One Engine on a Stationary Axis

A Cointegration-Native, DV01-Tight, Five-Leg Mean-Reversion Architecture for Treasury Yield-Curve-Spread Calibration — Q7 AEGIS AI S2

Haneef R. Haqq, MBA[1]

[1]Quant7 Alpha, LLC

Technical Whitepaper · v1.0 · Strategy S2 (Bond Yield-Curve Spreads)

June 30, 2026

Abstract

We describe S2, the bond-curve member of the Q7 AEGIS AI family, which trades the cointegration of US Treasury futures yield-curve spreads—the dollar-weighted residual between two leg prices—across four canonical spreads (5y/10y, 10y/30y, 5y/30y, 30y/ultra) and four intraday horizons (5, 15, 60, 240 minutes). The design premise is the mirror image of the directional members of the family: a cointegrated curve spread is, by construction, an approximately stationary process, so the four-engine regime partition that powers price-axis strategies *collapses*—trend, breakout, and sweep semantics fire precisely when the cointegration is failing, which is the condition the system must block. S2 is therefore a single mean-reversion specialist. It enters when the residual deviates beyond a calibrated number of standard deviations from its rolling mean, holds while the dislocation persists, and exits as it reverts. Onto that clean z-score signal we port only the framework patterns that are signal-substrate-agnostic: a five-leg adaptive exit indexed on z-score position, a continuous confluence gate, a spread-native four-mode regime router, and a five-state volatility-regime classifier. Sizing is DV01-tight and preventive—positions are scaled to a worst-case stop distance under the fleet-wide drawdown ceiling—and the account trades a single micro contract intraday with end-of-day flatten. A per-configuration, train-equals-serve hybrid-AI gate (gradient-boosted direction, financial-language sentiment, regime) filters entries. We detail the cointegration construction, the engine and exit logic, the sizing and gating layers, the hybrid-AI overlay, the risk cascade, and the calibration acceptance criteria, and we state the architecture's limitations explicitly. No realized-performance claims are made.

1. Edge Thesis

S2 trades the residual of a cointegrating relationship between two Treasury futures, not a directional price. For each spread —`SPRD_5_10`, `SPRD_10_30`, `SPRD_5_30`, `SPRD_30_UB`—the dollar-weighted residual between the long-leg and short-leg contract is approximately stationary on horizons from minutes to weeks, because Federal Reserve policy, supply auctions, and dealer hedging keep curve relationships tethered to fundamentals [1]. When the residual diverges, the same forces that anchor the curve pull it back; the edge is the reversion, and the holding period is the time the dislocation takes to decay. This is a structural property of the instrument, not a fitted pattern: cointegration is a testable relationship [1], and the mean-reverting residual it produces is the canonical statistical-arbitrage signal [5].

The decisive architectural consequence is what does *not* port. The directional members of the Q7 family run four engines that each target a microstructural regime of a trending price. On a stationary spread, three of the four are not merely weak —they are anti-edge. **Trend** on a cointegrated residual fires exactly when the cointegration is breaking, the failure mode the rest of the system exists to detect and block. **Breakout** on a curve spread coincides almost entirely with scheduled macro events (FOMC, CPI, refunding), which are news risk under an explicit blackout, not tradable expansion. **Sweep** requires prior-high/prior-low levels on the traded instrument, but a synthetic spread has no such level—the long-leg high and the short-leg high are different numbers pointing opposite ways, and their combination is not a level any participant leans into. **Mean-reversion** is the only engine whose semantics survive the price-to-spread substitution, and it is already the whole strategy. S2 is therefore a deliberate single-engine specialist; Section 3 develops the non-portability argument in detail.

2. Architecture

S2 is organized in three independently testable layers.

Signal. Per (spread, timeframe), two leg prices are fetched, combined into a DV01-weighted residual, and reduced to a rolling z-score; a cointegration-health check measures rolling-mean drift in sigma units and blocks entries when the relationship is degrading.

Intelligence. Five subsystems run in parallel and write into a confluence aggregator: the Λ-Vol volatility-regime classifier (Section 6); a curve-regime classifier (steepener/flattener states); a cointegration-health score; the hybrid-AI direction model; and the financial-sentiment channel (Section 8).

Execution. Regime router → confluence gate → DV01-tight preventive sizing → entry. Exits run the five-leg adaptive cascade with a break-even arm and an MFE ratchet; any hard-risk event flattens all remaining legs immediately.

