash protection; when calls cost more, it is paying up for upside. The series are quoted in volatility points, the units option prices are quoted in, and spliced to a near-constant maturity. They sit outside the mini-PCI and serve as a direct market price of directional appetite.

SMH's 25-delta risk reversal has been positive on every 2026 session: upside calls have cost more than downside puts all year. That call premium peaked at +0.107 volatility points on March 6, was still +0.087 on July 24, and has since collapsed to +0.010 as of August 28, the lowest print of the year. MU shows the same fade with a harder edge. Its risk reversal sits at −0.040, the 18th percentile of 2026, has been negative on every August session, and bottomed at −0.109 in late May: put buyers, not call buyers, currently set the price of MU protection. Read together with the ranking table, the two series repeat the split that runs through this report: the names score bullish because options are cheap, while the direct price of upside appetite is fading, not rising.

## 10.5 Extremes and the takeaway

The study's extreme flag marks a name whose 5-day-average percentile reaches 0.90 or higher, or 0.10 or lower, within its own 2026 history. Six of the 14 names flag today, every one on the bullish side: TSM, INTC, STX, WDC, LRCX and AMAT. Five of the six carry negative SPEC scores; only TSM gets there with the crowd actually buying. For the other five, the extreme is made of cheap options and dealer-absorbed direction, not call demand.

Semiconductor positioning is the concentrated version of the market read. The all-semis composite sits at +32.4, the 91st percentile of 2026, against +21.0 for the full-market PCI (85th percentile of 2026, 61st under the longer-history convention), and six of 14 names flag at bullish year-to-date extremes. Memory is the highest-scoring cluster, all four memory names rank in the top ten, and the memory-minus-AI-logic spread sits in the top fifth of its year. The composition carries the warning documented in Chapters 6 and 14: the lean rests on cheap option prices and supportive dealer positioning, speculative buying is below its 2026 midpoint for nine of 14 names, SMH's call-over-put premium is the year's lowest, and the least crowded name in the complex is NVDA. The complex is crowded by 2026's standard, and the crowding is built on complacency rather than chasing.

# 11. Short-Dated Versus Long-Dated: Positioning by Time to Expiry

An option that expires this afternoon and an option that expires next year are bought by different crowds for different reasons, and they can tell different stories on the same day. This chapter sorts the evidence as of August 28 by time to expiry (days to expiry, DTE; options traders call this tenor) and asks whether the short-dated and long-dated pictures agree. It also states plainly where the planned breakdown could not be reported.

Three study terms repeat below, in human terms: SPEC tracks what the option crowd is actually buying, DEALER tracks what market makers are forced to do, and HEDGE tracks what protection costs.

## 11.1 The layered design and delivered coverage

The research plan (section 5) splits the option board into five expiry buckets; section 7 specifies the layered read as descriptive parallel tracks, not a re-blend of the composite.

| Bucket | Definition | Role in the plan |
|---|---|---|
| 0DTE | expires the same day | headline layer |
| SHORT | 1 to 28 DTE | headline layer |
| MID | 29 to 91 DTE | headline layer |
| OTHER | 92 to 365 DTE | reported in tables, excluded from the four headline layers |
| LEAP | over 365 DTE | headline layer |

Table 11.1. DTE bucket partition as contracted in RESEARCH_PLAN section 5. Bucketing was to come from the expiry-breakdown endpoint (volume and open interest per expiry, daily), with the option-contract screener supplying the current-day cross-section.

Coverage limit. The full layered read specified in plan section 7 was not carried into the results files: no results JSON contains a per-bucket series, and the endpoint inventory (probe_endpoints.json) stores the option-contract screener probe as schema only (a 20-row sample, 21 to 90 DTE filter, field names with no retained values) and contains no expiry-breakdown probe at all. Under this report's number contract every figure must come from a results JSON, so no quantity is quoted for options that expire the same day (0DTE) or for options with more than a year to run (LEAPs), and the planned 0DTE-versus-LEAP contrast is not quantified. What follows is the tenor evidence embedded in the delivered series (Table 11.2).

## 11.2 Tenor windows in the delivered series

The flow window is bullish on the week. The 5-day average of SPEC sits at +21.3, the 77th percentile of 2026, higher than about three of every four trading days this year (60th on the multi-year rank). The

| Series | Latest | 2026 rank | Multi-yr rank |
|---|---|---|---|
| *All-expiry transacted flow (tenor-aggregated)* | | | |
| SPEC sub-index, 5-day average | +21.26 | 77% | 60% |
| S4 volume-vs-premium wedge (diagnostic) | +0.547 | 80% | n/a |
| *Nearest expiry and dealer book (snapshot 2026-08-28)* | | | |
| Max pain, 2026-08-31 expiry | 760.0 | n/a | n/a |
| Gamma-flip strike, interpolated | 769.62 | n/a | n/a |
| Spot distance to flip | −0.06% | n/a | n/a |
| Net SPY gamma exposure | −12,649.7 | 70% | 72% |
| NOPE, session close | −5.4673 | 13% | n/a |
| *25 to 55 DTE, spliced constant-maturity* | | | |
| H2 25-delta risk reversal | +0.02605 | 18% | 16% |
| *30d tenor* | | | |
| H7 implied-move component score | +98.80 | | score |
| *7d versus 90d tenor* | | | |
| H4 term slope, iv(90d) − iv(7d) | +0.041 | 78% | n/a |
| H4 component score | +55.42 | | score |
| *VIX complex (tenor-free cross-check)* | | | |
| H6 VIX iv rank | 5.45 | 1% | 4% |
| VIX-options mini-PCI, 5-day average | +43.30 | 89% | 70% |

Table 11.2. Tenor windows in the delivered results, as of 2026-08-28. The 2026 rank is where the latest value stands against every trading day of 2026 so far (165 sessions); the multi-year rank is against the full harvested history, whose depth varies by series. Component scores are already percentile transforms, score $=$ (pct $-$ 0.5) $\times$ 200, so no separate percentile is repeated for them. The multi-year rank is not distinct from the 2026 rank for H4 and H7 because interpolated IV exists YTD only; NOPE has no multi-year depth; snapshot metrics (max pain, flip, distance) have no YTD distribution. Net gamma exposure is quoted in raw endpoint units. Max pain is reported for the nearest expiry strictly after the as-of date. The VIX-options row is positioning on VIX options itself, where bullish means demand for volatility, not for stocks. (Source: m1.json, m2.json, m3.json, m5.json headline_numbers and series)

nearest-expiry mechanics no longer push against it. Dealer direction, the 5-day average of what market makers are forced to do, is +2.5, the 51st percentile of the year (55th multi-year), roughly the middle of the range (Chapter 5). Net gamma, a measure of whether market makers' hedging dampens market moves (positive) or amplifies them (negative), is still negative, and that is the 2026 norm rather than an anomaly: 116 of 165 sessions this year have printed negative SPY net gamma. But the book is closer to flat than on about seven of ten days this year (70th percentile), and the study's amplification-risk gauge, which ranks how strongly that hedging would magnify a selloff, has fallen from the 73rd percentile of the year on August 20 to the 33rd on August 28 (29th multi-year). Spot closed 0.06% below the gamma flip, the price level where dealer hedging switches from dampening moves to amplifying them, so the market sits essentially on top of that line. Max pain for the August 31 expiry, the strike where the largest amount of open option value would expire worthless, sits below the market at 760. The one soft mechanics print is NOPE, a measure of how much option trading is pushing around the stock itself: $-5.5$, the 13th percentile of the year.

The measured insurance tenors are cheap at the level and split underneath. Taken together, the seven protection components average +31.4 for the week, the 82nd percentile of 2026 (85th multi-year), the cheap-and-unbid read of Chapter 6. Every price-level measure agrees. The VIX iv rank, where the cost of options on VIX itself sits in its one-year range (0 = cheapest, 100 = most expensive), is 5.4, the 1st percentile of the year (4th multi-year). The 30-day implied-move component, which scores how small an index move the options market is pricing for the next month, is at +98.8, smaller than on nearly every other day this year. The term slope, whether longer-dated options cost more than short-dated ones (calm) or less (stress), is upward at +0.041, the 78th percentile; the 7-day point has exceeded the 90-day point on only 38 of 165 sessions this year (YTD extremes $-0.058$ on 2026-03-09 and +0.061 on 2026-01-16). The split shows up in demand. The 25-delta risk reversal, the price of upside call options versus equally out-of-the-money downside puts, where a falling reading means the market is paying up for downside protection, has been positive on all 158 YTD sessions with a valid splice (7 of 165 sessions fall between monthly rolls and are skipped; 13 roll dates YTD), so calls have out-priced puts all year. Today's +0.026 sits at the 18th percentile of that range (16th multi-year): the call tilt is thinner than on four of every five days this year, between the YTD extremes of +0.01773 on 2026-08-04 and +0.10753 on 2026-03-06. The second demand read runs the same direction. VIX-options positioning, the one cross-check with no tenor at all, is at +43.3, the 89th percentile, the second-highest positioning read in the 23-name cross-section behind GOOGL, up 61 points and ten rank places in a month (Chapter 9). Buyers are paying up for volatility exactly where volatility is priced cheapest.

## 11.3 What agrees and disagrees on 2026-08-28

Within the flow window, the week and the day disagree. Five of the eight blended SPEC components printed positive on August 28: the two put/call ratios, the net premium tilt, the bullish premium share, and the volume-surge score (Chapter 4). The three negatives are the two trade-initiation reads, which track whether buyers are hitting the seller's asking price (buyer-initiated) or sellers are hitting the bid, and the out-of-the-money tilt (S6 $-24.7$, S7 $-53.2$, S8 $-22.9$). The counts and the dollars say calls; the initiation reads say buyers were paying up for puts on the day. The daily SPEC print itself was $-3.1$, the 49th percentile, at the end of a six-session stretch that ran from $-47.4$ on August 20 (the 2nd percentile of the

year) to +35.2 on August 27 (the 91st). The bullish flow window is one week old, and the as-of day did not extend it. The dealer window shares the recency problem: its 5-day average is +2.5, but the August 28 daily print was −36.0, the only negative print among the last five sessions and the most negative of them.

The S4 wedge, the study's standing diagnostic for cheap short-dated contracts inflating contract counts relative to dollars (Chapter 3), is the tenor signature inside that flow. It sits at +0.547, the 80th percentile of 2026, and has been positive on each of the last six sessions (+0.354, +0.066, +1.247, +0.560, +0.603, +0.547). Contract counts are running hotter than premium dollars, the pattern that appears when small, cheap, mostly near-dated contracts carry the count. Across the year the wedge has been direction-mixed (91 positive days, 74 negative; YTD extremes +1.695 on 2026-01-02 and −1.581 on 2026-04-08), so a six-session positive run at the 80th percentile is a lean, not the usual back-and-forth.

The year's sharpest quantified same-day tenor divergence points at the same mixture. On 2026-08-04, the day SPEC printed its YTD peak of +61.1 (99th percentile), the 25 to 55 day risk reversal printed its YTD minimum of +0.01773: record crowd bullishness in the flow window against the year's strongest put bid in the mid-tenor window. The wedge that day was +0.913, the cheap-contract signature at nearly full strength. The August 28 configuration rhymes with it at lower intensity: SPEC 5-day at the 77th percentile rather than the 99th, the wedge at +0.547 rather than +0.913, the risk reversal at the 18th percentile rather than the floor.

At the weekly scale the tenor windows lean the same way: flow bullish (SPEC 77th percentile), the dealer book near the flat line rather than deep in amplification territory (amplification risk at the 33rd percentile, spot 0.06% below the flip), and insurance cheap at every measured tenor (VIX iv rank at the 1st percentile, 30-day implied move near its YTD smallest, term slope calm at the 78th). The disagreements live one level down. The 25 to 55 day skew sits in the bottom fifth of its 2026 range, the initiation reads inside the day's flow favored puts, the count-versus-dollar wedge is at the 80th percentile, and the VIX-options crowd holds the second most crowded position in the cross-section against the year's cheapest volatility pricing. That mixture, at higher intensity, marked the August 4 speculative peak. In maturity terms it restates the recency and demand-versus-level qualifiers that travel with the Chapter 14 verdict.

A fully layered read remains buildable from the harvested endpoints (expiry-breakdown bucket shares, a per-expiry 0DTE snapshot, 180d and 360d pairs extending H4); until then, claims about same-day expiries and LEAPs stay out of reach.

# 12. Validation: Stress-Testing the Composite

The composite now reads bullish. This chapter asks how much weight that reading can carry. It runs three tests: whether the three sub-indices measure different things, whether the composite's daily readings told you anything about the returns that followed, and whether the verdict survives removing any single ingredient.

Two boundaries frame everything here. Every number is built from the same Unusual Whales data as Chapters 1 through 11; Chapter 13 repeats the key exercise on public, non-vendor data reaching back to 1990. And forward returns can only use the vendor's one-year window of regular-session SPY closes (2025-08-28 to 2026-08-28, amendment A1), so the last entry date with a complete 63-day forward window is 2026-05-29. The registered power statement applies throughout: 165 trading days is enough to rank days against each other, and not enough to certify a return forecast.

## 12.1 Coherence: three dials, not one

A reasonable first worry is that the three sub-indices are three names for one signal. They are not. As a reminder of what each dial measures in human terms: SPEC tracks what the options crowd is actually buying, HEDGE tracks what downside protection costs, and DEALER tracks what market makers are forced to do. The average pairwise correlation of their daily scores this year is 0.043, essentially zero. A standard dimensionality check (principal component analysis, which finds the single pattern explaining the most day-to-day variation) agrees: across the 18 underlying components, the largest common pattern explains 36.4% of daily variance over the 158 complete days. The blend carries three genuinely separate readings, not one reading rebranded three times.

| | PC1 | PC2 | PC3 | PC4 | PC5 |
|---|---|---|---|---|---|
| Variance share | 36.45% | 22.81% | 6.87% | 5.88% | 4.69% |

| Correlation (daily scores, YTD) | PCI | SPEC | HEDGE | DEALER |
|---|---|---|---|---|
| PCI | 1.000 | 0.887 | 0.513 | 0.383 |
| SPEC | | 1.000 | 0.191 | 0.400 |
| HEDGE | | | 1.000 | −0.463 |
| DEALER | | | | 1.000 |

(Source: results/m6.json, series.coherence_corr and series.pca_var_share_top5)

Two facts in that table matter later. First, HEDGE and DEALER lean against each other (correlation −0.463): days when protection is cheap tend to be days when dealer positioning reads bearish, so those two dials rarely agree by accident. When they do agree, as they did in the week behind the current verdict, that is unusual, and Chapter 14 weighs how much to trust it. Second, the composite tracks SPEC closely (0.887),

a direct result of SPEC's 0.45 weight: the composite mostly moves when the crowd's buying moves.

