

AlphaQuery Intelligence — Flow-Anticipation Composite

Pipeline Specification v1.0 · Query Breadth Beta Flagship Strategy

Prepared: July 11, 2026 · **Target live (paper):** August 1, 2026 · **Cadence:** Weekly

1. Strategy statement

One sentence: *A weekly cross-commodity model that trades the gap between J.P. Morgan QDS's published positioning projections and last-reported CFTC positioning, conditioned on inventory tightness, term-structure carry, and spec crowding, with price momentum demoted to a gate rather than a signal.*

Edge hypothesis: the forecaster publishes its own flow forecast every Tuesday evening and does not trade it; the projection delta is perishable (decays over ~5 sessions) and is largely ignored by systematic CTAs, whose models are price-momentum-driven.

Discipline principle (Counterfactual Symmetry): the model acts identically on this week's signal regardless of last week's outcome. No discretionary overrides inside the composite; discretion lives only at the portfolio level (universe membership, kill switch).

2. Universe (v1)

Six markets, chosen for CFTC coverage + QDS projection coverage + IBKR liquidity + representation across sectors:

Market	IBKR symbol	Sector	Inventory proxy source
ICE Sugar No. 11	SB	Softs	Green Pool via JPM Fig 88–89 fallback; ISO quarterly

Market	IBKR symbol	Sector	Inventory proxy source
CBOT Wheat	ZW	Grains	USDA WASDE ending stocks / use
CBOT Corn	ZC	Grains	USDA WASDE ending stocks / use
CBOT Soybeans	ZS	Grains	USDA WASDE ending stocks / use
COMEX Gold	GC	Precious	N/A — S pillar weight redistributed (see §4.6)
NYMEX Natural Gas	NG	Energy	EIA weekly storage vs 5-yr seasonal

Deferred to v2: COMEX–LME copper spread (needs LME COTR ingestion + spread execution logic), ICE Coffee, ICE Cocoa (re-enters when momentum gate clears), Brent.

3. Component architecture

Five components, all on the existing Cloudflare stack:

None

C1

INGEST-CFTC

Worker + Cron (Fri 16:30 ET)

CFTC Socrata API → normalize → D1: positions

C2

INGEST-JPM

Worker + manual upload endpoint

PDF (Tue eve) → Table 1 extraction → D1: projections

C3

INGEST-MARKET

Worker + Cron (daily 17:15 ET)

IBKR MCP: settles F1/F2/F3 → D1: curves

EIA API (Thu), USDA WASDE (monthly) → D1: inventory

C4

SCORER

Worker + Cron (Wed 08:00 ET)

Computes S,T,O,R,M per market → composite → D1: scores

Emits target book → KV: target_positions

```
| C5 EXECUTOR           Worker, triggered by C4 completion |
|   Diffs target vs current → Kelly-Thorp tranche state   |
|   machine → IBKR paper orders → D1: fills, audit_log    |
```

Storage: D1 (relational history: positions, projections, curves, inventory, scores, fills, audit_log)
+ KV (current state: target_positions, tranche_state per market, kill_switch flag).

Kill switch: single KV boolean checked by C5 before any order. Flipped manually or by the drawdown guard (§7).

4. Pillar definitions

All pillars normalized to $[-1, +1]$ where $+1$ = maximally bullish for the market.

4.1 S — Situation (inventory), weight 0.20

Days-of-use (or stocks/use) z-score vs trailing 8 years of same-calendar-month observations (seasonal adjustment by construction), **inverted**:

$$S = -\text{clip}(z_{\text{inventory}}, -2, +2) / 2$$

Tight inventories (negative z) → positive S . Sources: EIA weekly (natgas), USDA WASDE (grains), ISO/Green Pool (sugar).

4.2 T — Tension (crowding), weight 0.30

Managed Money net position as percentile of its post-2018 distribution, **inverted**:

$$T = 1 - 2 \cdot \text{pctile}(\text{MM}_{\text{net}}, \text{post-2018 window})$$

Max spec short (pctile ≈ 0) → $T \approx +1$ (maximum squeeze fuel). Max spec long → $T \approx -1$.
Source: C1.

