One Position, One Dollar Regime

A DXY-Proxy-Gated, Single-Position Architecture for
FX-Futures Calibration — Q7 AEGIS AI S6

Haneef R. Haqq, MBA[1]

[1]Quant7 Alpha, LLC

Technical Whitepaper · v1.0 · Strategy S6 (FX Futures)

June 30, 2026

Abstract

We describe S6, the FX-futures member of the Q7 AEGIS AI family, which trades CME micro currency futures on four major USD pairs (6E/EUR, 6B/GBP, 6A/AUD, 6J/JPY) across four horizons (15 minutes, 1 hour, 4 hours, daily). The design premise is that the dominant systematic driver of the major-pair complex is the *U.S. dollar regime* itself: when the dollar trends, the majors co-move against it with high cross-sectional coherence, and when it is range-bound, intraday flow reverts to session value. S6 therefore runs *two* engines rather than a regime menu—a dollar-regime trend engine that is gated by a four-major weighted-log dollar proxy, and a session mean-reversion engine that fades session-VWAP extremes—and it holds a *single* position with a single adaptive exit, deliberately rejecting the multi-leg ladder used by the higher-frequency Q7 strategies. Position sizing is DV01-aware and equal-vol, floored by a hard per-trade dollar cap. The same per-configuration, train-equals-serve hybrid-AI gate used by the reference architecture filters entries. We detail the dollar-proxy construction, the two engine semantics, the single-position adaptive exit, the DV01 sizing and risk construction, and the calibration acceptance criteria, and we state the architecture's limitations explicitly. We also record candidly that S6 was re-ported to this faithful golden after an earlier over-built variant was found and rejected (Section 11). No realized-performance claims are made; the golden documents an approximate 11-of-16 design ceiling, which we present as a design *target*, not a result.

1. Edge Thesis

S6 trades the major USD currency-futures complex, where the canonical edge is not a menu of idiosyncratic per-pair patterns but the state of the dollar itself. The majors are, to first order, the reciprocal of a single latent factor—broad USD strength—so a strategy that reads that factor directly and trades the pairs as its expression captures a more stable signal than one that treats each pair as an independent instrument. S6 expresses this with two complementary engines. **FX Regime (FX_REGIME)**, the H1/H4/Daily workhorse, fires when a four-major dollar proxy is in a strong directional regime *and* the traded bar agrees with the USD-strong direction—trend persistence in the dollar is the edge condition. **Session Mean-Reversion (SESSION_MR)**, the M15 engine, fires when intraday price has stretched far from session value and the dollar regime is *not* strongly trending—range-bound dollar conditions are where session flow reverts. The two engines partition the dollar's regime space into trend and range; the proxy and the volatility context arbitrate which is admissible. This is the structural contrast with single-engine FX specialists, which carry one geometry into both states.

2. Architecture

S6 is organized in three independently testable layers.

Signal. Per (pair, timeframe), the active engine for that horizon computes an entry candidate each bar and writes a candidate direction plus an engine tag. FX_REGIME is the default engine on H1/H4/Daily; SESSION_MR is the default on M15. Arbitration admits the candidate only if its engine's gate clears, subject to a per-engine cooldown.

Intelligence. Three subsystems run in parallel: the dollar-proxy regime layer (Section 3); the single-position adaptive exit planner (Section 6); and the hybrid-AI overlay (Section 7).

Execution. Engine gate → regime admissibility → cooldown → hybrid-AI gate → DV01-aware equal-vol sizing under a hard dollar cap → entry. The position then runs a *single* adaptive exit—an ATR trail-from-entry ratchet, a time stop, and a regime-flip close—with no leg ladder.

Future	Pair	Horizons	Configs
6E	EUR/USD	M15·H1·H4·D	4
6B	GBP/USD	M15·H1·H4·D	4
6A	AUD/USD	M15·H1·H4·D	4
6J	JPY/USD	M15·H1·H4·D	4

The 16 configurations (4 pairs × 4 horizons) are each calibrated independently. Contract economics use each micro future's `POINT_VALUE` (the CME micro multiplier per pair), and commissions use a flat per-contract NinjaTrader schedule rather than a percentage rate—the futures analogue of the venue economics carried by the rest of the fleet.