**Coherence: correlation of PCI, sub-indices and components**

Daily Frame A scores, YTD 2026; PC1 = 36% of component variance (158 days x 18 components)

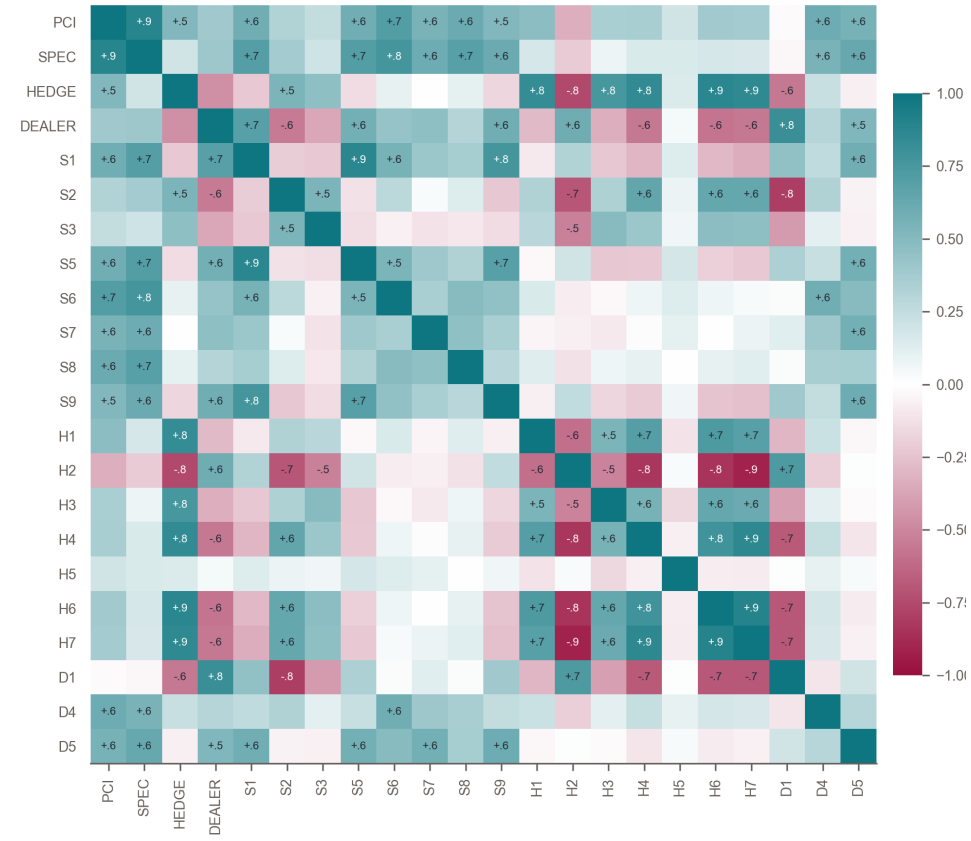


Source: Unusual Whales; WSP Market Positioning Study M6

Inside the SPEC block the components are less independent than the count of eight suggests. The adversarial review flagged S1, S5, and S9 as close duplicates (score correlations 0.873, 0.768, and 0.684 for the three pairs), and S9 takes its sign from S1 by construction, so the direction of net options premium effectively receives about 3/8 of the SPEC blend. S9 is also discontinuous by design: it jumps whenever net premium changes sign. Separately, D1's economic sign rests on the vendor's assumption about dealer inventory, which no second source can verify. None of this changes any sub-index sign, and the round-2 adversarial pass re-tested the trio against the refreshed verdict: collapsing the three duplicates into one voice does not change the verdict class. The honest summary is that SPEC has fewer than eight independent voices.

## 12.2 Forward returns by composite quintile

The demanding test is whether the composite's readings said anything about what SPY did next. All 165 sessions were sorted into five equal buckets of 33 by each day's composite rank within 2026, so Q5 holds the year's most bullish fifth of days and Q1 its most bearish fifth. Forward SPY returns were then measured 5, 10, 21, and 63 sessions ahead. Later horizons have fewer usable entry dates because the forward window must close before August 28: 160, 155, 144, and 102 entries. Error bars are built by resampling

whole 10-day stretches of the year rather than single days, which respects the fact that markets trend (a block bootstrap; 5,000 draws, reproducible under reruns). The FDR column marks cells that stay significant after correcting for the fact that 20 cells were tested at once (Benjamini-Hochberg at $q = 0.10$).

| Horizon | Q | $n$ | Mean % | Median % | Hit | $t$ | 95% CI | $p_{boot}$ | FDR |
|---|---|---|---|---|---|---|---|---|---|
| 5d | Q1 | 33 | 0.97 | 1.08 | 66.7% | 3.01 | [0.70, 1.95] | 0.0004 | yes |
| | Q2 | 33 | 0.42 | 0.40 | 57.6% | 1.18 | [−0.36, 1.29] | 0.2559 | |
| | Q3 | 32 | 0.27 | 0.28 | 59.4% | 0.94 | [−0.08, 0.94] | 0.1308 | |
| | Q4 | 33 | 0.01 | −0.39 | 42.4% | 0.03 | [−0.63, 0.78] | 0.9290 | |
| | Q5 | 29 | 0.08 | −0.10 | 48.3% | 0.29 | [−0.45, 0.13] | 0.2975 | |
| 10d | Q1 | 32 | 1.65 | 1.50 | 71.9% | 3.11 | [1.18, 3.40] | 0.0004 | yes |
| | Q2 | 31 | 0.73 | 0.08 | 51.6% | 1.30 | [−0.35, 2.56] | 0.1712 | |
| | Q3 | 32 | 0.92 | 0.54 | 59.4% | 2.16 | [−0.04, 2.27] | 0.0588 | |
| | Q4 | 32 | 0.55 | 0.00 | 50.0% | 1.10 | [−0.75, 2.19] | 0.3723 | |
| | Q5 | 28 | −0.40 | −0.81 | 28.6% | −1.10 | [−1.18, −0.19] | 0.0108 | yes |
| 21d | Q1 | 31 | 2.69 | 2.86 | 77.4% | 3.21 | [1.35, 6.07] | 0.0004 | yes |
| | Q2 | 29 | 2.95 | 2.95 | 79.3% | 4.04 | [1.31, 6.38] | 0.0004 | yes |
| | Q3 | 28 | 1.07 | 0.26 | 57.1% | 1.30 | [−1.40, 3.83] | 0.3319 | |
| | Q4 | 30 | 1.00 | −0.27 | 46.7% | 1.25 | [−1.64, 4.21] | 0.3951 | |
| | Q5 | 26 | 0.32 | −0.31 | 42.3% | 0.57 | [−1.03, 0.76] | 0.9150 | |
| 63d | Q1 | 20 | 8.57 | 9.48 | 100.0% | 8.80 | [7.27, 12.02] | 0.0004 | yes |
| | Q2 | 20 | 8.07 | 7.35 | 95.0% | 8.02 | [8.07, 11.59] | 0.0004 | yes |
| | Q3 | 21 | 6.08 | 4.73 | 100.0% | 7.74 | [6.04, 8.89] | 0.0004 | yes |
| | Q4 | 24 | 6.26 | 6.45 | 91.7% | 7.18 | [4.80, 10.16] | 0.0004 | yes |
| | Q5 | 17 | 3.82 | 2.88 | 76.5% | 3.13 | [1.83, 7.71] | 0.0004 | yes |

(Source: results/m6.json, series.fwd_table_uw; mean forward SPY return per cell, UW regular-session closes. Hit = share of entries with a positive forward return.)

Ten of the 20 cells clear the multiple-testing gate (ten at the stricter $q = 0.05$, eleven at $q = 0.20$), and the raw pattern runs contrarian: the most bearish days preceded the strongest forward returns at every horizon, and the most bullish days preceded the weakest. The only significant negative cell is Q5 at 10 days: after the year's most bullish readings, SPY averaged −0.4% over the next two weeks and rose only 29% of the time. Q5 minus Q1 spreads are negative at every horizon (−0.9% at 5 days, −2.1% at 10, −2.4% at 21, −4.8% at 63) and carry the same sign on the other index tickers: at 21/63 days, QQQ −4.2%/−7.5%, IWM −0.9%/−5.2%, DIA −0.5%/−5.0%. (Source: results/m6.json, series.fwd_spreads_uw and series.fwd_spreads_other_tickers)

The reading behind the current verdict sits in exactly the bucket this table is hardest on: the five sessions from August 21 through August 27 all rank in the year's most bullish fifth, while the August 28 close itself falls back to the middle fifth. The next section asks whether the in-sample Q5 pattern deserves forecast weight; the answer is mostly no.

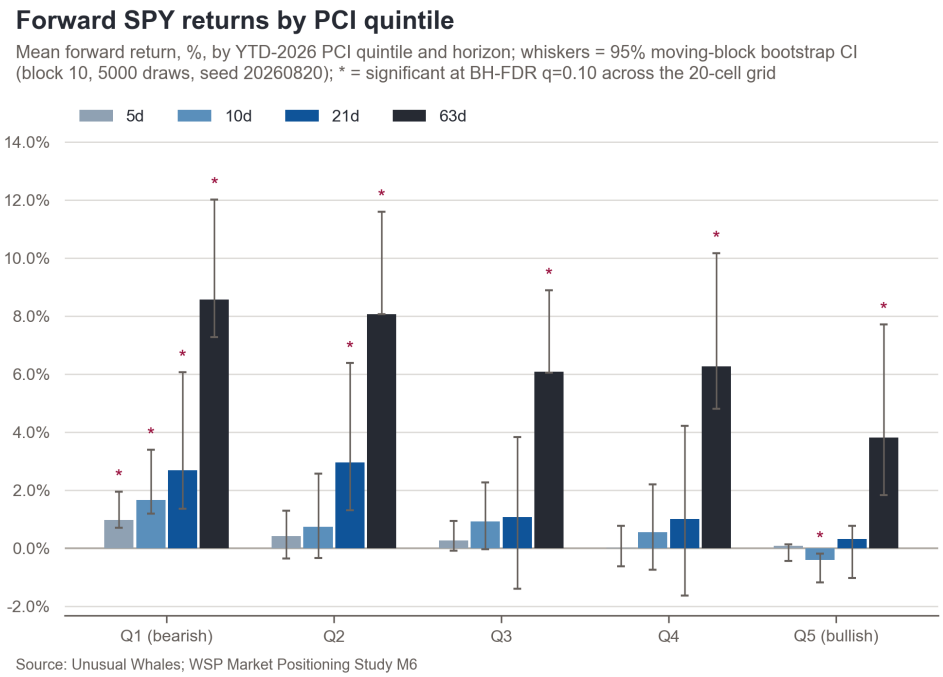


## 12.3 Placebo, benchmark, and the one-episode problem

Before reading the quintile table as a tradable contrarian signal, run two reference tests. The placebo reassigns the quintile labels to random dates 1,000 times and asks how often chance alone produces a spread that large. The benchmark is a rule that knows nothing about options: it sorts days only by whether SPY fell over the trailing 63 sessions.

| Horizon | PCI Q5−Q1 | Placebo mean | Placebo sd | $p$ (two-sided) | Momentum spread | $n$ pos/neg |
|---|---|---|---|---|---|---|
| 5d | −0.886 | 0.012 | 0.468 | 0.065 | −0.443 | 129/31 |
| 10d | −2.055 | 0.006 | 0.682 | 0.001 | −1.405 | 124/31 |
| 21d | −2.366 | −0.039 | 1.090 | 0.039 | −5.348 | 113/31 |
| 63d | −4.754 | −0.007 | 1.401 | 0.001 | −7.511 | 71/31 |

(Source: results/m6.json, series.placebo and series.momentum_benchmark; momentum spread = mean forward return after trailing-63d-down days minus after trailing-63d-up days)

The placebo p-values look strong. The benchmark column undercuts them: at 21 and 63 days the naive momentum rule produces larger spreads than the composite (−5.3% versus −2.4%, and −7.5% versus −4.8%). A rule built from nothing but the fact that the market just fell beats the composite at its own game at the horizons that matter. Whatever the quintile table captures, it is mostly the fact that 2026 mean-reverted.

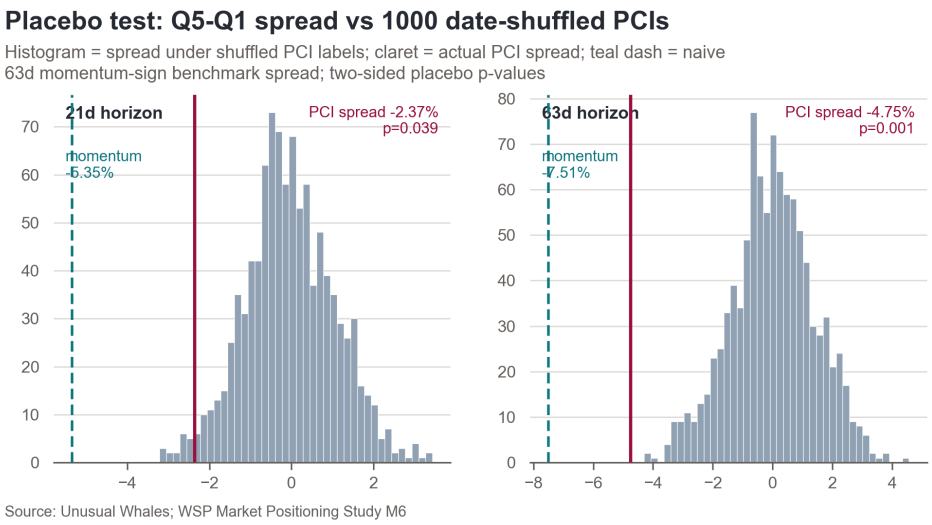


Decomposing the year shows why. 2026 contains one dominant decline and rebound: the composite peaked on January 27 (+34.2), SPY then fell 9.1% into the March 30 trough, and the recovery ran into early May (the regime segmentation in Chapter 7 puts the boundary at May 4). The bearish-consensus bucket lives inside that event: 16 of Q1's 33 days fall between the January 27 peak and May 4, so buying the average Q1 day largely meant buying the dip of the year's one big decline. The 63-day horizon cannot escape the episode at all: its entry dates end on May 29, and the same slide sets the worst case in every bucket. A year like this contains roughly seven non-overlapping 21-day windows and offers no independent second event to test against. The date-shuffle placebo destroys exactly this episode-and-trend structure, so its p-values, and the 10-of-20 significance count, overstate how much evidence one year can supply.

