For Alpha

Ai-Powered Investment Replication

Strategy Spotlight: Double CTA Decoding with Dynamic Multi-Horizon Setting

November 2025, Ai For Alpha Team

Abstract

We introduce a two-stage *Double CTA Decoding* for replication and attribution. Stage 1 decomposes the SG CTA Trend index into five horizon-controlled trend sleeves (20d, 60d, 125d, 250d, 500d) plus a raw market-return sleeve. Stage 2 projects the benchmark onto these six sleeves to form a *dynamic, interpretable factor mix* that is directly investable.

Across evaluation windows (Since Inception, Last 10Y, Last 5Y), **Double Decoding strikes the best balance of fidelity and efficiency**: it sustains near-benchmark correlation while stabilizing Sharpe and Return/MaxDD relative to any single sleeve. For context, single-sleeve monthly correlations to SG CTA Trend since inception are 0.84 (No-Medium-Term), 0.81 (250d), 0.80 (Augmented Short-Term), and 0.78 (125d).[1] Efficiency leaders by window are 500d in the last 5Y (Sharpe 0.75, Return/MaxDD 0.83), Augmented Short-Term in the last 10Y (Sharpe 0.54), and No-Medium-Term since inception (Sharpe 0.81). Stage 2 blends these strengths through a time-varying horizon mix.

Methodologically, all weights are *time-varying* and estimated with a Bayesian graphical model under sequential (no look-ahead) information, ensuring stability and interpretable horizon attribution.

[1] Monthly correlations from inception; see Table 6.

Background and Motivation

CTA replication is typically approached either bottom-up (market level) or top-down (factor level). Our earlier notes showed the value of horizon structuring and the cost of over-weighting medium-term bands:

- *Medium-Term Trend Replication (Spotlight)*
- *Multi-Horizon CTA Replication (Spotlight)*

This paper merges those lines into a two-stage *Double CTA Decoding* with *Dynamic Multi-Horizon* weights that evolve over time and remain directly interpretable. See also other related work on graphical decoding and trend-factor geometry (Benhamou et al., 2025; Etienne et al., 2025; Ohana et al., 2024; Benhamou et al., 2024).

Investment Universe and Cost Framework

Our tests run on a liquid, diversified futures universe that matches standard CTA exposures across equities, rates, FX, and commodities. Implementation costs are applied consistently across all strategies (single sleeves, blends, and the Stage B Horizon Decoding replicator). We use three layers of costs; contract-level inputs appear in Table 1.

- **Transaction cost (Tx. cost).** Round-turn execution expense that bundles bid–ask, brokerage, exchange and clearing fees, plus a small slippage buffer. Quoted in basis points of notional per round-turn and applied on filled trades. Orders are *netted across sleeves* at the instrument level (one ticket per market per day).
- **Replication (roll) cost.** Systematic carry/roll drag when the front contract is rolled to the next maturity. Estimated as the 2005–2025 average front-to-next calendar spread for each contract and applied on scheduled rolls.
- **Management fee.** Flat 50 bps per annum on AUM. Where noted, results can be shown gross or net of this fee.

Asset class	Costs (Tx, Roll)	Instruments (exchange)
Commodities	2 / 15 bps	GC (COMEX); CL, NG (NYMEX); CO (ICE Europe); HG (COMEX)
Equity Indices	2 / 15 bps	ES, NQ (CME); NK (OSE); FESX (Eurex); Z (ICE Europe); E-mini EM (CME)
Fixed Income (Rates)	2 / 10 bps	TU, TY (CBOT); RX (Eurex); G (ICE UK Gilts); JGB (OSE); XM (ASX)
FX (vs USD)	2 / 2 bps	EUR, JPY, GBP, AUD, CAD (CME)

Table 1: Futures universe and cost assumptions (Tx = round-turn transaction cost; Roll = average front-to-next spread). Costs are expressed in basis points of notional and applied consistently across Stage A sleeves and the Stage B Horizon blend.