Spread	Construction	Horizons	Configs
5y / 10y belly	ZN − 0.55·ZF	5·15·60·240m	4
10y / 30y (NOB)	ZB − 0.45·ZN	5·15·60·240m	4
5y / 30y curve	ZB − 0.25·ZF	5·15·60·240m	4
30y / ultra-long	UB − 0.80·ZB	5·15·60·240m	4

The 16 configurations (4 spreads × 4 horizons) are each calibrated independently. The leg-weight coefficient is the DV01 hedge ratio that renders the residual approximately duration-neutral; it is a property of the curve segment, not a searched parameter.

3. The Mean-Reversion Engine and Cointegration

The signal is constructed in three steps. First, two `request.security` calls fetch the long-leg and short-leg closes (e.g. ZN and ZF for the 5y/10y belly). Second, the DV01-weighted residual is formed,

$$\mathtt{spread}_t = \mathtt{long}_t - \mathtt{w}_{\mathtt{DV01}} \cdot \mathtt{short}_t,$$

where $\mathtt{w}_{\mathtt{DV01}}$ is the per-segment hedge ratio (0.55 / 0.45 / 0.25 / 0.80) chosen to neutralize the bulk of duration risk so the residual reflects curve shape rather than the level of rates. Third, the residual is reduced to a rolling z-score over a per-config window,

$$z_t = (\mathtt{spread}_t - \mu_w) / \sigma_w,$$

and the entry condition is a deviation beyond a calibrated `entry_z` (typically 1.5–3.0 σ), with exit as the residual reverts toward the mean and a catastrophic stop at `stop_z`. This is the Ornstein–Uhlenbeck reversion that statistical-arbitrage practice formalizes on a cointegrated residual [5], specialized to the curve.

Cointegration health. The reversion edge exists only while the cointegrating relationship holds, so S2 runs a continuous health check: the rolling-mean drift of the residual, measured in sigma units over a 60–400-bar window, is compared to a maximum-drift bound, and entries are blocked when drift exceeds it. A persistently drifting residual is, definitionally, a spread whose anchor is failing—the regime in which a naive reversion trade is most dangerous. The health check is both an entry gate and a confluence component (Section 7), and a sustained breach raises a structural defect flag.

Why the other engines are anti-edge. A trend model fit on spread data would train almost entirely on curve-dislocation tails—2008, the 2013 taper tantrum, March 2020, the 2023 regional-banking episode—and would therefore emit maximum-conviction signals at exactly the worst moments for a mean-reverting book, in direct opposition to the cointegration check. A breakout model fires on range expansion, which on curve spreads is dominated by scheduled events already removed by the Fed blackout; outside the blackout it fires during organic vol expansion, the regime in which the spread is reverting and a breakout trade is counter-trend. A sweep model has no defined level to sweep. Forcing the four-engine port would deliver negative edge; S2 keeps the single MR engine and ports only the substrate-agnostic framework patterns of Sections 4–7.

4. The Five-Leg Adaptive Exit

The framework's largest contribution to the reversion signal is replacing a binary z-revert (close all when the residual returns to `±exit_z`) with a graduated five-stage profit-capture sequence indexed on z-score position. This matters because a mean-reversion trade rarely has a single optimal exit point: the residual decays through a corridor, and capturing that decay in stages dominates an all-or-nothing close.

Let `span = |entry_z| − |exit_z|` be the distance from the entry extreme back to the exit band on the same side of zero. For a long-spread setup (entered at `z ≤ −entry_z`), the four scale-out legs fire at

$$\mathtt{leg}_k\mathtt{\_z} = -\mathtt{entry\_z} + \mathtt{span} \cdot f_k, \quad f = (0.30, 0.55, 0.80, 1.00),$$

with the fourth leg at the mean and a fifth *runner* leg that captures overshoot a calibrated distance (default 0.30 z) past the mean; the short-spread convention is mirrored. The legs close 30 / 25 / 20 / 15 / 10 percent of the position respectively, the first time the residual reaches each target on the correct side. The allocation front-loads confirmation (the first leg pays for the trade) while keeping the majority of size working into the median and full-reversion zones, and reserves a tail for the overshoot that mean-reversion frequently produces.