3. The Dollar-Proxy Regime Layer

The organizing signal of S6 is a synthetic broad-dollar index built from the four majors themselves, so that the regime read is available on the same data the system trades and requires no external index feed. The proxy is a weighted sum of log prices,

$$\mathrm{dxy} = -0.595\cdot\ln(\mathrm{EUR}) - 0.120\cdot\ln(\mathrm{GBP}) - 0.139\cdot\ln(\mathrm{AUD}) + 0.146\cdot\ln(\mathrm{JPY}),$$

standardized over a rolling window to a z-score `dxy_z = (dxy − μ)/σ`. The negative loadings on EUR, GBP and AUD and the positive loading on JPY reflect quoting convention—EUR/USD, GBP/USD and AUD/USD rise when the dollar weakens, while the JPY future (USD-base) rises when the dollar strengthens—so a positive `dxy_z` is an unambiguous USD-strong reading across the basket, and the magnitudes approximate the major-currency weights of the conventional dollar index. `dxy_z` drives a three-way regime label: *USD-strong* when it sits above a calibrated positive band, *USD-weak* below the symmetric negative band, and *neutral* in between. FX_REGIME is admissible only in the two strong states and only when the traded bar's direction agrees with the dollar-implied direction for that pair; SESSION_MR is admissible only in the neutral state. The construction is grounded in the dollar-index and network-momentum literature [4, 5], which documents that a single broad-dollar factor carries the majority of cross-pair co-movement and that a leader/aggregate signal supplies durable directional information—the same premise, instantiated here as a deterministic weighted-log proxy rather than a learned lead-lag graph.

4. The FX-Regime Engine

FX_REGIME is the directional workhorse on H1, H4 and Daily. It enters in the direction implied by the dollar regime—short the EUR/GBP/AUD futures and long the JPY future when `dxy_z` is USD-strong, and the reverse when USD-weak—conditioned on the traded bar confirming that direction, so the engine never fights the prevailing dollar trend. Because the proxy is computed from the basket, the engine reads a market-wide state rather than a single pair's noise, which is what gives the trend signal its stability across the four configurations that share a horizon. The engine carries a calibrated regime-band threshold (how strong `dxy_z` must be to admit a trade) and a per-engine cooldown; it does not size by a fixed multiplier (Section 5).

5. The Session Mean-Reversion Engine and DV01 Sizing

SESSION_MR is the M15 engine. It computes a session VWAP and the standardized distance of price from it—a session-VWAP-distance z-score—and fades extremes: a sufficiently large positive deviation is a short candidate and a sufficiently large negative deviation a long candidate, admitted only when the dollar regime is neutral so that the engine reverts intraday flow rather than standing in front of a dollar trend. The reversion premise is the standard intraday one—session value acts as an anchor that stretched price returns toward when no broad directional regime is driving the basket.

Both engines size identically and *not* by a fixed leverage multiplier. S6 uses DV01-aware equal-vol sizing: contracts are chosen so that each position carries a comparable dollar risk per unit of price movement—a risk-percent or DV01 target translated through the pair's `POINT_VALUE` and current ATR into a contract count—so that a high-volatility pair and a low-volatility pair contribute equal risk rather than equal notional. The computed size is then floored by a *hard per-trade dollar cap*, a structural ceiling that the optimizer cannot exceed, which bounds single-trade loss independently of the volatility estimate. Equal-vol normalization across the basket is the same risk-parity discipline the network-momentum literature applies to multi-asset trend portfolios [5].