The forward-risk table points the same way. If the year's most bullish days had genuinely flagged fragile tape, Q5 should show meaningfully deeper drawdowns and wilder markets over the next month. It barely does: the mean worst drawdown deepens from −3.1% (Q1) to −3.7% (Q5), a gap of half a point, and forward realized volatility shows no ordering at all (13.3% to 14.3% annualized, with Q2 the highest). At 63 days a gradient exists (Q1 −5.6% to Q5 −7.3%), but every quintile shares the identical −9.1% worst case, because every quintile's windows cross the same episode.

| 21d forward risk | $n$ | Mean max DD % | Worst DD % | Fwd RV % ann. |
|---|---|---|---|---|
| Q1 | 31 | −3.12 | −7.93 | 13.3 |
| Q2 | 29 | −3.25 | −5.79 | 14.3 |
| Q3 | 28 | −3.37 | −8.01 | 13.6 |
| Q4 | 30 | −3.42 | −7.76 | 13.5 |
| Q5 | 26 | −3.66 | −7.76 | 14.1 |

(Source: results/m6.json, series.fragility_table; max forward drawdown within the 21d window, forward realized vol annualized)

The quintile table reads contrarian: the year's most bullish positioning days preceded its weakest forward returns. That reading does not survive decomposition. It rests on one drawdown-and-rebound episode, and a naive rule that simply fades recent weakness produces larger spreads at 21 and 63 days. This chapter therefore validates the composite as a ranking of positioning, not as a return forecast. The verdict in Chapter 14 inherits that boundary.

## 12.4 Break robustness and leave-one-out

The break audit in Chapter 3 flagged 52 series whose history shifts level hard enough to distort a multi-year ranking. Zero of the 52 enter the market composite: the index-ticker flags sit on the vanna and charm series (shown descriptively only), the gamma flags feed only the amplification-risk axis (whether dealer hedging would amplify a move rather than dampen it, reported alongside the composite but never blended into it), and the rest sit on single megacap names outside the SPY/QQQ/IWM/DIA blend. Excluding any one of the 52 therefore reproduces the composite exactly and the verdict cannot move by construction: an identity, not robustness earned under stress, and recorded as such.

Because that check cannot fail, the module runs a real one: rebuild the composite with one ingredient removed and the remaining weights re-scaled, once for each sub-index and each SPEC and HEDGE component, plus a SPY-based approximation for the three DEALER components (the module stores no per-component market-level DEALER scores). Each variant is then scored by the study's pre-registered verdict rule, whose input $P$ is the composite's five-day average rank within 2026: bullish consensus at or above the 80th percentile of the year (higher than 8 of every 10 trading days of 2026), bearish at or below the 20th, neutral between.

| Variant | Rule input $P$ | Class |
|---|---|---|
| Baseline (M4, weights .45/.35/.20) | 0.8494 | bullish consensus |
| Exclude any single HARD-break series (52 variants) | 0.8494 | bullish (identity) |
| Drop SPEC sub-index | 0.8494 | bullish |
| Drop HEDGE sub-index | 0.7470 | neutral (class change) |
| Drop DEALER sub-index | 0.8735 | bullish |
| Drop one SPEC component (8 variants) | 0.8373 to 0.8614 | bullish, all |
| Drop one HEDGE component (7 variants) | 0.8193 to 0.8735 | bullish, all |
| Drop one DEALER component, SPY proxy (3 variants) | 0.8012 to 0.8494 | bullish, all |

(Source: results/m6.json, series.breakrob_table)

Across the 21 leave-one-out variants, $P$ stays between 0.747 and 0.8735, and exactly one variant leaves the bullish class: dropping the entire HEDGE sub-index lands at $P = 0.747$, which the rule reads as neutral. That result says two things. First, the bullish class is carried by cheap protection: strip out the HEDGE dial and what remains, the crowd's buying plus dealer positioning, still ranks around the 75th percentile of the

year, close below the bullish gate but no longer above it. Second, nothing comes anywhere near bearish: no variant drops $P$ below 0.747 against a bearish gate of 0.20. The direction of the read is robust; the class boundary depends on one sub-index. No single-component drop changes the class, and the closest calls are HEDGE components tied to the level of implied volatility (dropping H1 or H6 lands at 0.8193, still above the gate). HEDGE is the single series to watch, the role Chapter 15 assigns it, and its load-bearing position here feeds the confidence discussion in Chapter 14.

## 12.5 Power and single-vendor exposure

Both limits bind harder than the tables suggest. On power: 165 sessions rank days against each other with adequate sample, but the forward-return grid thins to 102 entry dates at 63 days with cells as small as $n = 17$ (Q5), the windows overlap heavily, and one episode does most of the work. The registered statement stands: ranking claims are well powered, forecasting claims are underpowered. On exposure: every number in Chapters 1 through 12 comes from a single data vendor, and the two vendor endpoints that could have cross-checked the tape sat behind a higher subscription tier during the run. Chapter 13 is the compensating check: it rebuilds the study's insurance read from public data back to 1990 and repeats the quintile logic there. What it cannot do is verify this vendor's options tape against another vendor's. The confidence grade attached to the verdict in Chapter 14 prices both limits in.

Scope and data flags for this chapter, in one place. (1) Forward returns use the vendor's regular-session SPY closes, which reach back exactly one year (2025-08-28 to 2026-08-28, amendment A1); quintiles are cut on each day's composite 2026 rank at 0.2/0.4/0.6/0.8, 33 sessions per bucket. (2) Two softer level-shifts sit inside the 2026 ranking window with no adjustment applied: QQQ put premium steps up about 2.1x on 2026-01-28 (it reaches the composite through one of eight SPEC components at QQQ's 0.30 market weight, roughly 3.75% of SPEC) and QQQ call delta steps up about 2.0x on 2026-04-09 (it feeds D1). The refreshed break audit adds no new flags on any series inside the market blend. (3) H3, the gap between what options charge for future volatility and the volatility that actually shows up, changes definition mid-stream: the official series ends 2026-07-31, the extension from 2026-08-03 measures realized volatility backward rather than forward, and the five sessions behind the current verdict all sit on the extension side of that seam. (4) Module checks pass: quintile sizes reproduce, the resampled error bars are byte-identical under reruns, and the sub-index blends reconstruct from stored component scores to within 0.00625 (SPEC) and 0.00007 (HEDGE).

# 13. Long-History Context: A Non-UW Cross-Check

Everything before this point measures 2026 against itself. This chapter checks the year against 36 years of market history, using the one piece of positioning with public data that deep: the price of insurance. The check matters for the verdict. A reading that is hot inside 2026 can still be ordinary by history, and that distinction decides how much weight the bullish call deserves.

NON-UW CHAPTER. Everything here comes from public, non-Unusual-Whales sources: FRED (VIXCLS, VXVCLS), Cboe (the SKEW index), and yfinance (SPY closes). No UW series enters this chapter and nothing here feeds Chapters 1 to 12. Per amendment A1 this is the study's designated home for long-history replication; per amendment A3 the Cboe equity put/call archive was inaccessible (HTTP 403), which limits the proxy set to the VIX level, the VIX3M−VIX term slope, and SKEW, plus SPY closes for forward returns. "Today" throughout this chapter is 2026-08-27: the FRED and Cboe series lag the study's 2026-08-28 as-of date by one session, and the last date carrying the complete [VIX, VIX3M−VIX, SKEW] triple is used.

## 13.1 What a long history can and cannot check

The composite of Chapters 1 to 12, the study's Positioning Consensus Index (PCI), exists only on the 165-session 2026 window. Chapter 12 showed what that costs: the year's quintile result is built on one drawdown-and-rebound episode, and no test inside a single year can tell a real contrarian signal from a year that happened to mean-revert. This chapter trades precision for history. It builds a simple stand-in for the composite from three public series with decades of data: the VIX (the market's benchmark price of one-month S&P 500 insurance), the gap between the three-month VIX and the one-month VIX (whether longer-dated protection costs more than short-dated, the calm configuration, or less, the stressed one), and the Cboe SKEW index (how much extra investors pay for protection against a deep crash). Each series is ranked against its own full history and the three ranks are averaged, using the sign convention the study applies everywhere: expensive insurance counts as fear, cheap insurance as complacency. The slope series begins 2007-12-04; before that the proxy averages the two remaining components.

Two limits bound every claim that follows. The proxy sees only the insurance axis. It is a stand-in for HEDGE (what protection costs); neither SPEC (what the crowd is actually buying) nor DEALER (what market makers are forced to do) has any public history reaching 1990. And the proxy is not the PCI: it can contextualize the cheap-insurance leg of the verdict, not the composite as a whole.

## 13.2 Today against 36 years of insurance pricing

The table ranks each reading against 2026 alone and against its full history. Both columns are percentiles of daily readings, where 100 means the highest on record in that window.

| Series (2026-08-27) | Level | 2026 pctile | Long-run pctile | History |
|---|---|---|---|---|
| VIX close | 14.51 | 2.4 | 28.5 | 1990+ |
| VIX3M−VIX slope | 3.05 | 68.5 | 82.0 | 2007+ |
| Cboe SKEW | 144.05 | 59.6 | 93.0 | 1990+ |
| Proxy composite | +7.10 | 84.8 | 55.4 | 1990+ |

(Source: results/m6.json, headline_numbers ch12_vix_today, ch12_slope_today, ch12_skew_today, ch12_proxy_today)

The proxy composite prints +7.1 on its −100 (maximum fear) to +100 (maximum complacency) scale. That is the 85th percentile of 2026, meaning more complacent than on roughly 8 of every 10 trading days this year, but only the 55th percentile since 1990: just above the middle of 36 years, nowhere near an extreme. That gap between the two frames is this chapter's headline. The Chapter 14 verdict rests on the same shape, a composite at the 85th percentile of its year, and the fuller-history figure quoted there (61st percentile since 2024) is an average of component ranks, not a rank of the composite itself, because the composite has no path before 2026. The proxy here is the study's only true multi-decade rank: one series measured against its own 36-year past. It reads today the same way the verdict's qualifier does: hot inside 2026, ordinary in the long view.

The composition carries a second lesson. The VIX at 14.51 is one of the calmest closes of 2026 (2nd percentile of the year) yet unremarkable across decades: everyday one-month insurance has been cheaper than this on roughly three of every ten days since 1990. The slope is genuinely steep by history (82nd percentile since 2007): three-month protection costs 3.05 volatility points more than one-month, the calm-market shape, but a rich version of it. And SKEW at 144.05 sits at the 93rd percentile since 1990: deep-crash protection is more expensive than on nine of every ten days in the sample. Inside the proxy, the cheap VIX and the steep slope push the score toward complacency while the rich SKEW pulls it back, and the blend lands just above the middle. The same split runs through the study's own protection-cost sub-index (Chapters 6 and 14): what is cheap is the level of everyday volatility; demand for downside protection is not absent.

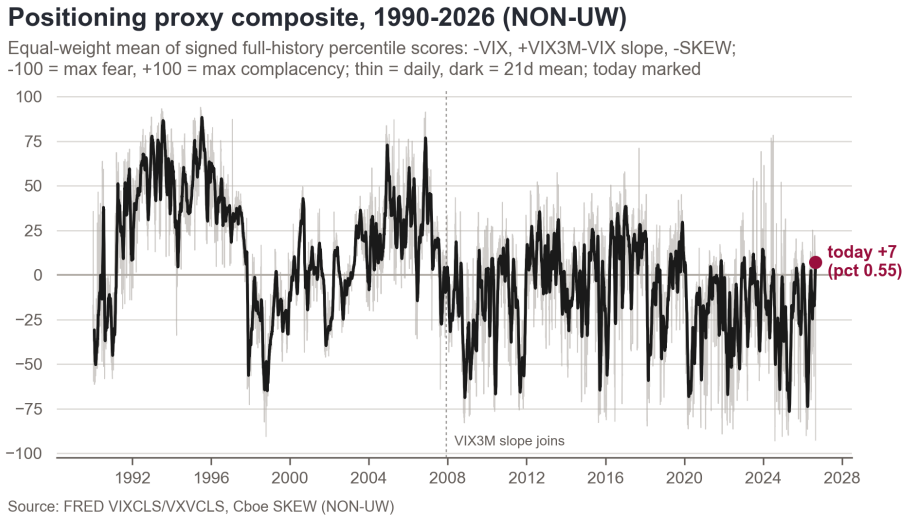

## 13.3 The closest analogues since 2007

To ask what followed similar insurance pricing, every session since 2007-12-04 (when the slope series begins) was scored on how close its [VIX, slope, SKEW] triple sits to today's, with each series standardized so the three scales are comparable (z-scores; closeness is Euclidean distance, the $z$-dist column). A date qualifies only if it already has both 21-day and 63-day forward SPY returns, which mechanically removes the most recent 63 sessions. The ten nearest dates:

| Date | VIX | Slope | SKEW | $z$-dist | Fwd 21d % | Fwd 63d % |
|---|---|---|---|---|---|---|
| 2023-06-05 | 14.73 | 3.17 | 143.18 | 0.0937 | +3.32 | +5.57 |
| 2024-08-16 | 14.80 | 3.01 | 142.55 | 0.1275 | +1.58 | +7.37 |
| 2026-01-07 | 15.38 | 3.21 | 143.57 | 0.1301 | +0.15 | −1.13 |
| 2023-06-16 | 13.54 | 3.21 | 143.07 | 0.1553 | +3.58 | +1.30 |
| 2025-07-10 | 15.78 | 3.14 | 143.32 | 0.1616 | +1.82 | +7.85 |
| 2026-05-29 | 15.32 | 3.34 | 144.18 | 0.1647 | −1.03 | +1.96 |
| 2023-06-02 | 14.60 | 3.32 | 142.61 | 0.1727 | +3.94 | +5.83 |
| 2026-01-23 | 16.09 | 3.06 | 144.02 | 0.1796 | −0.27 | +3.87 |
| 2023-06-06 | 13.96 | 3.08 | 146.27 | 0.1912 | +2.84 | +4.64 |
| 2023-07-21 | 13.60 | 2.80 | 142.55 | 0.1982 | −2.84 | −5.36 |
| Mean | | | | | +1.31 | +3.19 |

(Source: results/m6.json, series.ch12_analogue_table; forward SPY total-close returns, NON-UW)

Similar configurations have mostly resolved higher, but not reliably. Eight of the ten dates saw SPY higher 63 trading days later (the hit rate, the share of dates with a positive forward return), with a mean of +3.2% at 63 days and +1.3% at 21 days. The dispersion is wide: from +7.9% (July 2025, when that rally simply kept going) to −5.4% (July 2023, the top of that summer's rally and the start of a three-month correction). Three of the ten dates come from 2026 itself and replay the year in miniature: the early-January date sat just ahead of this year's drawdown (−1.1% over the next 63 days), while late January (+3.9%) and late May (+2.0%) resolved higher.