Implementation notes. (i) All sleeves are volatility-targeted; costs are applied to *filled* notional after sleeve netting. (ii) Roll timing follows exchange calendars; roll costs are applied only when an actual calendar switch occurs. (iii) Because orders are netted per instrument across factor sleeves (20/60/125/250/500/MKT) *before* routing, reported turnover and costs reflect the *true* implementation of Stage A *and* the Stage B Horizon Decoding.

Look-Back Straddles as Horizon-Controlled Trend Factors

A look-back option pays on the *path* of prices, not just the close. Holding a call *and* a put—the **look-back straddle**—accumulates value whenever price makes a large up *or* down excursion over a window h. The straddle's *delta* is positive near recent highs, negative near recent lows, and close to zero in ranges—precisely the sizing logic of breakout CTAs.

From path to factor. For market i with price $P_{t,i}$ and excess return $m_{t,i}$, define running extremes on window h:

$$H_{t,i}^{(h)} = \max_{1\le\ell\le h} P_{t-\ell,i}, \qquad L_{t,i}^{(h)} = \min_{1\le\ell\le h} P_{t-\ell,i},$$

and a robust scale (e.g. ATR or rolling stdev) $S_{t,i}^{(h)}$. A simple trend score that mimics the straddle delta is

$$s_{t,i}^{(h)} = \frac{P_{t-1,i} - \frac{1}{2}(H_{t-1,i}^{(h)} + L_{t-1,i}^{(h)})}{S_{t-1,i}^{(h)}}.$$

We map $s^{(h)}$ to a bounded position $\pi_{t,i}^{(h)} \in [-1,1]$ via a saturating link (e.g. $\pi = \tanh(\gamma s)$ or hard sign with banding). The *per-market trend factor return* is then

$$g_{t,i}^{(h)} = \pi_{t-1,i}^{(h)} \, m_{t,i},$$

rescaled to a common volatility target per horizon to make factors comparable across markets. Intuitively: long (short) into strength (weakness), flat in noise, convex in large moves.

The six canonical factors. We use five horizon-controlled trend factors and one non-trend baseline:

$$\mathcal{H} = \{20, 60, 125, 250, 500\}, \qquad \text{and} \quad \text{MKT}: \ g_{t,i}^{(\text{MKT})} \equiv m_{t,i}.$$

- **TF-20 / TF-60**: fast sleeves (1–3 months); react early, diversify reversals.
- **TF-125 / TF-250**: medium sleeves; classical CTA core but potentially crowded.
- **TF-500**: slow sleeve; persistent macro trends, shallow drawdowns.
- **MKT**: raw market-return sleeve; "no-trend" baseline.

Table 2: Look-back straddle mechanics and CTA analogues

Concept	**Intuition**	**CTA Analogue**
Path-dependent payoff	Rewards *max/min* excursions over h.	Trend P&L depends on the journey, not just the close.
Straddle symmetry	Long call+put captures both tails.	Long in up-trends, short in down-trends.
Delta as trend score	$+$ near highs, $-$ near lows, ≈ 0 in ranges.	Breakout sizing by distance to extremes.
Window h	Sets reaction speed and holding horizon.	Diversify horizon risk (20,60,125,250,500).
Convexity	Positive skew from large moves.	CTAs exhibit crisis convexity ("crisis alpha").

The result is a clean library of *per-market* factor returns $\{g_{t,i}^{(h)}\}_{h\in\mathcal{H}}$ plus the market-return baseline $g_{t,i}^{(\text{MKT})}$, ready to be fed into the decoding stages below.

Methodology

Overview

We implement a two-layer, fully *sequential* (no look-ahead) state-space pipeline that maps the SG CTA Trend daily excess return r_t^{Bench} into:

1. **Stage A — Six single-factor decoders (run in parallel):** for each factor $k \in \{20, 60, 125, 250, 500, \text{MKT}\}$ we regress the benchmark on the *per-market* realizations of that factor to obtain a time-varying cross-sectional mix and a factor-specific replication $\widehat{r}_t^{(k)}$.
2. **Stage B — Multi-horizon blend:** we regress the benchmark on the six replications $\{\widehat{r}_t^{(k)}\}_k$ to obtain a dynamic, interpretable horizon mix that is directly investable.