Two protective mechanisms ride the cascade. After leg 1 fires, the stop moves to the entry spread (break-even); after leg 2, an MFE ratchet advances the stop toward the mean by a calibrated fraction (default 0.50) of the remaining distance, locking accumulated P&L while leaving the back legs room to develop. The leg machinery operates only on the normal reversion path: any catastrophic exit—hard stop, time stop, or a fired risk halt—closes all remaining legs immediately with the corresponding reason code.

5. Sizing and the Curve

S2 sizes preventively from the dollar distance to the catastrophic stop rather than from a fixed lot. The risk target is a calibrated fraction of equity, bounded by a hard per-trade dollar cap, and the unit count follows from the stop distance:

$$\text{risk\$} = \min(\ eq\cdot dv01\_pct,\ eq\cdot hardStopPct\ ),$$

$$n_{long} = \max(\ 1,\ \min(\ max\_units,\ \text{risk\$}\ /\ stop\_dist\$\ )\ ).$$

The `min` against the hard cap is essential: it bounds per-trade dollar risk regardless of how tight the stop is, defeating the very-tight-stop pathology in which a small stop distance sizes the position up beyond catastrophic recoverability. The short leg is then sized by the DV01 hedge ratio, deliberately underhedged by a small cap multiplier (default 0.95) so the position retains positive expected drift in the direction of the reversion while neutralizing the bulk of outright rate risk. Because S2 runs a micro-account mandate—a single contract, intraday, with end-of-day flatten so no curve position is held overnight—the sizing layer operates against a small, fully specified risk budget, and commissions are modeled cash-per-contract on the venue's round-trip schedule.

6. Regime Routing and the Λ-Vol Classifier

S2 carries two regime layers that condition the reversion engine without ever overriding it. The **Λ-Vol classifier** is a deterministic five-state machine driven by the short-versus-long volatility ratio of the residual and its mean drift in sigma units:

$$\text{ratio} = \sigma_{short} / \sigma_{long}, \quad \text{drift} = |\mu_{short} - \mu_{long}| / \sigma_{long}.$$

In priority order it classifies `STRESS_HIGH` (ratio above a high threshold → hard block), `TREND_ACTIVE` (elevated ratio and drift → reversion weakened), `TREND_CALM` (low ratio and drift → fast clean reverts), `TRANSITION` (an ambiguous band that soft-weights into confluence), and `MEAN_REVERT` (the favorable default). The motivation for an explicit volatility-state layer over a single static geometry is the same one that motivates regime conditioning generally—the residual's reversion speed and reliability are themselves regime-dependent, and the documented dynamics of realized-volatility surfaces show level, slope, and convexity modes that a one-size exit serves poorly [3].

A **spread-native router** consumes the Λ-Vol state plus cointegration health and selects one of four modes—`MR_NORMAL`, `MR_OVERSHOOT`, `MR_TIGHT`, `MR_DEFENSIVE`—each applying multipliers to the base `entry_z`, `exit_z`, `stop_z`, and timestop. `MR_OVERSHOOT` demands a deeper extreme and grants wider stops when vol is active (deeper extremes revert more reliably); `MR_TIGHT` catches faster reverts and locks in quicker when vol is calm; `MR_DEFENSIVE` blocks new entries entirely and accelerates the exit of existing positions under stress or unhealthy cointegration. A separate curve-regime classifier (bull/bear steepener/flattener) acts as a directional gate—blocking long-spread entries in a bear-flattener and short-spread entries in a bull-steepener—and contributes a confluence component.

7. Confluence Scoring and Gating

S2 replaces a stack of binary boolean filters with a single continuous confluence score, so that losing one component partially does not kill a trade the others support. Six components are each computed on a 0–100 scale and combined by calibrated weights summing to one:

$$\text{conf} = \Sigma_i\ w_i \cdot \text{score}_i, \quad \Sigma_i\ w_i = 1.$$

Component	Reads	Weight
z-score	min(100, |z|/entry_z × 100)	0.30
HAI	direction probability × 100	0.20
curve regime	aligned with trade side	0.15
cointegration	1 − drift/max_drift	0.15
slippage	1 − est_slip/budget	0.10
Λ-Vol	state favorability	0.10

The entry gate fires when `conf ≥ confMin_effective`, where the per-config floor carries a small directional adjustment—short-spread entries require slightly higher confluence to compensate for the funding-cost asymmetry of bond shorts—under a hard feasibility cap that prevents any Optuna trial from driving the gate into infeasibility. The slippage component is illustrative of the design: a marginal pre-trade slippage estimate (intra-bar range plus a vol term, compared to a per-leg basis-point budget [4]) erodes confluence weight rather than hard-blocking, so the gate degrades smoothly instead of cliff-edging.