Two structural caveats. The ten dates collapse into six distinct episodes, with June 2023 alone contributing four sessions, so the effective sample is six, not ten. And the match is on insurance pricing only: nothing constrains what the crowd was buying, or how dealers were positioned, to have resembled today on those dates.

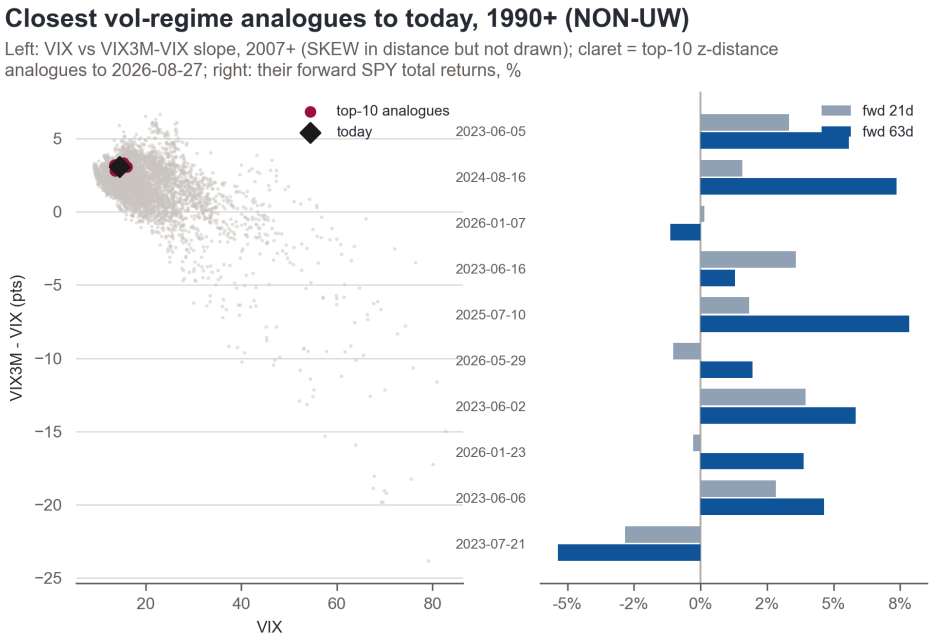


## 13.4 The same test on 26 years of data, 2000 to 2026

Chapter 12 asked whether days when the composite read bullish were followed by weaker returns, and one year of data could not settle it. The same design runs here on 6,704 sessions from 2000-01-03 (the start of the SPY close history). Each day's proxy score is sorted into five equal buckets (Q1 = the most fearful fifth, Q5 = the most complacent fifth, which is the most bullish under the study's convention), and forward SPY returns are measured 5, 10, 21, and 63 sessions ahead. Error bars are built by resampling whole stretches of days rather than single days, which respects the fact that markets trend (a block bootstrap: blocks of 10 sessions, 5,000 draws, seeds fixed deterministically from 20260820). Significance is judged after correcting for the fact that the grid runs 20 tests at once (BH-FDR at $q = 0.10$). For each bucket the table shows the average and median forward return, the hit rate, a standard $t$-statistic, the resampled 95% interval, the bootstrap $p$-value, and whether the cell survives the correction.

| Horizon | Q | $n$ | Mean % | Median % | Hit | $t$ | 95% CI | $p_{boot}$ | FDR |
|---|---|---|---|---|---|---|---|---|---|
| 5d | Q1 | 1339 | 0.43 | 0.75 | 61.0% | 4.33 | [0.04, 0.79] | 0.0308 | yes |
| | Q2 | 1341 | 0.12 | 0.43 | 57.4% | 1.76 | [−0.09, 0.31] | 0.2767 | |
| | Q3 | 1338 | 0.13 | 0.32 | 57.3% | 2.01 | [−0.08, 0.33] | 0.2076 | |
| | Q4 | 1340 | 0.22 | 0.33 | 59.4% | 4.43 | [0.06, 0.39] | 0.0100 | yes |
| | Q5 | 1341 | 0.06 | 0.17 | 55.4% | 1.48 | [−0.08, 0.20] | 0.3787 | |
| 10d | Q1 | 1339 | 0.79 | 1.48 | 64.5% | 5.95 | [0.16, 1.38] | 0.0136 | yes |
| | Q2 | 1341 | 0.30 | 0.76 | 60.0% | 3.35 | [−0.05, 0.62] | 0.0984 | |
| | Q3 | 1337 | 0.27 | 0.61 | 59.2% | 3.26 | [−0.07, 0.62] | 0.1136 | |
| | Q4 | 1336 | 0.30 | 0.59 | 61.5% | 4.10 | [0.01, 0.59] | 0.0408 | yes |
| | Q5 | 1341 | 0.23 | 0.48 | 59.7% | 4.04 | [−0.02, 0.48] | 0.0764 | |
| 21d | Q1 | 1339 | 1.60 | 2.60 | 68.4% | 9.28 | [0.70, 2.53] | 0.0016 | yes |
| | Q2 | 1341 | 0.65 | 1.65 | 65.8% | 4.75 | [0.01, 1.25] | 0.0472 | yes |
| | Q3 | 1335 | 0.48 | 1.09 | 61.4% | 3.85 | [−0.09, 1.05] | 0.1052 | |
| | Q4 | 1331 | 0.68 | 1.31 | 64.5% | 6.38 | [0.15, 1.16] | 0.0140 | yes |
| | Q5 | 1337 | 0.55 | 0.95 | 62.9% | 6.63 | [0.13, 0.94] | 0.0092 | yes |
| 63d | Q1 | 1322 | 5.07 | 6.00 | 77.1% | 21.14 | [3.72, 6.45] | 0.0004 | yes |
| | Q2 | 1330 | 2.02 | 3.55 | 68.7% | 9.74 | [0.81, 3.15] | 0.0012 | yes |
| | Q3 | 1326 | 1.30 | 3.01 | 66.2% | 6.06 | [0.10, 2.41] | 0.0368 | yes |
| | Q4 | 1326 | 1.40 | 3.21 | 65.2% | 6.49 | [0.18, 2.55] | 0.0220 | yes |
| | Q5 | 1337 | 1.95 | 2.71 | 68.1% | 13.21 | [1.16, 2.68] | 0.0004 | yes |

(Source: results/m6.json, series.ch12_fwd_table_nonuw; mean forward SPY return per cell, NON-UW closes, 2000-01-03 onward)

Thirteen of the 20 cells stay significant after the correction (nine at the stricter $q = 0.05$, seventeen at the looser $q = 0.20$), and the complacent-minus-fearful spread (Q5−Q1) is negative at every horizon. Restricting to 2007 onward, where all three components exist, keeps the sign and roughly the magnitude:

| Q5−Q1 spread (mean fwd SPY return, %) | 5d | 10d | 21d | 63d |
|---|---|---|---|---|
| Full window, 2000+ | −0.37 | −0.56 | −1.05 | −3.12 |
| 3-component subsample, 2007+ | −0.23 | −0.46 | −0.97 | −3.36 |

(Source: results/m6.json, series.ch12_fwd_spreads_nonuw and series.ch12_fwd_spreads_nonuw_2007plus)

This is the sign confirmation the 2026 window could not supply. One year contains one drawdown-and-rebound; 26 years contain many independent episodes, so a negative spread here cannot be the single-episode artifact the Chapter 12 spread might be. Three qualifications keep the result in proportion. First, the information lives at the fear extreme, not the complacency extreme. Q1 is the strongest cell at every horizon: after the most fearful fifth of days, SPY averaged +5.1% over the next quarter, with gains 77% of the time. Q5, the complacent fifth, still averaged +2.0%, more than the middle buckets. Expensive insurance reliably preceded strong returns; cheap insurance did not reliably precede losses, it preceded smaller gains. Second, the momentum benchmark that explained most of the year's spread in Chapter 12 was not replicated on this window, so whether fading insurance extremes adds anything beyond fading

recent price weakness is untested here. Third, entry dates overlap heavily within each horizon, so cell counts near 1,340 overstate the independent evidence, and the resampled error bars only partially correct for that.

The long history contributes one genuine result and one piece of context. The result: on a 26-year, multi-episode sample, days of expensive insurance were followed by stronger returns and days of cheap insurance by weaker ones (Q5−Q1 negative at all four horizons, 13 of 20 cells significant after the multiple-test correction, sign stable in the 2007+ subsample). That is the confirmation the single-episode 2026 window cannot establish, and its information sits almost entirely at the fear extreme, not the complacency extreme. The context: today's proxy reads the 85th percentile of 2026 but only the 55th since 1990, in the flat middle of the grid rather than at either informative extreme, and its ten nearest analogues mostly resolved higher (8 of 10 positive at 63 days, mean +3.2%). Set against the Chapter 14 verdict (bullish consensus, at the overly-bullish boundary), the long view adds the discipline: the bullish reading is a statement about 2026's own range; by 36 years of insurance pricing today is ordinary, and even genuine complacency extremes have historically meant smaller gains rather than reliable losses. The proxy sees only the insurance axis, so this context bounds the cheap-insurance leg of the verdict, not the crowd or dealer legs.

# 14. The Verdict

This chapter turns thirteen chapters of measurement into one call. It runs the classification rule fixed before any data were scored (Research Plan section 8) on the August 28 close, shows what drives the result and what followed similar readings earlier in 2026, and prices the call against the second adversarial review.

Bullish consensus, at the boundary of the study's own pre-registered overly-bullish trigger, printed at medium confidence. For the first time this year all three dials point the same way on the week's average: the crowd is buying calls, dealer positioning has flipped supportive, and insurance is cheap and unbid. Two qualifiers travel with the call. First, it is hot within 2026's own range (85th percentile of the year) but ordinary against the fuller history (61st percentile since 2024). Second, it is one week old: speculative demand jumped from the 21st to the 77th percentile of the year in the final six sessions of August. The year as a whole still shows no statistical trend; what changed is the last week, not the whole tape. The boundary, spelled out: the trigger fires on the cost-of-protection leg alone, by 1.9 percentile points (about three trading days' rank); removing any one of three of its seven components un-fires it; cheap option prices carry it while the two components tracking actual protection buying lean the other way; and the other leg, amplification risk, is dead at the 33rd percentile, down from the 73rd six sessions earlier.

## 14.1 The rule and its inputs

The verdict runs on the Positioning Consensus Index (PCI), the study's single score for how bullish positioning is, blending three dials scored $-100$ to +100 within 2026: SPEC, what the crowd is actually buying; DEALER, what market makers' hedging is forced to do; HEDGE, what protection costs. $P$ is where the PCI's latest 5-day average ranks against every trading day of 2026 so far: at or above 0.80 the class is bullish consensus, at or below 0.20 bearish, otherwise neutral. A bullish class escalates to overly bullish only if a second dial confirms: $F$, amplification risk, at or above 0.70, or the HEDGE rank at or above 0.80. Amplification risk means this: when market makers are net short gamma, their hedging sells as the market falls and buys as it rises, amplifying moves in both directions; $F$ ranks that exposure within 2026. (The mirror arm reads overly bearish at $P \leq 0.20$ with HEDGE at or below 0.20.) Confidence is high if all three dials agree in sign and the quality checks are clean, medium if one dissents, low if two dissent or a load-bearing series is broken.

| Input | Value | 2026 rank | Multi-year rank | Window |
|---|---|---|---|---|
| $P$ (PCI, 5-day average) | +21.04 | 0.8494 | 0.6126 | rank of 5d mean in 2026 dailies |
| PCI, latest day | +1.61 | 0.5301 | 0.5227 | 2026-08-28 |
| $F$ (amplification risk) | 0.3301 | 0.3301 | 0.2947 | percentile of $-$net gamma |
| HEDGE, 5-day average | +31.35 | 0.8193 | 0.8494 | M3, SPY-based market series |
| Trend test (Mann-Kendall $z$) | $-0.681$ | | | $p = 0.4961$, $n = 165$ |
| Trend slope (Sen) | $-0.0219$ | | | score points per session |

(Source: results/m4.json, m2.json, m3.json headline_numbers. The multi-year rank of the blended composite is a convention (weighted average of component ranks); the blend has no pre-2026 path to rank against.)

$P = 0.8494$ clears the 0.80 gate: the week ranks at the 85th percentile of the year, stronger than roughly 8 of every 10 trading days of 2026. The escalation arm fires through HEDGE alone (0.8193 against 0.80) with the amplification-risk leg dead at 0.3301; per the round-2 review, this report always prints the firing leg, its margin, and the other leg's value.

## 14.2 The class is real, and it is young

The class itself is the most robust result of the refreshed run: all three admissible readings of the rule sentence fire (0.8265 to 0.9815; on the strictest, this is the third-highest smoothed reading of the year), dropping any single day from the window keeps $P$ at 0.8253 or higher (four of the five days printed +23.0 to +31.5, the 88th to 95th percentile: a week-long state, not one print), and 20 of 21 leave-one-out variants keep the class ($P$ 0.747 to 0.8735; collapsing the crowd dial's three near-duplicate components leaves 0.8373). (Source: results/m6.json breakrob_*; results/ADVERSARIAL_LOG2.md findings 1, 5, 6)

Three qualifiers cut the other way. First, this is a 2026 statement: against the study's longer component history the same week sits at the 61st percentile, where the rule would read neutral, and that multi-year figure is itself a convention (table note above). Second, it is three sessions old: the class first fired on August 26, $P$ moved from 0.3735 to 0.8494 in the six sessions after August 20, and the August 28 close itself printed +1.6, the 53rd percentile, with two of the three dials negative on the day. The verdict describes the August 21 to 27 tape. Third, the study's slower clocks disagree on purpose: the regime table labels July 28 to August 28 "neutral, flat" (average PCI +7.7), and a standard statistical trend test (Mann-Kendall) finds no 2026 trend ($p = 0.50$). No contradiction, different windows: a one-week surge inside a flat month inside a trendless year. (Source: results/m4.json headline_numbers; results/ADVERSARIAL_LOG2.md finding 1)

## 14.3 Why "overly" is only a boundary call

The round-2 review killed the bare qualifier. The HEDGE leg fires, but not by enough to stand alone:

- **Margin.** 1.9 percentile points above the gate: 3.2 rank positions out of 165 days, or 1.68 score points (5-day average +31.35 versus the 80th-percentile daily value +29.68).
- **Component fragility.** Removing any one of three of the seven HEDGE components un-fires it (0.7952,

0.7349, 0.7530); the other four removals keep it (0.8313 to 0.8614).