4.3 O — Options/carry (term structure), weight 0.15

Annualized front-spread roll yield, squashed:

$$\text{roll} = (F1 - F2) / F2 \times (365 / \text{days_between}) ; 0 = \tanh(\text{roll} / 0.15)$$

Backwardation → positive O. This is the direct port of the GPU term-structure lens:
backwardation = the physical market paying you to be long. Source: C3/IBKR.

4.4 R — Recommendation-flow (the QDS projection gap), weight 0.25

The differentiated input. From Table 1: $\text{gap} = \text{projection} - \text{last_reported}$ (USD mm), normalized by that market's total OI (Table 2 sector or market OI):

$$R = \tanh\left(\text{gap} / \text{OI} / 0.02\right)$$

A projected inflow of 2% of OI → $R \approx +0.76$. **Freshness decay:** R is multiplied by $\max(0, 1 - \text{sessions_since_publication}/5)$. Stale projection → R contributes nothing. Source: C2.

4.5 M — Mechanics (momentum gate), weight 0.10

From internal replication of a simple dual-lookback momentum z (long ≈ 350–500 sessions, short ≈ 60–100, per JPM Fig 18 conventions):

- If sign(short-term signal) agrees with sign(composite-ex-M): $M = \pm 1$ (adds its 0.10).
- If it disagrees: $M = 0$ (abstains — momentum never vetoes on direction alone).
- **Extreme override:** if $|z_{\text{short}}|$ or $|z_{\text{long}}| >$ that market's extreme threshold (± 1.5 to ± 3 per Fig 18 conventions; we adopt ± 2 uniformly in v1 for simplicity), **the entire composite is zeroed** for that market this week. This is the exhaustion brake — it is what benches cocoa-type setups automatically.

4.6 Missing-pillar handling

If a pillar is unavailable (e.g., no meaningful inventory series for gold), its weight is redistributed pro-rata across the remaining pillars. Log every redistribution to audit_log.

5. Composite, activation, sizing

$\text{Score} = 0.20 \cdot S + 0.30 \cdot T + 0.15 \cdot O + 0.25 \cdot R + 0.10 \cdot M$ (subject to §4.5 zeroing and §4.6 redistribution)

Activation (hysteresis band):

- Open long: $\text{Score} \geq +0.60$ · Open short: $\text{Score} \leq -0.60$
- Hold while $|\text{Score}| > 0.20$ · Exit when $|\text{Score}| \leq 0.20$ or sign flips
- Re-entry requires re-crossing ± 0.60 (no churn inside the band)

Sizing (fractional Kelly via AQ Kelly Sizer):

- Base fraction: 0.25 Kelly (quarter-Kelly, standard house conservatism)
- Conviction scalar: $k_{\text{eff}} = 0.25 \times (|\text{Score}| - 0.60) / 0.40$ mapped onto [0.10, 0.25] Kelly — a 0.62 score sizes near 0.10 Kelly, a 1.00 score at full quarter-Kelly
- Per-market notional cap: 15% of paper NAV · Portfolio gross cap: 60% of NAV
- Volatility input to Kelly: 60-day realized vol of the front contract

Entry mechanics (Kelly-Thorp tranche state machine):

- Tranche in thirds: 1/3 at signal (Wed open), 1/3 at Wed close if not $>1\sigma$ adverse, final 1/3 Thursday open
- Exit: full position at exit-band trigger, at next Wednesday's scoring run (no intraweek exits except kill switch or drawdown guard)

6. Table 1 extraction (component C2) — the critical path

The QDS projection lives in a PDF table. Design for robustness, not elegance:

Pipeline: manual upload endpoint (drag the PDF in; do not build scraping of a licensed research portal) → **pdfplumber** extraction targeting the page containing header string "Net investor positioning across tracked commodity markets" → locate columns by header anchors ("Last reported value", "Latest projection", "Projection vs last reported") → row-match on canonical market names (fuzzy match ≥ 0.85 against a fixed alias table: "CBT Corn"→ZC, "ICE Sugar"→SB, etc.).