6. Single-Position Adaptive Exit

S6 holds *one* position per configuration and exits it with a single adaptive mechanism—there is no four-leg ladder, no scale-out cascade, and no runner. Three rules govern the exit, and the first to trigger closes the position. (i) A **trail-from-entry ATR ratchet** anchored on a Wilder ATR: the stop trails the favorable excursion at a calibrated ATR multiple from entry, ratcheting only in the favorable direction so realized gains are protected without giving the trade a fixed target. (ii) A **time stop**: a maximum holding duration in bars, after which the position is closed regardless of excursion, so a stalled trade does not tie up risk budget. (iii) A **regime-flip close**: when the dollar regime turns neutral or flips against the open position, the position is closed—*flat, not reversed*. The regime-flip rule is the engine's honesty mechanism: S6 trades a dollar regime, so when that regime is no longer present the thesis for the trade has expired and the position is exited rather than carried on inertia or flipped into an un-validated counter-trade. Any hard-risk event (drawdown halt, per-trade cap, macro gate) closes the position immediately on the same single path.

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

S6 filters rule-based entries through the same hybrid-AI overlay as the reference architecture—a fusion of three models, a supervised directional learner, a financial-language sentiment model, and a market-regime classifier—in the spirit of the multi-modal architecture of [9], adapted to per-configuration, intraday, train-equals-serve operation, with a per-engine entry threshold. The motivation is the one that paper documents empirically: trend and mean-reversion logic fail through regime transitions and are blind to news shocks, and a fusion of structured ML with sentiment and regime awareness is materially more robust than any isolated predictor. The gate is applied independently to FX_REGIME and SESSION_MR entries with its own calibrated threshold per engine.

7.1 Directional model — gradient-boosted trees (XGBoost)

The directional learner is an XGBoost classifier that predicts whether the forward move is favorable. The reference design [9] frames this as a binary next-period-direction problem (label 1 if the next-period return is positive) over a compact technical feature set, with ~200 boosting rounds, tree depth 6, and a 0.05 learning rate on a 70/30 split. S6 adapts that recipe to its FX, per-(pair, horizon) setting: labels are the forward-return sign thresholded at a small favorable/unfavorable band, and the feature set is *direction-conditional*—separate long and short framings (roughly a dozen technical and dollar-regime features, including `dxy_z`, the session-VWAP distance, ATR context and an is-long indicator) are computed so one model scores both sides. Mechanically, XGBoost is an additive ensemble of shallow regression trees fit by stagewise 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/column subsampling and L1/L2 leaf-weight penalties; the ensemble outputs a calibrated probability. Two disciplines distinguish S6's use from the reference: the model is trained *in-cycle* on the current cycle's own accumulated trades (no stale prior-cycle model), and it is admitted only if it clears a per-direction discriminability floor (AUC $\geq$ 0.55 long *and* short) and an adversarial leakage bound (adversarial AUC < 0.85). Its entry threshold is itself a calibratable parameter, and the gate is calibratable on/off—if the optimizer selects it off, the rule-based stack carries the decision.

7.2 Sentiment model — FinBERT

The sentiment channel is FinBERT [10], a BERT-family transformer fine-tuned on financial text. Each news 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 negation, hedging, and finance-specific jargon that generic sentiment models miss. Per pair per day the article-level scores are aggregated into a normalized polarity,

$$S_t = (1/N)\ \Sigma_i \left(P_{positive}^{(i)} - P_{negative}^{(i)} \right), \quad S_t \in [-1, 1].$$

Following [9], S6 consumes S_t as a *risk filter and soft threshold skew, never as a standalone entry signal*: a sufficiently negative S_t (below a calibrated floor) suspends new entries and can flatten exposure, shielding positions against adverse-news shocks—central-bank surprises, rate decisions, macro releases—that purely technical logic is blind to. *Honesty note:* in a historical backtest without a time-aligned, point-in-time news 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.

7.3 The tandem and the regime channel

The third model is the market-regime classifier—the dollar-proxy regime layer (Section 3) at the strategy level, and a regime label at the overlay level—which conditions both the gate and the exit. The three combine as a tandem: the directional model proposes a probability, the sentiment channel can veto or attenuate it, and the regime label selects whether FX_REGIME or SESSION_MR is even admissible. This is the multi-modal fusion that [9] shows outperforms any single isolated predictive model 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.