- **Cheap insurance, still-active buying.** The three components measuring how cheap options are carry the label (+74.0, +92.5, +88.7; the middle one is the VIX's own option prices in the cheapest 5% of their one-year range, itself in the bottom 1.2% of 2026 readings). The two components tracking what protection buyers actually do dissent: the price of upside calls versus equally out-of-the-money downside puts shows relative put demand near its strongest of the year ($-55.2$), and the short-volatility ratio is elevated ($-47.7$). The disagreement is larger than the margin.
- **Short half-life.** The leg holds next session only if the day prints a HEDGE score of +8.8 or better (41.8% of 2026 days did); the August 28 print by itself ranks 0.7711, below the gate.

(Source: results/ADVERSARIAL_LOG2.md finding 2; results/m3.json H6_VIX_iv_rank_latest)

The dead leg moved too. Negative net gamma, the state in which dealer hedging amplifies moves rather than dampens them, remains the year's norm (116 of 165 sessions), with spot 0.06% below the SPY flip level of 769.62, the price above which hedging would switch to dampening. Ranked within 2026, though, the dial reads calm, and it is unstable: 21st percentile to 73rd to 33rd within the last ten sessions. That instability is the design flaw the review flagged: in two consecutive runs the either-or trigger fired on whichever leg sat just above its gate while the other was dead (August 20: amplification risk 0.7194 against 0.70 with HEDGE at 0.90; August 28: HEDGE 0.8193 against 0.80 with amplification risk at 0.33), and seven 2026 days sit inside the 0.78 to 0.82 HEDGE band alone. The label carries no margin information, so this report prints the anatomy, and Chapter 15 restates the triggers the same way. (Source: results/m2.json headline_numbers; results/ADVERSARIAL_LOG2.md findings 2 and 8)

## 14.4 Attribution: what makes the read

The waterfall decomposes the composite exactly (contributions sum to the day's PCI; stored check difference 0.0):

| Component | Latest contribution | Contribution pct (2026) | August monthly mean |
|---|---|---|---|
| SPEC ($0.45\times$ score) | $-1.381$ | 0.494 | +1.256 |
| HEDGE ($0.35\times$ score) | +10.195 | 0.771 | +10.672 |
| DEALER ($0.20\times$ score) | $-7.201$ | 0.175 | $-4.905$ |
| PCI | +1.61 | 0.530 | +7.024 |

(Source: results/m4.json, headline_numbers waterfall_* and series.waterfall_monthly)

The day and the week disagree, and the class lives in the week. On August 28 itself the crowd dial is flat ($-1.4$, its 49th percentile of the year), cost of protection carries the read (+10.2, 77th), and dealer positioning drags ($-7.2$, 17th). The 5-day averages behind the verdict are all positive: SPEC +21.3, HEDGE +31.4, DEALER +2.5. (Source: results/m1.json, m3.json, m2.json headline_numbers)

## 14.5 What followed readings like this in 2026

$P = 0.8494$ sits in the top fifth of the year's readings (quintile cuts at 0.2/0.4/0.6/0.8). Forward SPY returns after the 2026 days in that band:

| Horizon | $n$ | Mean % | Hit | 95% CI | FDR |
|---|---|---|---|---|---|
| 5d | 29 | 0.08 | 48.3% | [−0.45, 0.13] | |
| 10d | 28 | −0.40 | 28.6% | [−1.18, −0.19] | yes |
| 21d | 26 | 0.32 | 42.3% | [−1.03, 0.76] | |
| 63d | 17 | 3.82 | 76.5% | [1.83, 7.71] | yes |

(Source: results/m6.json, series.fwd_table_uw; forward SPY returns, UW regular-session closes. Hit = share of entries with a positive forward return. CI = error bars built by resampling whole stretches of days, which respects that markets trend. FDR = significant after correcting for running many tests at once.)

The grid is contrarian throughout: the most bullish fifth underperforms the most bearish fifth at every horizon (10-day spread −2.1 points, 21-day −2.4, 63-day −4.8; 10 of 20 cells survive the correction). Chapter 12 shows the pattern is dominated by one drawdown-and-rebound episode in the first half of the year, and the 63-day column has no entries after late May (the forward window runs off the sample), which is why its means are large and positive. Read the table as a description of 2026, not a forecast.

The rule's own history is the second base rate. Its default output is neutral: only 14 of 165 days in 2026 classify bullish, in five clusters (the first week of January, January 27 to 28, June 1, August 3 to 5, August 26 to 28). The late-January firing came immediately before a 9.1% SPY drawdown that ran to March 30; the early-August firing preceded the mid-August fade ($P$ back to 0.37 by August 20). Across prior bullish-class days, average SPY returns were −0.13% over the next 10 sessions (positive 18% of the time) and −0.56% over 21 (positive 38%): eight to eleven overlapping windows, descriptive only, and round 1 killed the causal version of this pattern in both directions. The narrow point stands: in 2026 this state has marked local tops, not launches. (Source: results/ADVERSARIAL_LOG2.md finding 1)

## 14.6 Why the printed confidence is medium

The contracted grade tests the sign of each dial's 5-day average. On that statistic all three are positive (+21.3, +31.4, +2.5), zero dissent, and the mechanical output is high. The round-2 review killed the bare grade; this report prints medium, for four reasons:

1. **Zero dissent rests on DEALER at +2.5** on a −100 to +100 scale: 0.07 of that series' typical daily swing (standard deviation 36.7), the 50.6th percentile of the year, indistinguishable from zero. The average stays positive next session only if the next daily beats +13.2 (38.8% of 2026 days did). One ordinary dealer session flips the grade.
2. **The equally valid daily-sign reading is medium, with a negative majority.** On the August 28 prints, SPEC −3.1, HEDGE +29.1, DEALER −36.0. The newest DEALER daily is the most negative of the scored window, with NOPE (a measure of how much option trading is pushing around the stock itself) at

its 12.7th percentile of the year. Replaying each of the last ten daily prints as the next session, nine of ten scenarios lose at least one element of "bullish, overly, high confidence": eight drop the grade, three lose the class.

3. **The "quality checks clean" condition was never evaluated**; as coded, the grade tests signs only. Inside the scored window: the official volatility-risk-premium feed (the gap between what options charge for future volatility and the volatility that shows up) ends July 31, so all five scored values of that component come from the study's own extension; and every in-sample number in Chapters 1 to 12 is single-vendor, with round 1's two soft data breaks still inside the unrestricted 2026 window. (Source: results/ADVERSARIAL_LOG2.md findings 3 and 7)
4. **The class leans on one dial.** Of the 21 leave-one-out variants, only removing the entire HEDGE sub-index changes the class, to neutral at $P = 0.747$ (the remaining two dials still rank the week at the 75th percentile). Milder than the August 20 run, when removing HEDGE flipped the read to bearish, but the call still stands or softens with a single dial. (Source: results/m6.json, breakrob_loo_flips, breakrob_P_min/max)

Implementation disclosures on the confidence rule. (1) As coded in M4, the "low if two dissent" branch cannot fire: dissent is counted against the majority sign of three non-zero sub-index signs, yielding zero or one dissenter by construction, and the "quality checks clean" and "broken load-bearing series" conditions are not implemented. That is why the stored module string says high while this report prints medium. (2) Multi-year percentiles mix conventions across modules: blended composites use the weighted component-average convention, the HEDGE headline uses its own rank-based convention, and the DEALER multi-year rank blends one component with full history against two that exist only in 2026. Cross-module multi-year comparisons are approximate.

## 14.7 What this read is, and is not

This is a ranking statement, well powered at 165 sessions: where positioning and option prices sit within their 2026 distributions. It is not a return forecast; Chapter 12 shows the in-sample forward-return pattern is one mean-reversion episode. Use the verdict as worded: a one-week-old bullish consensus at the boundary of the overly-bullish trigger, carried by cheap insurance while insurance buying itself continues, printed at medium confidence. The dial to watch is HEDGE, the only component whose removal changes the class; the negative net gamma backdrop that would amplify a shock remains the year's norm even while the amplification-risk dial reads calm; Chapter 15 lists the pre-registered triggers that would change the read. This is research, not investment advice; nothing here is a recommendation to buy, sell, or hold any security, and the authors are not acting as investment advisers.

# 15. What Would Change the Read

The Chapter 14 verdict comes from thresholds written down before any data ran, so new data can overturn it mechanically, with no re-interpretation. This chapter lists those tripwires, names the single series most likely to trip them, and closes with the monitoring checklist. The exits matter: the read is three sessions old, and its strongest qualifier clears its gate by about two percentile points.

## 15.1 Falsification triggers

The rule (RESEARCH_PLAN Section 8) takes three inputs, recomputed daily. $P$ is where the 5-day average of the composite ranks against every trading day of 2026 so far: 0.8494 on August 28, the 85th percentile, higher than about 8 of every 10 days this year. $F$ is amplification risk. Net gamma measures whether market makers' hedging dampens market moves (positive) or amplifies them (negative); when it is negative, hedging sells into falls and buys into rallies, and $F$ ranks how strong that amplifying setup is against the year: 0.3301, down from 0.73 six sessions earlier. The third input is the 2026 rank of HEDGE, the sub-index tracking what protection costs: 0.8193. (The other two: SPEC tracks what the crowd is actually buying, DEALER what market makers are forced to do.)

| Trigger | Contract rule (Sec. 8) | Current input (Aug 28) | Status |
|---|---|---|---|
| Bullish consensus | $P \geq 0.80$ | $P = 0.8494$ | met since Aug 26 |
| . . . overly bullish | also $F \geq 0.70$ or HEDGE pct $\geq 0.80$ | $F = 0.3301$; HEDGE $= 0.8193$ | HEDGE leg only, by 0.019 |
| Bearish consensus | $P \leq 0.20$ | $P = 0.8494$ | far side |
| . . . overly bearish | also HEDGE pct $\leq 0.20$ | HEDGE $= 0.8193$ | far side |
| Trend (trajectory) | Mann-Kendall trend test $p < 0.05$, two-sided | $p = 0.4961$; Sen slope $-0.0219$/session | not met |

(Source: m4.json: verdict_P, mann_kendall_p, sen_slope_per_session; m2.json: fragility_F_latest; m3.json: HEDGE_latest_5d; first-firing date from ADVERSARIAL_LOG2.md)

The "overly bullish" row needs its margins printed with it. The qualifier is an either-or across two noisy dials, and it has now fired on each in turn while the other lay dead: on data through August 20, amplification risk sat at 0.7194, just 0.019 over its 0.70 gate, with the protection-cost dial at 0.90; eight sessions later the legs have swapped, with HEDGE 0.019 over its 0.80 gate (3.2 rank positions out of 165 days, or 1.7 points on the $-100$ to +100 score scale) and $F$ at 0.33. The dials are fast: $F$ went 0.21 to 0.73 and back to 0.33 within ten sessions, and seven 2026 days sit inside the 0.78 to 0.82 HEDGE band. This report therefore prints the firing leg, its margin, and the other leg's value wherever the label appears, and phrases the verdict as "at the overly-bullish boundary", not a bare "overly bullish". (Source: ADVERSARIAL_LOG2.md findings 2 and 8; m2.json: fragility_F_latest; m3.json: HEDGE_latest_5d)

The bullish class itself is one week old. $P$ first crossed 0.80 on August 26 after rising from 0.3735 on

August 20, and no single day's removal drops it below 0.8253. But the as-of day printed +1.61, the 53rd percentile, and the 5-day average follows the dailies with a lag. Against the study's longer multi-year component history the same average ranks at the 61st percentile, where the rule would read neutral (a component-average approximation; the blend has no pre-2026 path to rank against). Base rates advise restraint: 14 of 165 sessions this year classify bullish, the late-January cluster immediately preceded the year's deepest drawdown, the August 3 to 5 cluster preceded the mid-August fade, and SPY averaged $-0.1\%$ over the ten days after a bullish-class day with gains 18% of the time (descriptive, eleven overlapping windows). The study's own regime table still labels July 28 through August 28 "neutral, flat" (mean +7.72), and the trend test finds no 2026 trend ($p = 0.4961$). The read is one hot week, and one ordinary week un-makes it. (Source: m4.json: verdict_P, pci_latest_day, regime_6_mean_pci, mann_kendall_p; ADVERSARIAL_LOG2.md finding 1)

The confidence stamp is thinner still. The contract grants "high" when all three sub-indices agree in sign and QC is clean; the sign test runs on 5-day averages, where SPEC +21.3, HEDGE +31.4, and DEALER +2.5 give zero dissent. The DEALER agreement, though, is +2.5 on a $-100$ to +100 scale: 0.07 of that series' typical daily swing, the 51st percentile of the year, statistically zero. On the as-of day's daily prints the majority sign is negative (SPEC $-3.1$, HEDGE +29.1, DEALER $-36.0$) and the same rule says medium. The 5-day DEALER average stays positive tomorrow only if the session beats +13.2, done by 39% of 2026 sessions, and the QC-clean condition is not evaluated in code. The honest statement: high by the mechanical sign rule on a +2.5 DEALER margin; medium on the daily-sign reading; one ordinary dealer session flips it. (Source: m1.json: spec_latest_5d; m3.json: HEDGE_latest_5d; m2.json: dealer_dir_latest5d; m4.json: waterfall_spec/hedge/dealer_contrib_latest; ADVERSARIAL_LOG2.md finding 3)

## 15.2 The single series to watch: HEDGE

Chapter 12's robustness grid rebuilds the composite 21 times, each run leaving out one ingredient. $P$ stays between 0.747 and 0.8735, and exactly one variant changes the class: removing the entire HEDGE sub-index drops $P$ to 0.747 and the read to neutral, not bearish. Even ignoring protection costs entirely, crowd buying plus dealer positioning ranks near the 75th percentile of 2026. No single-component removal changes the class, and the break-exclusion check is an identity (0 of 52 flagged series enter the market composite). (Source: m6.json: breakrob_table, breakrob_loo_flips, breakrob_P_min, breakrob_P_max, breakrob_hard_in_pci)

HEDGE also carries the overly-bullish qualifier alone, held up by the level of option prices rather than an absence of demand for protection. The three price-level components sit at or near 2026 floors: SPY's IV rank (H1), where option prices sit in their one-year range (0 = cheapest, 100 = most expensive), scores +74.0 on the 5-day average; the VIX IV rank (H6), the same measure for options on the VIX itself, is 5.4491 of 100, its 1st percentile of 2026, scoring +92.5; and the 30-day implied-move percentile (H7) scores +88.7. Both demand components dissent: the 25-delta risk reversal (H2), the price of upside calls versus equally out-of-the-money downside puts, is +0.02605, its 18th percentile of the year, scoring $-55.2$: put protection is relatively well bid; and the SPY short-volume ratio (H5) scores $-47.7$. The rest are mid-pack (volatility risk premium +34.2, term slope +33.0). Removing any one price-level component un-fires the qualifier (HEDGE rank 0.7952 without H1, 0.7349 without H6, 0.7530 without H7); removing any of the

other four leaves it firing (0.8313 to 0.8614). Holding the gate tomorrow needs a daily HEDGE print of at least +8.8 (42% of 2026 sessions); August 28's own print ranks 0.7711, below the gate. The leg is not frame-fragile (all three percentile conventions fire; multi-year component-average rank 0.8494); the exposure is margin plus component mix. (Source: m3.json: HEDGE_latest_5d, H2_RR_latest, H6_VIX_iv_rank_latest, series score_H1..H7_SPY_frameA; m4.json: waterfall_hedge_contrib_latest; ADVERSARIAL_LOG2.md finding 2)

HEDGE is the hinge: the only live leg of the overly-bullish qualifier (0.8193 against a 0.80 gate, 3.2 rank positions of margin), the only sub-index whose removal changes the verdict class ($P$ 0.8494 to 0.747, bullish to neutral), and split inside, price levels at 2026 floors against dissenting demand components. Three of its seven component removals un-fire the qualifier; the disagreement inside the index is larger than the margin over the gate. Watch the VIX IV rank first; at the year's floor, its only room is up.