Validation gates (all must pass or the run flags for manual review, and R defaults to 0 for affected markets):

1. Row count within expected range (30–36 markets)
2. Internal consistency: $\text{projection} - \text{last_reported} \approx \text{published delta column}$ within ± 1 USD mm for every row
3. Sector subtotals sum to within 0.5% of published totals
4. Publication date parsed from header matches expected Tuesday

Fallback: a minimal manual-entry form (six numbers per week for the v1 universe) so a parsing failure never kills the week. Store both raw extraction and validated values in D1 with a **source** flag (**parsed** | **manual**).

Licensing note: extracted values are used internally as model inputs only — no redistribution of JPM data in client-facing outputs, consistent with the report's terms. Client notes reference the framework, not the vendor's numbers.

7. Risk guards

- **Drawdown guard:** if strategy paper P&L drawdown > 6% of allocated NAV, kill switch flips; manual review required to re-arm
 - **Data-staleness guard:** any pillar input > 10 sessions stale → pillar weight redistributed and event logged; if T or R is stale, market is skipped entirely (they are the model)
 - **Fat-finger guard in C5:** reject any order > 2× the largest historical tranche for that market
 - **Event blackout:** no new entries within 24h of WASDE (grains), EIA storage (natgas), FOMC (gold). Exits always permitted.
-

8. Logging & attribution (what makes the beta evaluable)

Every Wednesday run writes to D1 **scores**: all five pillar values, weights actually applied, composite, gate decisions, and the QDS gap raw value — per market, every week, whether or not a trade results. Every order writes intended price, fill, slippage. This enables three attribution cuts at review time:

1. **Signal attribution:** which pillar drove each entry (decompose Score at entry)
 2. **R-pillar isolation:** hypothetical P&L of R-only vs full composite — the direct test of the edge hypothesis
 3. **Gate value:** P&L of trades the extreme-override *blocked* (shadow book) — did the exhaustion brake earn its keep?
-

9. Operating calendar (weekly cycle)

Day	Time (ET)	Action	Component
Fri	15:30	CFTC COT release	—
Fri	16:30	Ingest CFTC	C1 (cron)
Tue	evening	JPM publication → upload PDF	C2 (manual)

Day	Time (ET)	Action	Component
Tue	on upload	Extract + validate Table 1	C2
Wed	08:00	Score all markets, emit targets	C4 (cron)
Wed	09:30	Tranche 1 entry	C5
Wed	close	Tranche 2 (conditional)	C5
Thu	09:30	Tranche 3	C5
Thu	10:30	EIA natgas storage ingest	C3
Daily	17:15	Curve settles ingest	C3

10. Build sequence to August 1

Window	Deliverable	Exit criterion
Jul 11–15	C1 + C3 (public data ingestion), D1 schema	Backfilled 3 years of CFTC + curves; schema frozen
Jul 14–18	C2 extraction + validation harness	Parses the Jul 7 report (this document) and two prior weeks' PDFs with all gates passing
Jul 18–22	C4 scorer + backfill scores 12 weeks	Hand-check 3 weeks of pillar math against source docs
Jul 21–24	C5 + Kelly Sizer integration, IBKR paper hookup	Round-trip a manual test order
Jul 22 (Tue)	Live rehearsal cycle #1	Full Tue→Thu cycle on paper, human-shadowed

Window	Deliverable	Exit criterion
Jul 29 (Tue)	Live rehearsal cycle #2	Cycle runs with zero manual intervention
Aug 1	Beta live	Two clean cycles logged; kill switch tested

Two rehearsal cycles is operational validation, not statistical validation — the statistical case gets made by the attribution logs over the following quarter. Review checkpoint: first Friday of September, using the three attribution cuts in §8.

11. Open decisions (flagged, not blocking)

1. Uniform ± 2 extreme threshold vs per-market thresholds from Fig 18 conventions — v1 uses uniform; revisit after 8 weeks
 2. Whether soybeans stays in the universe (its R and M inputs structurally conflict often; may be a chronic no-trade)
 3. Copper spread (v2) — requires LME COTR data licensing decision
 4. Whether the working paper publishes before or after one quarter of paper track record
-

AlphaQuery Intelligence · Internal — model specification. Not investment advice; paper-trading scope only. QDS/JPM data used as internal model inputs under research license; not for redistribution.