8. Calibration and Acceptance

Each configuration is accepted into a cycle only on a simultaneous five-criterion gate—win rate, gross profit factor, trade frequency, drawdown (a soft, logged criterion), and a commission-to-gross-profit bound (a hard block)—evaluated on the hybrid-AI-gated system, with the hybrid-AI discriminability bounds enforced separately. The commission bound is material for S6 because the flat per-contract NinjaTrader fee is a fixed cost per round trip rather than a notional-proportional one, so low-edge high-frequency configurations are screened out by construction. Cycle eligibility additionally requires full coverage with no blocked configuration, positive aggregate result, and in-cycle model freshness. Promotion toward live deployment requires consecutive consistent cycles, specification-to-implementation parity within tolerance, and an independent re-derivation—the governance described in the program's pipeline whitepaper.

9. Risk and Venue

Risk is enforced preventively at the sizing layer: positions are sized so that a worst-case adverse move, measured from the monotonically ratcheting equity peak, cannot carry the account past the hard drawdown ceiling, which is the fleet-wide structural constant—a **24% ceiling**—identical across the specification, the calibrator, and the executor. Because a reactive kill cannot stop an overnight or weekend gap, the ceiling is held by *size-to-headroom*: as the account draws down toward the ceiling, position size is reduced so that the remaining headroom cannot be breached by a single adverse move. The DV01-aware equal-vol sizer and the hard per-trade dollar cap (Section 5) are the instruments of this preventive control. The execution venue is CME micro FX futures via NinjaTrader; contract economics use each pair's `POINT_VALUE` multiplier, commissions use the flat per-contract schedule with a conservative round-trip assumption, and overnight/weekend exposure is governed by the time stop and the regime-flip close rather than carried blindly.

10. Limitations

The dollar proxy is a fixed-weight, deterministic instrument: its loadings approximate the major-currency dollar weights but are not re-estimated online, and the network-momentum literature [4, 5] indicates that a learned cross-asset lead-lag graph would supply orthogonal alpha and is the principled replacement for the fixed-weight proxy—this is a roadmap item, not a present claim. The proxy reads only the four majors S6 trades, so a dollar move driven by pairs outside the basket is seen only through its projection onto these four. The regime classifier is deterministic and can saturate in a persistent trend; the hybrid-AI gate constrains but does not certify causal or durable structure; and execution, liquidity, gap, and venue/regulatory risks are not eliminated by architectural integrity. The golden documents an approximate **11-of-16** design ceiling—we present that as the design *target* the architecture is built to reach, not a realized result. As with every Q7 deliverable, no realized-performance representation is made, and the acceptance criteria above are calibration gates, not outcomes.

11. Integrity Note: Re-Port to the Faithful Golden

We record candidly that the S6 architecture described here is a deliberate re-port. An earlier S6 variant was over-built relative to its golden specification—it carried a bespoke own-pair signal, a cloned multi-leg exit ladder imported from the higher-frequency strategies, and additional overlay machinery that the golden never sanctioned. That variant was a stripped-and-grown shadow rather than a faithful port, and it was rejected on review. The faithful golden is the single-position RX01 design documented above: two engines (FX_REGIME and SESSION_MR), the dollar-proxy regime layer, one position with one adaptive exit, and DV01-aware sizing under a hard cap. We surface this because architectural integrity in the Q7 program means the implementation faithfully realizes its golden specification—no stripped controls, no un-goldened additions—and a candid account of a rejected over-build is part of that standard, not an exception to it.

References

[1] R. Cont, A. Kukanov, S. Stoikov. *The Price Impact of Order Book Events.* 2011.

[2] R. Cont, A. de Larrard. *Price Dynamics in a Markovian Limit Order Market.* 2011.

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

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

[5] *Network Momentum across Asset Classes.*

[6] *A Primer on the Signature Method in Machine Learning.*

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

[8] *Quantum-Probability-Theoretic Asset Return Modeling* (multi-modal return distributions at high-volume regimes). 2024–2025.

[9] 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 Equity Strategies.* ComSIA 2026 (Springer LNNS).

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