## 15.3 Monitoring checklist

PELT below is an algorithm that finds the dates where a series shifts to a new level; Mann-Kendall with the Sen slope is a standard trend test.

| Cadence | Series | Latest (2026-08-28) | What a move would mean |
|---|---|---|---|
| daily | composite index (PCI), daily and 5-day average | +1.61 daily (53rd pct); +21.04 5-day ($P = 0.8494$) | $P$ below 0.80 ends the bullish class; the daily print leads and is already mid-range |
| daily | SPY gamma-flip distance | $-0.06$% (flip 769.62); 116 of 165 sessions negative YTD | the flip is the price where dealer hedging switches from dampening moves to amplifying them; above = out of the year's amplifying norm, lower = $F$ rebuilds |
| daily | 25-delta risk reversal (H2) | +0.02605 (18th pct; YTD floor 0.01773, Aug 4) | a break below the floor = crash protection getting bid; drags the 5-day HEDGE average toward its 0.80 gate |
| daily | VIX IV rank (H6) | 5.4491 (1st pct) | strongest carrier of the qualifier, at the year's floor; any lift reprices insurance (dropping H6 alone un-fires the gate, 0.7349) |
| weekly | regime re-fit (PELT) | 6 regimes; current since Jul 28, mean +7.72, "neutral, flat" | a new breakpoint or first non-neutral label would square the slow regime read with the fast 5-day rule |
| weekly | trend test (Mann-Kendall + Sen) | $p = 0.4961$; slope $-0.0219$ per session | $p < 0.05$ would turn one hot week into a year-long trend; it has not |

(Source: m4.json: pci_latest_day, pci_latest_5d, verdict_P, n_regimes, regime_6_mean_pci, mann_kendall_p, sen_slope_per_session; m2.json: gamma_flip_level_SPY, gamma_flip_distance_pct, days_negative_gamma_ytd; m3.json: H2_RR_latest, H2_RR_min_YTD, H6_VIX_iv_rank_latest)

Three caveats bound every trigger in this chapter. First, every monitored series except the Chapter 13 long-history proxies is single-vendor (Unusual Whales), with no in-sample cross-source validation. Second, the volatility-risk-premium input changes definition inside the scored window: the official forward-looking series ends July 31 and all five verdict sessions use the trailing-volatility extension; dropping the component leaves the qualifier firing (0.8434), so the splice is disclosed, not decisive. Third, two soft data breaks sit inside the unrestricted 2026 window (QQQ put premium, January 28; QQQ call delta, April 9), about 3.75% of the SPEC blend. Sources: m3.json notes; mod0_breaks.json; ADVERSARIAL_LOG2.md finding 7.

This dossier is research, not investment advice. The triggers above say when the study's description of positioning would change, not when to trade, and the power caveat from Chapters 2 and 12 applies in full: 165 sessions rank a day within the year; they cannot forecast returns.

# A. Data Tables

This appendix holds the reference tables behind Chapters 1 to 8: the month-by-month composite summary, the latest-session readings for every component of the three sub-indices, and the full regime segmentation. All values are as of the 2026-08-28 close, from the module result files. Scores run from −100 (extreme bearish consensus) to +100 (extreme bullish consensus) after each component's sign convention. The three sub-indices, in human terms: SPEC is what the crowd is actually buying, DEALER is what market makers are forced to do, and HEDGE is what protection costs. Percentiles anchor a value inside a distribution; the 85th percentile of 2026 means higher than 85 of every 100 trading days this year. Frame A ranks a value against every trading day of 2026 so far (the 2026 rank, 165 sessions from 2026-01-02 to 2026-08-28); Frame B ranks it against the full multi-year history available for that series (the multi-year rank).

Frame B conventions differ across modules because not every series has a real multi-year path. The M1, M4, and M5 blended composites have no pre-2026 history as blends, so their Frame B percentile is a weighted average of component percentiles, not a true multi-year rank. The M3 HEDGE headline ranks the 5-day average within the Frame B daily series. The M2 DEALER Frame B blends full-history inputs (D1) with inputs that exist only in 2026 (D4, D5). Cross-module Frame B comparisons are therefore indicative, not exact (Chapter 2). Table A.1's PCI minimum and maximum columns are computed from the daily `PCI_A` series in m4.json.

## A.1 Monthly composite summary

| Month | Sessions | PCI mean | PCI min | PCI max | SPEC | HEDGE | DEALER |
|---|---|---|---|---|---|---|---|
| 2026-01 | 20 | +11.086 | −22.26 | +34.15 | +2.167 | +9.331 | −0.411 |
| 2026-02 | 19 | −0.680 | −43.73 | +27.57 | −2.066 | −2.289 | +3.676 |
| 2026-03 | 22 | −4.538 | −30.20 | +30.16 | −1.133 | −12.276 | +8.870 |
| 2026-04 | 21 | −8.385 | −48.06 | +16.36 | +1.070 | −9.846 | +0.391 |
| 2026-05 | 20 | +0.804 | −30.38 | +33.50 | +2.366 | +3.785 | −5.347 |
| 2026-06 | 21 | −5.334 | −34.37 | +25.29 | −1.536 | +0.319 | −4.118 |
| 2026-07 | 22 | +0.933 | −25.53 | +34.88 | −1.901 | +1.628 | +1.206 |
| 2026-08 | 20 | +7.024 | −22.75 | +46.15 | +1.256 | +10.672 | −4.905 |

Table A.1. Monthly composite summary, Frame A. SPEC, HEDGE, and DEALER columns are monthly means of the daily weighted contributions (0.45, 0.35, and 0.20 times the respective sub-score); they sum to the PCI mean by construction. Session counts sum to the 165-session YTD window. The August row covers 20 sessions through 2026-08-28; it contains the +46.15 YTD peak (2026-08-03) and is the strongest HEDGE month of the year, meaning protection was cheapest and least bid. April contains the −48.06 YTD trough (2026-04-21). (Source: m4.json series waterfall_monthly, PCI_A; headline_numbers)

## A.2 Latest component readings

Component rows quote the 2026-08-28 session as Frame A scores, since per-component Frame B scores are not stored for every component (Chapter 2). Raw-value percentiles and sub-index aggregates quote

both frames where they exist.

### A.2.1 SPEC components (M1)

| Id | Component | Sign | Score (A) |
|---|---|---|---|
| S1 | Net premium tilt | + | +16.33 |
| S2 | Premium C/P ratio | + | +1.75 |
| S3 | Volume C/P ratio | + | +27.83 |
| S4 | Ratio gap $z_A$(S3) − $z_A$(S2) [raw z-gap, not in blend] | unsigned | +0.547 |
| S5 | Bullish premium share | + | +13.61 |
| S6 | Call ask-side share | + | −24.70 |
| S7 | Put ask-side share | − | −53.19 |
| S8 | Out-of-the-money (OTM) net-premium tilt, market-wide [as implemented] | + | −22.89 |
| S9 | Volume vs 30d average, signed by sign(S1) | sign(S1) | +16.69 |
| SPY sub-index (weight 0.40) | | | +15.81 |
| QQQ sub-index (weight 0.30) | | | +2.26 |
| IWM sub-index (weight 0.15) | | | −25.90 |
| DIA sub-index (weight 0.15) | | | −41.27 |
| Market SPEC, daily (daily rank A 49.4%) | | Frame A | −3.07 |
| Market SPEC, daily | | Frame B | −1.98 |
| Market SPEC, 5d mean (pct A 77.2%, B 60.4%) | | | +21.26 |

Table A.2. SPEC latest readings, 2026-08-28. Positive scores mean the crowd is leaning toward upside bets for this year; negative means it is not. Ask-side share counts trades that hit the seller's asking price, i.e. buyer-initiated trades. S4 is the raw z-score gap between the volume and premium call/put ratios (Frame A percentile 80.1%), reported as a diagnostic and excluded from the blend. S1, S5, and S9 form a near-duplicate block (YTD score correlations 0.871, 0.762, 0.677), so net-premium direction carries roughly 3 of 8 units of blend weight; S9 also takes its sign from S1 by construction, so it jumps discontinuously whenever S1 crosses zero (three S1 sign flips in the last 10 sessions); see Chapter 4. QQQ component Frame B windows are restricted per the break audit (Chapter 3). (Source: m1.json series S1_A..S9_A, SUB_*_A, SPEC_A/B, SPEC_A_pct; headline_numbers; adversarial_log.json for the redundancy correlations and S1 flip count)

### A.2.2 DEALER components (M2)

### A.2.3 HEDGE components (M3)

### A.2.4 Composite (M4)

## A.3 Regime segmentation

| Id | Series | Raw value | Pct (A) | Pct (B) |
|---|---|---|---|---|
| D1 | Customer delta demand, SPY | −88,059,059.7 | 31.3% | 37.1% |
| D2 | Net gamma (GEX), SPY [amplification axis, not in PCI] | −12,649.7 | 69.9% | 72.0% |
| D4 | NOPE daily close, SPY | −5.4673 | 12.7% | n/a |
| D5 | Directional delta flow, SPY | 496,162.3 | 47.6% | n/a |
| D5 | Directional delta flow, QQQ | −1,669,184.3 | n/a | n/a |
| D6 | Vanna, SPY [descriptive only] | −631,711,312.8 | 5.4% | 36.9% |
| D6 | Charm, SPY [descriptive only] | 245,193,417.2 | 67.5% | 64.8% |
| Direction score, SPY (Frame A) | | −38.96 | | |
| Direction score, QQQ (Frame A) | | −38.55 | | |
| Direction score, IWM (Frame A) | | +18.07 | | |
| Direction score, DIA (Frame A) | | −77.11 | | |
| Market DEALER direction, daily (Frame A / B) | | −36.00 / −22.03 | | |
| Market DEALER direction, 5d mean | | +2.48 | 50.6% | 54.8% |

Table A.3. DEALER latest readings, 2026-08-28. Net gamma (GEX) measures whether market makers' hedging dampens market moves (positive) or amplifies them (negative). NOPE is a measure of how much option trading is pushing around the stock itself. Vanna and charm are second-order dealer-hedging exposures, shown for description only and excluded from all scores under the break audit (Chapter 3); their Frame B figures are post-break segment percentiles (vanna from 2023-09-12, charm from 2024-08-28). The direction blend equal-weights D1, D4, and D5 per ticker, then applies cap-tier weights. D1's positive orientation rests on the vendor's dealer-inventory assumption (customer delta as the negated dealer book), which cannot be verified in-sample (Chapter 5). D4 and D5 harvests exist only for 2026, so they have no distinct multi-year history. The 5-day average (+2.48) sits at the 50.6th percentile, while the as-of session itself printed −36.0, the most negative day of that 5-day window. Scores in the lower block are oriented values, not raw units. (Source: m2.json series and headline_numbers)

| Amplification risk and mechanics metric | Value | Note |
|---|---|---|
| Amplification risk F, market | 0.3301 | Frame B 0.2947 |
| Amplification risk F, SPY | 0.3012 | Frame B 0.28 |
| Gamma-flip level, SPY | 769.62 | snapshot, no YTD distribution |
| Flip distance (spot − flip)/spot | −0.06% | spot just below the flip |
| Negative net-gamma sessions, YTD | 116 of 165 | the 2026 norm |
| Max pain, SPY next expiry (2026-08-31) | 760.0 | snapshot |

Table A.4. Dealer amplification risk and mechanics, 2026-08-28. F is the percentile of negative net gamma. High F means dealer hedging is positioned to amplify moves rather than dampen them: hedging forces dealers to sell as the market falls and buy as it rises. Low F means that pressure has eased. F feeds only the qualifier leg of the pre-registered verdict rule (Chapter 14), never the PCI blend. The reading collapsed in late August: F stood at 0.73 six sessions earlier (2026-08-20) and 0.33 at the window close, and the two frames now agree (Frame A 0.3301, Frame B 0.2947). The gamma flip is the price level where dealer hedging switches from dampening moves to amplifying them; spot closed 0.06% below it. Max pain is the strike where the most option value would expire worthless at the next expiry. (Source: m2.json headline_numbers; series F_market_A_pct)

| Id | Component | Sign | Raw | Pct (A) | Pct (B) | Score (A) |
|---|---|---|---|---|---|---|
| H1 | IV rank (iv_rank_1y), SPY | − | n/a | n/a | n/a | +97.59 |
| H2 | 25d risk reversal, 30-45 DTE splice | + | 0.02605 | 17.6% | 15.8% | −64.78 |
| H3 | VRP [spliced series] | − | 0.01373 | 46.4% | 44.1% | +7.23 |
| H4 | Term slope iv(90d) − iv(7d) | + | 0.041 | 77.7% | n/a | +55.42 |
| H5 | Short-vol ratio, SPY | − | n/a | n/a | n/a | −87.95 |
| H6 | VIX iv rank | − | 5.4491 | 1.2% | 4.0% | +97.59 |
| H7 | Implied move 30d percentile | − | n/a | n/a | n/a | +98.80 |
| HEDGE SPY, daily (Frame A / Frame B) | | | | | | +29.13 / +28.12 |
| HEDGE SPY, 5d mean (pct A 81.9%, B 84.9%) | | | | | | +31.35 |
| HEDGE QQQ, 5d mean, reduced form (pct A 89.8%) | | | | | | +34.80 |
| HEDGE IWM, 5d mean, reduced form (pct A 91.6%) | | | | | | +46.25 |

