

Felipe Sinisterra

Market Positioning Study Prompt

Copy and paste the complete prompt below into a new agent session.

```
# Goal

Create a roughly 50-page PDF answering: **Is the market overly bullish, neutral, or overly bearish right now, and how has that changed during 2026?**

Before collecting data, define the universe, endpoints, analytical modules, chapter structure, and report style. Do not rely on any external specification or private project files.

# Data workflow

Use Unusual Whales as the only source for options, flow, skew, gamma, and implied-volatility data. FRED or Polygon may be used only for a clearly labeled long-history comparison using VIX, SKEW, and CBOE put/call data.

Use the MCP to explore the API, but check `get_public_api_docs` before assuming the visible tools cover everything. The following require REST access:

- `iv-rank`
- `historical-risk-reversal-skew`
- `variance-risk-premium`
- `volatility/realized`
- `vix-term-structure`
- `nope`
- `spot-exposures`
- `options-pulse/total`

Create or use a REST client for bulk collection. Stay below 35,000 calls against the 40,000-call daily limit.

Before calculating historical percentiles, run CUSUM or PELT change-point detection on every series. Tabulate every structural break and decide whether to rescale, restrict, or drop the affected history.

Build three separate indicators:

- SPEC: speculative demand
- DEALER: market mechanics and gamma
- HEDGE: skew, variance risk premium, and IV rank

Calculate them across 0DTE, one-to-four-week, one-to-three-month, and LEAP maturities. Blend them into a Positioning Consensus Index from -100 to +100, with a component-attribution waterfall. Preserve disagreements between indicators.

Use the 2026 YTD percentile as the headline comparison. Show the multi-year percentile beside it as context.

# Output

Produce:

- `Market_Positioning_Consensus.pdf`
- Machine-readable `results/*.json` files
- `outputs/Positioning_Data_Appendix.xlsx`

State whether positioning is overly bullish, neutral, or overly bearish. Include confidence, component attribution, and historical base rates. Do not provide price targets or investment recommendations.

# Evals

Every number must trace to a results JSON. Seed all bootstraps.

Run a blind re-derivation with zero failures and an adversarial review that tries to disprove the verdict. Record both surviving and rejected conclusions.

Publish the break audit and clearly state that the YTD ranking is well powered while the forward-return analysis is not.
```