8. The Hybrid-AI Gate (Train = Serve)

S2 filters rule-based reversion entries through the same hybrid-AI overlay as the rest of the family—a fusion of a supervised directional learner, a financial-language sentiment model, and a market-regime classifier—in the spirit of the multimodal architecture of [7], adapted to per-configuration, intraday, train-equals-serve operation. The motivation is the one that paper documents empirically: rule-based logic is blind to news shocks and degrades through regime transitions, and a fusion of structured ML with sentiment and regime awareness is materially more robust than any isolated predictor.

8.1 Directional model — gradient-boosted trees (XGBoost)

The directional learner is an XGBoost classifier that, per spread × timeframe, outputs the probability of a successful revert given the current state. The reference design [7] frames direction as a binary next-period problem over a compact technical feature set with shallow trees and stagewise boosting; S2 adapts that recipe to the spread axis, where the natural features are *reversion-state* features: the z-score and its magnitude, z velocity and acceleration, the short/long volatility ratio and drift, time-of-day and day-of-week, and the prior-day spread direction. Mechanically, XGBoost is an additive ensemble of shallow regression trees fit by gradient boosting on a regularized logistic objective: each new tree fits the negative gradient of the loss given the current ensemble, scaled by a shrinkage rate, with row and column subsampling and L1/L2 penalties on leaf weights; the ensemble outputs a calibrated probability that becomes the `hai_score` confluence component. Two disciplines distinguish S2's use from the reference. First, the model is trained *in-cycle* on the current cycle's own accumulated trades with strict walk-forward, no-look-ahead labeling—no stale prior-cycle model is ever served. Second, it is admitted only if it clears a per-configuration discriminability floor (a minimum cross-validated AUC) and a leakage bound; weak configurations are flagged for retraining rather than deployed on a coin-flip probability. The entry threshold is itself a calibratable parameter, and the gate is calibratable on/off—if the optimizer selects it off, the rule-based reversion stack carries the decision.

8.2 Sentiment model — FinBERT

The sentiment channel is FinBERT [8], a BERT-family transformer fine-tuned on financial text. Each item is tokenized and contextualized through stacked self-attention layers and mapped by a classification head to probabilities over {positive, neutral, negative}; domain fine-tuning is what lets it parse the hedging and central-bank jargon that generic sentiment models miss. For S2 the relevant polarity is the *hawkish–dovish* axis of monetary-policy language, aggregated into a normalized score

$$S_t = (1/N)\ \Sigma_i\ (\ P_{dovish}^{(i)} - P_{hawkish}^{(i)}\ ),\quad S_t \in [-1, 1].$$

Following [7], S2 consumes S_t as a *risk filter and soft threshold skew, never as a standalone entry signal*: a small polarity adjustment skews the HAI threshold, while a sufficiently hawkish reading (below a calibrated floor) hard-blocks long-spread entries, since an overtly hawkish macro regime is hostile to the steepening that a long-spread reversion is betting on. *Honesty note:* in a historical backtest without a time-aligned, point-in-time policy-text feed for the evaluated window, this channel is run neutral under a signed no-feed exemption (honest-empty) rather than fabricating sentiment; the exemption is recorded in the cycle dossier, and train=serve is preserved by deploying the same neutral configuration that was calibrated.

8.3 The tandem and the regime channel

The third model is the market-regime classifier—the Λ-Vol router (Section 6) at the strategy level and a regime label at the overlay level—which conditions both the gate and the exit geometry. The three combine as a tandem: the directional model proposes a revert probability, the sentiment channel can veto or attenuate it, and the regime label selects the operating envelope (normal / overshoot / tight / defensive). This is the multi-modal fusion that [7] shows outperforms any single isolated predictor across changing regimes. The decisive engineering choice is that all three are computed in-cycle from the same data the system trades—the overlay calibrated is the overlay deployed, so the gate cannot be a flattering in-sample artifact of a model the live system never runs.