Table A.5. HEDGE latest readings, 2026-08-28. Scores are post-orientation: positive means insurance is cheap or unbid, which reads as complacency and bullish consensus. In plain terms: IV rank is where today's option prices sit in their one-year range (0 = cheapest, 100 = most expensive); the 25-delta risk reversal is the price of upside call options versus equally out-of-the-money downside puts, and when puts cost more the market is paying up for crash protection; VRP is the gap between what options charge for future volatility and the volatility that actually shows up; the term slope asks whether longer-dated options cost more than short-dated ones (calm) or less (stress). The sub-index's strength sits in the option-price-level components (H1, H6, H7, all near their year extremes) while the two demand-side components lean the other way (H2 −64.78, H5 −87.95): index option prices are near year lows, but the relative bid still favors the put side (Chapters 6 and 14). Raw values are stored for H2, H3, H4, and H6 only. H3 splices the official risk-premium series (forward-window realized volatility, through 2026-07-31) with a trailing 20d extension after; the two definitions differ at the 2026-08-03 seam (Chapter 6). H4's interpolated IV exists only for 2026, so its Frame B is not distinct. QQQ and IWM run reduced form (H1-H4, H7; no H5/H6), and DIA is excluded from ticker-level HEDGE. (Source: m3.json series score_H*_SPY_frameA, HEDGE_* series; headline_numbers)

| Series | Value | Pct (A) | Pct (B) |
|---|---|---|---|
| PCI, daily 2026-08-28 | +1.61 | 53.0% | 52.3% |
| PCI, 5d mean (verdict input P) | +21.04 | 84.9% | 61.3% |
| Contribution, SPEC (0.45 × score) | −1.381 | 49.4% | n/a |
| Contribution, HEDGE (0.35 × score) | +10.195 | 77.1% | n/a |
| Contribution, DEALER (0.20 × score) | −7.201 | 17.5% | n/a |
| Rolling 20d slope | +0.437 | 72.1% | n/a |

Table A.6. Composite readings, 2026-08-28. Contribution percentiles rank each contribution within its own YTD series. The 5-day Frame A percentile (0.8494) is the verdict input P of Chapter 14. The same reading at the multi-year convention is 0.6126, which the verdict rule would call neutral; the composite's Frame B is a component-average convention, not a true multi-year rank (Chapter 2). The as-of session itself printed +1.61, the 53rd percentile of the year: the elevated 5-day reading is carried by the final week of August, not by the latest day. The rule's overly-bullish check reads the HEDGE 5-day percentile (0.8193, Table A.5) and amplification risk F (0.3301, Table A.4) against gates of 0.80 and 0.70; only the HEDGE leg clears, by 0.019 (Chapter 14). (Source: m4.json headline_numbers; m3.json and m2.json for the gate legs)

| # | Start | End | Days | Mean PCI | Pct (A) | Sen slope | Label | Event at start |
|---|---|---|---|---|---|---|---|---|
| 1 | 2026-01-02 | 2026-02-02 | 21 | +10.98 | 71.7% | −0.370 | neutral, flat | YTD start |
| 2 | 2026-02-03 | 2026-03-04 | 21 | +0.16 | 50.6% | +0.320 | neutral, flat | none |
| 3 | 2026-03-05 | 2026-05-04 | 42 | −7.85 | 37.4% | −0.111 | neutral, flat | NFP 03-06 |
| 4 | 2026-05-05 | 2026-06-04 | 22 | +3.74 | 57.8% | +0.901 | neutral, rising | NFP 05-08 |
| 5 | 2026-06-05 | 2026-07-27 | 35 | −4.97 | 42.8% | +0.164 | neutral, flat | NFP 06-05 |
| 6 | 2026-07-28 | 2026-08-28 | 24 | +7.72 | 64.5% | +0.454 | neutral, flat | FOMC 07-29, PCE 07-30 |

Table A.7. Full regime table from the change-point step (PELT, an algorithm that finds the dates where a data series shifts to a new level; rbf cost, minimum segment 20 sessions, penalty 0.355, 6 segments). Labels combine the level rule (mean PCI at or above +20 bullish tilt, at or below −20 bearish tilt, else neutral) with the within-segment trajectory from the standard statistical trend test (Sen slope); all six regimes classify neutral on level (five flat, one rising), with means spanning −7.85 to +10.98. Pct (A) ranks each regime mean within the YTD daily PCI distribution. Event tags are FOMC/CPI/NFP/PCE/PPI releases within 3 calendar days of the regime start. The current regime (2026-07-28 to 2026-08-28) classifies neutral and flat on its 24-day mean (+7.72) even though the final five sessions rank at the 84.9th percentile; the segmentation reads levels over a window, not the latest week (Chapter 14).
(Source: m4.json series regime_table; headline_numbers regime_*_mean_pci)

# B. Endpoint Inventory

This appendix records every data feed the study pulled, what each one returned, and what the full build cost in requests. It exists so a reader can see exactly what the evidence rests on and rebuild the dataset from scratch. Three words repeat below: an endpoint is one URL pattern on the vendor's data service, a call is one request to it, and the lake is the local file store the downloads land in. Everything comes from a single vendor's feed, so the single-vendor caveat from Chapter 2 travels with every row.

The build ran in three passes. On 2026-08-20 we probed the vendor's REST API with 64 test calls across 30 recorded endpoints (`results/probe_endpoints.json`), establishing per family the response status, historical depth, row shape, and field list, and surfacing the paths our subscription tier blocks. The production harvest that day logged 3,852 calls. The 2026-08-30 refresh (Amendments A4 and A5) then brought every series to the 2026-08-28 close and added the semiconductor universe of Chapter 10, at a further 435 calls. The lake now holds 1,051,351 rows across 4,092 endpoint pulls; cumulative logged calls stand at 4,287, or 4,351 counting the probes.

## B.1 Probe results

A probe is a small test call made before any real harvesting, to learn how an endpoint behaves. Status 200 means the endpoint answered with data; 403 means it is blocked at our subscription tier; 404 means the path does not exist; 422 means it rejected the parameters. The table is the probe-day record (2026-08-20). All paths are relative to `/api`. Probes ran against SPY (and the Technology sector tide); depth and shape generalize across the universe per the contract.

| Path | Params | Status | Rows | Fields | Coverage / shape |
|---|---|---|---|---|---|
| `stock/SPY/iv-rank` | `timespan=1y` | 200 | 251 | 5 | 2025-08-21 to 2026-08-20 |
| `stock/SPY/iv-rank` | none | 200 | 5 | 5 | 2026-08-14 to 2026-08-20 |
| `stock/SPY/historical-risk-reversal-skew` | `timeframe=1Y`, `expiry=30`, `delta=25` | 422 | n/a | n/a | param combination rejected |
| `stock/SPY/historical-risk-reversal-skew` | none | 200 | 11 | 4 | 2026-08-06 to 2026-08-20 |
| `stock/SPY/volatility/variance-risk-premium` | none | 200 | 231 | 5 | 2025-08-21 to 2026-07-23 (feed lag) |
| `stock/SPY/volatility/realized` | `timeframe=1Y` | 200 | 251 | 5 | 2025-08-21 to 2026-08-20 |
| `stock/SPY/volatility/stats` | none | 200 | 1 | 9 | snapshot |
| `stock/SPY/volatility/term-structure` | none | 200 | 34 | 7 | 2026-08-20, per-expiry |
| `volatility/vix-term-structure` | `history_days=400` | 403 | n/a | n/a | plan-gated |
| `stock/SPY/nope` | none | 200 | 390 | 10 | 2026-08-20 intraday, 1-min |
| `stock/SPY/spot-exposures` | none | 200 | 472 | 14 | current day, 1-min |
| `stock/SPY/spot-exposures/strike` | `limit=50` | 200 | 50 | 37 | 2026-08-20, per-strike |

| Path | Params | Status | Rows | Fields | Coverage / shape |
|---|---|---|---|---|---|
| options-pulse/total | none | 403 | n/a | n/a | plan-gated |
| market/total-options-volume | limit=500 | 200 | 500 | 5 | 2024-08-22 to 2026-08-20 |
| market/oi-change | limit=50 | 200 | 50 | 25 | current-day top movers |
| market/market-tide | interval_5m=true | 200 | 81 | 5 | 2026-08-20, 5-min bars |
| market/market-tide | interval_5m=true, otm_only=true | 200 | 81 | 5 | 2026-08-20, 5-min, OTM only |
| market/Technology/sector-tide | none | 200 | 391 | 5 | 2026-08-20 intraday, 1-min |
| market/SPY/etf-tide | none | 200 | 390 | 6 | 2026-08-20 intraday, 1-min |
| screener/option-contracts | limit=20, order=premium, min_dte=21, max_dte=90 | 200 | 20 | 45 | current-day snapshot |
| stock/SPY/greek-exposure | timeframe=3Y | 200 | 752 | 9 | 2023-08-22 to 2026-08-20 |
| stock/SPY/greek-exposure | timeframe=5Y | 200 | 1,102 | 9 | 2022-03-30 to 2026-08-20 |
| stock/SPY/max-pain | none | 200 | 33 | 6 | current day, per-expiry |
| stock/SPY/interpolated-iv | none | 200 | 9 | 5 | 2026-08-20, 9-tenor grid |
| shorts/SPY/volume-and-ratio | none | 200 | 500 | 4 | 500 daily rows (key si) |
| darkpool/SPY | limit=20 | 200 | 20 | 18 | latest prints |
| news/headlines | limit=20 | 200 | 20 | 8 | latest headlines |
| market/events | none | 404 | n/a | n/a | not found; see Amendment A2 |
| stock/SPY/stock-state | none | 200 | 1 | 9 | snapshot |
| stock/SPY/ohlc/1d | limit=500 | 200 | 756 | 8 | 2025-08-20 to 2026-08-20 |

(Source: results/probe_endpoints.json; probed 2026-08-20)

Six behaviors found in the probes shaped the harvest design and the refresh:

- iv-rank returns a fixed 251-row one-year span; timespan values above 1y are ignored. The same one-year cap applies to volatility/realized.
- ohlc/1d returned 756 rows over 252 regular sessions because each session carries three market_time buckets (pr, r, po: pre-market, regular, post-market); only regular-session rows are used (Amendment A1). The same one-year price cap is why the forward-return tests in Chapter 12 run on vendor closes covering 2025-08-28 to 2026-08-28 (the 63-day horizon loses the final quarter's tail), and why the long-history replication in Chapter 13 uses non-vendor closes.
- variance-risk-premium runs about a month behind the tape: the probe found it stalled at 2026-07-23, and the refreshed pull reaches only 2026-07-31 against the 2026-08-28 as-of date. The date parameter is ignored, so the lag is a hard right edge. The live tail is computed instead as implied minus realized volatility from iv-rank plus realized (the seam of 2026-08-03 disclosed in Chapter 6).
- historical-risk-reversal-skew rejects the documented timeframe parameters (422) and requires a real expiry date; with one supplied it returns roughly 250 daily rows for that expiry's 25-delta risk reversal, the price of upside calls versus equally out-of-the-money downside puts. A single expiry ages one day per day, so the constant-maturity 30-to-45-day series in Chapters 6 and 10 are stitched from 12

monthly-expiry pulls per ticker, 60 pulls in all across SPY, QQQ, IWM, SMH, and MU (contract Section 6.3).

- `market/total-options-volume` rejects `limit` above 500 (422), fixing the family's depth at a rolling 500-day window; refreshed, that window covers 2024-08-30 to 2026-08-28.
- The per-date families return one trading session per call: `nope`, `market-tide`, `sector-tide`, `greek-flow`, `interpolated-iv`, `expiry-breakdown`, `max-pain`, and `oi-change`. This is what drives the call budget in Section B.3.

## B.2 Gated, failed, and excluded paths

Two endpoints sit behind a higher subscription tier, one documented path does not exist, and a few reject their documented parameters. None of them feeds a score.

| Path | Status | Disposition |
|---|---|---|
| `volatility/vix-term-structure` | 403 | Blocked at this subscription tier; excluded outright and noted here per contract Section 4 |
| `options-pulse/total` | 403 | Blocked at this subscription tier; excluded outright |
| `market/events` | 404 | Chronology routed through the vendor's MCP tool interface instead (Amendment A2) |
| `stock/{t}/greeks?date=D` | 200 | Returns zero rows for historical dates; current-only, excluded from all history |
| `stock/VIX/ohlc`, `stock/VIX/stock-state` | 422 | VIX is an index, not a stock; VIX spot taken from the `iv-rank close` field |
| `historical-risk-reversal-skew` with timeframe params | 422 | Requires a real expiry date parameter; monthly-expiry stitch used instead |
| `market/total-options-volume` with `limit` above 500 | 422 | Depth capped at 500 rows |

(Source: results/probe_endpoints.json; RESEARCH_PLAN.md Section 4)

No gated endpoint is load-bearing for any score: the two 403 paths would have served only descriptive context, and each has an in-scope substitute (the interpolated-IV grid and the total-options-volume series respectively).

## B.3 Call budget

The initial 2026-08-20 harvest logged 3,852 calls. The 2026-08-30 refresh added 435 more: 138 to extend the per-date families across the six new sessions (2026-08-21 to 2026-08-28), 159 to re-pull the one-call and snapshot families fresh, 36 to rebuild the market risk-reversal stitches, and 102 for the semiconductor additions of Amendment A5 (13 tickers across six families, plus 24 SMH and MU expiry pulls). Re-pulled families overwrite their manifest keys, which is why the lake's 4,092 stored pulls sit 195 below the 4,287 cumulative calls. The per-date families dominate the count, 3,795 of the 4,092 pulls, because each call returns a single session; sector tide alone is 1,815 pulls, 11 sectors times 165 sessions.

| Family | Calls | Rows | Call structure |
|---|---|---|---|
| sector_tide | 1,815 | 708,891 | 11 GICS sectors × 165 sessions, intraday |
| interp_iv | 660 | 5,940 | 4 index tickers × 165 sessions, per-date IV grid, 9 horizons |
| expiry_breakdown | 495 | 16,043 | 3 tickers × 165 sessions, per-date days-to-expiry buckets |
| greek_flow | 330 | 133,762 | 2 tickers × 165 sessions, intraday |
| market_tide | 330 | 26,830 | 165 sessions × 2 variants (all trades, out-of-the-money only) |
| nope | 165 | 64,350 | SPY × 165 sessions, intraday |
| rr_skew | 60 | 7,985 | 12 monthly expiries × 5 tickers for the constant-maturity stitches |
| greek_exposure | 36 | 38,677 | one call per harvested ticker, 5Y depth |
| iv_rank | 36 | 9,036 | one call per harvested ticker |
| options_volume | 36 | 17,876 | one call per harvested ticker |
| realized | 36 | 9,030 | one call per harvested ticker |
| vol_stats | 36 | 36 | one snapshot per harvested ticker |
| vrp | 36 | 8,312 | one call per harvested ticker |
| flow_per_expiry | 4 | 115 | current-day snapshots, 4 index tickers |
| max_pain | 4 | 116 | current-day snapshots, 4 index tickers |
| screener | 4 | 750 | current-day snapshots by days-to-expiry bucket (same-day bucket empty at the re-pull) |
| spot_gex_strike | 4 | 1,502 | current-day strike maps, 4 index tickers |
| short_vol | 3 | 1,500 | SPY, QQQ, IWM; 500 rows each |
| misc | 2 | 600 | news headlines; market total options volume |
| **Lake total** | **4,092** | **1,051,351** | 4,287 logged calls including re-pulls; 4,351 with the 64 probes |

(Source: harvest manifest, `manifest.sqlite` in the data lake; probe count per RESEARCH_PLAN.md Section 4)

The 36-call families map one call to each harvested ticker: the 23 core names (5 index complex, 11 sector ETFs, 7 mega-caps) plus the 13 semiconductor additions of Amendment A5 (the SMH ETF and 12 stocks; NVDA was already in the core set). Lake layout and manifest schema are in contract Section 11; reproducibility details, including the deterministic seed and the byte-identical rerun check, are in Appendix C.

# C. Reproducibility and Verification

This appendix records how the study was run and what three independent checking passes found: a blind re-derivation of every headline number, a sign-convention audit, and two rounds of adversarial attack, one on the original August 20 verdict and one on the refreshed August 28 verdict this report prints. Throughout, SPEC measures what the option-buying crowd is actually buying, DEALER what market makers are forced to do, and HEDGE what protection costs; the PCI is their weighted blend, the study's single positioning score. The verification artifacts ship with the study (`QC_REPORT.md`, `results/sign_audit.json`, and the two adversarial logs under `results/`).

## C.1 Run configuration and modules

The study first ran on 2026-08-20 and was fully refreshed on 2026-08-30 to the 2026-08-28 close (amendment A4): six more sessions, for a 2026 window of 165 trading days (2026-01-02 to 2026-08-28). The seed for every step that involves randomness is 20260820, fixed before the first run, unchanged through the refresh, and stored in each module JSON. Inputs sit in the local parquet lake (`C:\uw-data\positioning`, indexed by `manifest.sqlite`); each module writes `results/mN.json` with series, headline numbers (value, the 2026 rank, the multi-year rank), figures, notes, and self-checks. Figures use `ftstyle.py`; endpoint inventory and call budget: Appendix B.

| Module | Role | Output | Note |
|---|---|---|---|
| M0 | Break audit | `mod0_breaks.json` | 226 series after the refresh: 52 HARD, 79 SOFT, 95 clean. PELT (an algorithm that finds where a series shifts to a new level) on $\log_{10}\lvert x\rvert$, confirmed by CUSUM; HARD at a 10x jump in the segment median, SOFT at 2x |
| M1 | SPEC sub-index | `m1.json` | S1 to S9, four index tickers plus the market blend |
| M2 | DEALER sub-index | `m2.json` | D1 to D6, the amplification-risk axis, the gamma flip (the price where dealer hedging switches from dampening moves to amplifying them) |
| M3 | HEDGE sub-index | `m3.json` | H1 to H7; SPY full form, QQQ/IWM reduced (H1 to H4, H7) |
| M4 | Composite, trend, verdict | `m4.json` | PCI blend, trend test (Mann-Kendall/Sen), PELT regimes, the sec-8 verdict rule |
| M5 | Cross-section | `m5.json` | 23-name mini-PCI, sectors, concentration |
| M6 | Validation and long history | `m6.json` | Coherence, quintiles, break robustness (Ch. 12); the non-UW series behind Ch. 13 |
| M7 | Semis and memory | `m7.json` | 14-name semis universe (amendment A5, Ch. 10); same reduced-form mini-PCI as M5 |

**Determinism.** M1 through M5 and M7 contain no random steps; M4's trend test is implemented directly with tie-corrected variance and is deterministic. M6 is the only module with randomness. Its error bars are

built by resampling whole stretches of days, which respects the fact that markets trend (a block bootstrap: overlapping blocks of length 10, 5,000 draws). Each cell derives its seed from the master seed (20260820 + 1000 × horizon + quintile; 500,000 offset for the non-UW grid); the 1,000-shuffle placebo is seeded the same way. M6's stored reproducibility checks confirm that re-running the seeded grids on the refreshed data reproduces the stored results identically.

## C.2 Blind QC re-derivation

A QC pass rebuilt every headline number of the original 2026-08-20 run from the raw lake using only the written definitions in the research contract and the modules' own notes. No module code was opened; that is what makes the pass blind. Of 89 headline keys: 79 PASS (within 0.5% relative or 0.01 absolute), 1 RECONCILED, 0 FAIL, 9 SKIPPED (seeded bootstrap grids and checks needing artifacts outside the blind input set). The one RECONCILED row is the DEALER multi-year percentile slot, a documented presentation choice. That run's verdict rule reproduced independently; no value, percentile, or conclusion changed. The refresh re-ran the same audited code on six more sessions with definitions unchanged, and the internal identity checks pass on the refreshed data (the contribution waterfall sums exactly to the PCI; the SPEC and HEDGE blends reconstruct from stored component scores). Scratch code: `code/qc/`.
(Source: QC_REPORT.md; m4.json and m6.json checks)

Three structural observations from the QC pass remain true after the refresh:

1. **S8 contract wording.** The plan's wording for SPEC component S8 does not describe the implemented quantity, the out-of-the-money (OTM) net-premium tilt; the implementation is documented in the module notes, passed the sign audit, and a wording amendment is recommended (Section C.6).
2. **The multi-year rank mixes conventions.** The multi-year (Frame B) slot is computed three different ways across modules; for the blended composite it is a component-average convention, not a rank in any real multi-year history, because the blend has no pre-2026 path. Round 2 promoted this caveat into the verdict print itself (Chapter 14).
3. **Mixed tie handling.** Some inputs are stored at coarse precision, which creates rank ties, and the modules break those ties in more than one way. This is the source of every residual sub-0.01 percentile wiggle; the loosest reproduction in the pass, the trend test's z statistic at 0.9% relative distance, traces to the same tie ambiguity and changes no conclusion.

## C.3 Sign-convention audit

Every component sign was walked from contract to code to an economic spot check on the SPEC peak and trough dates (2026-08-04 and 2026-02-12, both unchanged by the refresh): 24 checks, 23 OK, 1 FLAG. The FLAG is the confidence rule as coded in M4. Dissent is counted against the majority sign of the three sub-indices, so the count is always 0 or 1 and the rule's low-confidence branch can never fire; the contract's other conditions for high confidence (a clean QC and no broken load-bearing series) are not wired in at all. Immaterial in round 1, load-bearing in round 2: the refreshed run printed high confidence on the sign count alone, and the adversarial pass killed that print partly because the QC-clean condition was never

evaluated (Section C.5). Two smaller notes: M1's field map for S9 omits the premium fields feeding sign(S1), a multi-year-mapping nuance; and a latent bug in the shared library's close for NOPE (a measure of how much option trading pushes around the stock itself), reverse-sorted files returning the open bar, was caught and overridden in M2. (Source: results/sign_audit.json)

## C.4 Adversarial review, round 1 (the retired August 20 verdict)

An adversarial pass attacks the report's own headline claims with independent recomputation from the lake and forces edits where a claim fails. Round 1 ran on the original August 20 verdict, which then read neutral with a bearish tilt. Nine claims were attacked: three survived, four softened, two killed. The refresh retired that verdict, and its verdict-specific findings do not transfer: dropping the entire HEDGE sub-index then pushed the composite rank from 0.406 to 0.075, bearish; on the refreshed data the same drop gives 0.747, still neutral. The structural findings carry forward and are applied throughout this report:

- **Negative net gamma is the 2026 norm, not news.** When market makers are net short gamma, their hedging sells as the market falls and buys as it rises, amplifying moves. That has been the state on 116 of 165 sessions this year, so the report calls it the year's normal condition, never elevated fragility.
- **The validation result is episode-driven.** The contrarian forward-return spreads lean on the January-to-March drawdown and rebound, and a naive momentum benchmark produces larger spreads. Chapter 12 presents it that way.
- **One vendor.** Every number in Chapters 1 to 12 is Unusual Whales data, with no independent in-sample cross-check; the H3 series splices two definitions of the volatility premium (official feed to 2026-07-31, a trailing-volatility extension after). Disclosed in Chapter 2 and next to the verdict.
- **Three SPEC components overlap.** S1/S5/S9 substantially repeat one signal, net-premium direction. Disclosed where SPEC is built; round 2 re-tested it against the new verdict (below).
- **Directional tilt statements are weight-conditional** and are labelled as such wherever they appear.

(Source: results/ADVERSARIAL_LOG.md; refreshed counts from m2.json and m6.json)

## C.5 Adversarial review, round 2 (the refreshed verdict)

Round 2 ran after the refresh, on the verdict this report prints. Its main input is P, where the five-day average of the composite ranks against every trading day of 2026 so far, written as a fraction: P $= 0.8494$ means higher than about 85 of every 100 days this year. The attacker replicated every verdict input exactly before attacking (composite rebuild within $5 \times 10^{-5}$). Eight claims: three survived, two softened, three killed; the killed and softened rows dictated the exact wording in The Read and Chapter 14. (Source: results/ADVERSARIAL_LOG2.md)

| Claim | Outcome | Finding and disposition |
|---|---|---|
| Verdict class: bullish consensus (P = 0.8494) | Softened | Robust inside 2026: all three rule readings fire, any single-day drop keeps P at or above 0.8253, 20 of 21 leave-one-out variants keep the class. Four caveats print with it: (i) the multi-year rank reads neutral under the same rule (0.6126, a component-average convention); (ii) the class is 3 sessions old, P swung 0.3735 to 0.8494 in six sessions, the as-of day printed the 53rd percentile; (iii) 14 of 165 days classified bullish this year and the earlier clusters preceded the April drawdown and mid-August fade (next-10-day SPY mean −0.13%, hit 18%, descriptive n = 11); (iv) the regime table calls the current regime neutral and flat (mean +7.7, trend p = 0.496). |
| “Overly bullish” qualifier | Killed | Fires on one leg: HEDGE 0.8193 against the 0.80 gate, the other leg (F, the amplification-risk gauge) dead at 0.3301. Margin 1.9 points = 3.2 rank positions; dropping H1, H6, or H7 un-fires it (0.7952 / 0.7349 / 0.7530); price-level components carry it (H1 +74.0, H6 +92.5, H7 +88.7) while both demand components dissent (H2 −55.2, H5 −47.7); the as-of day’s HEDGE rank is 0.7711, below the gate; one-session survival about 42%. The report prints “at the overly-bullish boundary” with the margin, the un-fires, the level-versus-demand split, and the dead F leg, never the bare qualifier. |
| Confidence high (0 of 3 dissent) | Killed | The zero-dissent count rests on DEALER’s 5-day mean of +2.5 on a −100 to +100 scale, about 7% of its typical daily swing. On the as-of day’s values the majority sign is negative (SPEC −3.1, HEDGE +29.1, DEALER −36.0) and the rule reads medium; the newest DEALER print is the window’s most negative; 9 of 10 next-session replays lose the label; the QC-clean condition was never evaluated (the sign-audit FLAG). The report prints medium-grade confidence and names the statistic used. |
| Print-pipeline consistency | Killed | The refreshed verdict existed only in the JSONs: every report page still said neutral, fourteen files carried retired numbers, the semis chapter did not exist, and two module note strings contradicted their own stored checks (159 sessions versus 165; quintile sizes 31 to 32 versus the actual 33). Fix: rebuild every verdict-bearing file against the new verdict, write Chapter 10 from `m7.json`, regenerate the note templates. This document is that rebuild. |
| P convention ambiguity | Survives | All three admissible readings of the percentile sentence fire (0.8494, 0.8265, 0.9815); the strictest is the most extreme (third-highest smoothed reading of the year). The attack strengthened the class. |
| S1/S5/S9 redundancy drives the class | Survives | The late-August SPEC move is broad-based: collapsing the overlapping trio into one voice raises SPEC (21.3 to 23.7) and leaves P at 0.8373. The round-1 disclosure stands; the artifact claim failed. |
| H3 seam / flagged-series exposure | Survives | All five verdict-window H3 values come from the trailing-volatility extension, a splice inside the scored window; but H3 is mid-pack (+34.2) and dropping it changes nothing (qualifier still fires at 0.8434). The refreshed break audit adds no flag inside the market blend (new SOFT flags are sector and megacap series). Disclosure carries, next to the verdict. |
| “Overly” OR-gate design | Softened | Both rounds fired on whichever leg sat just above its gate while the other was dead: F = 0.7194 against 0.70 in round 1 (killed as a frame artifact), HEDGE 0.8193 against 0.80 in round 2 with F at 0.33. F traversed 0.21 to 0.73 to 0.33 within ten sessions; 7 days this year sit inside the 0.78 to 0.82 HEDGE band. The report always prints the firing leg, its margin, and the other leg; Chapter 15 carries the design note. |

Across two verification rounds the arithmetic held and the words kept needing discipline. The blind QC pass reproduced every re-derivable headline with zero failures, the sign audit passed 23 of 24 checks, and both adversarial passes replicated the verdict inputs exactly before attacking. What failed was framing: round 1 killed two claims attached to a neutral verdict; round 2 killed the bare “overly bullish” qualifier, the bare “high confidence” stamp, and a print pipeline still carrying the retired verdict. This report prints the post-attack versions.

## C.6 Amendment log

- **A1 (2026-08-20).** UW daily price history caps at one year, so Chapter 12’s forward returns run on the UW window (2025-08-20 onward; the 63-day horizon loses the final quarter’s tail); the long-history replication moved to Chapter 13 on non-UW closes, labelled.
- **A2 (2026-08-20).** The market-events REST path returns 404 (not found); the chronology came through the MCP market-events tool (684 rows archived to the lake).
- **A3 (2026-08-20).** The Cboe equity put/call archive returns 403 (access denied); the Chapter 13 proxies are limited to the VIX, VIX3M, and SKEW series plus SPY closes.
- **A4 (2026-08-30).** Full data refresh through the 2026-08-28 close: six more sessions (165 YTD), one-call families re-pulled, the break audit and all modules re-run in dependency order on the same seed. Report language rewritten in plain English; definitions unchanged.
- **A5 (2026-08-30).** Semiconductor and memory section added as Chapter 10 with new module M7 (14 names, about 102 extra UW calls); later chapters renumbered; report author set to Felipe Sinisterra.
- **Proposed (from QC).** Amend the contract’s S8 wording to the implemented OTM net-premium tilt; recorded as a recommendation only.