9. Risk, Calibration, and Acceptance

Risk. Risk is enforced preventively at the sizing layer (Section 5): positions are scaled so that a worst-case adverse move from the monotonically ratcheting equity peak cannot carry the account past the hard drawdown ceiling, a structural constant identical across the specification, the calibrator, and the executor. Behind that preventive primary sits a defense-in-depth halt cascade—a daily-loss cap, a rolling-drawdown halt off the high-water mark, a kill switch, and a per-trade dollar cap—each independently togglable and each emitting a coded structural-defect signal when breached, plus a post-loss cooldown that lengthens after stop-outs to break revenge-trade clusters. A Fed event blackout blocks new entries within a window around FOMC, CPI, and first-Friday NFP; existing positions are *not* force-flattened at the boundary, since the catastrophic-stop logic handles event gaps more efficiently than a pre-event flatten that would lock in unrealized P&L at suboptimal points.

Calibration. Each of the 16 configurations gets an independent TPE Bayesian search over eight parameters—`entry_z`, `exit_z`, `stop_z`, time-stop bars, z-window, risk percentage, max units, and `confMin`—warm-started from the prior cycle's best trials and rebuilt from a clean baseline each cycle (never patched). The five-leg allocations and z-revert fractions, the confluence weights, and the Λ-Vol thresholds are held global at this stage to keep the search tractable. Critically, the objective uses *flat hard gates*, not negative penalties: a configuration that violates the commission, trade-count, drawdown, or net-result bound returns a flat zero rather than a negative score, because penalty gradients in the negative region lead the sampler to *explore* the violating region to escape, whereas a flat rejection has no gradient and the search avoids it entirely.

Acceptance. A configuration is accepted only on a simultaneous gate—win rate, gross profit factor, trade frequency per timeframe, a commission-to-gross bound (a hard block), and a drawdown bound—evaluated on the hybrid-AI-gated system, with the HAI discriminability floor enforced separately. Cycle eligibility additionally requires aggregate positive result, no fatal structural-loss or commission-spiral flags, and in-cycle model freshness. Promotion toward live deployment requires consecutive consistent cycles, Pine↔Python parity within tolerance (trade count within a few percent, net result within a tight band, drawdown timing aligned), TradingView treated as the authoritative backtest, and a paper-trading fill-rate check —the governance described in the program's pipeline whitepaper. Every threshold in this section is a calibration *gate*, not a reported outcome.

10. Limitations

The cointegration that grounds the edge is a statistical relationship that can and does break; the health check detects degradation but cannot anticipate a structural regime change in curve dynamics, and the architecture treats a sustained drift breach as a block, not a tradable signal. The DV01 hedge ratios that render the residual duration-neutral are segment-level constants and approximate the true time-varying hedge; under a sharp convexity or financing shift the residual carries more outright rate risk than the static weight implies. The volatility-spillover literature indicates that curve legs co-transmit volatility through correlated shocks [2], and a learned dual-adjacency regime view is the principled upgrade to the rule-based Λ-Vol and curve-regime layers—a roadmap item, not a present claim; the same literature is explicit that such spillover structure improves regime and risk classification but *not* return prediction, so it is deliberately excluded from the direction model. The sentiment channel depends on a point-in-time policy-text feed that is run neutral when unavailable. The hybrid-AI gate constrains but does not certify causal or durable structure, and execution, liquidity, financing, and venue/regulatory risks are not eliminated by architectural integrity. As with every Q7 deliverable, no realized-performance representation is made, and the acceptance criteria above are calibration gates, not outcomes.

References

[1] R. F. Engle, C. W. J. Granger. *Co-integration and Error Correction: Representation, Estimation, and Testing.* Econometrica, 1987.

[2] B. Mallory. *Two-Step Regularized HARX for Volatility Spillovers.* 2026.

[3] R. Cont, J. da Fonseca. *Dynamics of Implied Volatility Surfaces.* Quantitative Finance, 2002.

[4] M. Gould et al. *Limit Order Books* (survey). Quantitative Finance, 2013.

[5] M. Avellaneda, J.-H. Lee. *Statistical Arbitrage in the U.S. Equities Market.* Quantitative Finance, 2010.

[6] *Follow the Leader: Enhancing Systematic Trend-Following Using Network Momentum.* 2024.

[7] V. N. K. Pillai, A. Ajith, S. K. J. *Generating Alpha: A Hybrid AI-Driven Trading System Integrating Technical Analysis, Machine Learning and Financial Sentiment for Regime-Adaptive Strategies.* ComSIA 2026 (Springer LNNS).

[8] D. Araci. *FinBERT: Financial Sentiment Analysis with Pre-trained Language Models.* 2019.