This "double decoding" delivers six transparent sleeves from Stage A and a parsimonious, dynamically tilted blend in Stage B.

Notation. Markets $i = 1, \ldots, N$, horizons $h \in \mathcal{H} = \{20, 60, 125, 250, 500\}$. Let $m_{t,i}$ be market excess returns and $g_{t,i}^{(h)}$ the per-market trend-factor returns (Section 3). Define the factor-k cross section at t as

$$\mathbf{x}_t^{(k)} = \begin{cases} \big(g_{t,1}^{(h)}, \ldots, g_{t,N}^{(h)}\big)^\top & \text{if } k = h \in \mathcal{H}, \\ \big(m_{t,1}, \ldots, m_{t,N}\big)^\top & \text{if } k = \text{MKT}. \end{cases}$$

Conventions. Returns are excess, sleeves and benchmark are vol-targeted and de-meaned on an expanding window, and all state updates are forward-only (sequential) with no look-ahead.

Stage A: Six single-factor graphical decoders

For each $k \in \{20, 60, 125, 250, 500, \text{MKT}\}$ we fit a time-varying linear observation equation

$$r_t^{\text{Bench}} \;=\; \mathbf{x}_t^{(k)\top} \boldsymbol{\alpha}_t^{(k)} \;+\; \varepsilon_t^{(k)}, \qquad \varepsilon_t^{(k)} \sim \mathcal{N}(0, \sigma_k^2), \tag{1}$$

with Gaussian random-walk states

$$\boldsymbol{\alpha}_t^{(k)} \;=\; \boldsymbol{\alpha}_{t-1}^{(k)} \;+\; \boldsymbol{\eta}_t^{(k)}, \qquad \boldsymbol{\eta}_t^{(k)} \sim \mathcal{N}(\mathbf{0}, \mathbf{W}_t^{(k)}). \tag{2}$$

We enforce light structure for stability and interpretability:

$$\boldsymbol{\alpha}_t^{(k)} \succeq \mathbf{0}$$

implemented by projecting each Graphical model update onto the simplex. The one-step MAP update is the exponentially weighted ridge problem

$$\widehat{\boldsymbol{\alpha}}_t^{(k)} = \arg \min_{\alpha \succeq 0,\, \mathbf{1}^\top \alpha = 1} \Big\{ \sum_{\tau=1}^{t} \delta_k^{t-\tau} \big(r_\tau^{\text{Bench}} - \mathbf{x}_\tau^{(k)\top} \alpha\big)^2 + \lambda_k \|\alpha - \widehat{\boldsymbol{\alpha}}_{t-1}^{(k)}\|_2^2 \Big\}, \quad \delta_k \in (0, 1]. \tag{3}$$

The *factor-specific replication* (Stage A output) is

$$\widehat{r}_t^{(k)} \;=\; \mathbf{x}_t^{(k)\top} \widehat{\boldsymbol{\alpha}}_t^{(k)}. \tag{4}$$

Collect these six streams as $\widehat{\mathbf{r}}_t = \big(\widehat{r}_t^{(20)}, \widehat{r}_t^{(60)}, \widehat{r}_t^{(125)}, \widehat{r}_t^{(250)}, \widehat{r}_t^{(500)}, \widehat{r}_t^{(\text{MKT})}\big)^\top$.

Why "six decoders" instead of one market mix for all horizons? Running a dedicated cross-section for each factor lets the benchmark reveal *where* in the universe each time-scale lives (e.g., 500d may lean to rates/FX, 20d to equities/commodities). Stage A thus gives six clean, investable replications with distinct cross-sectional fingerprints.

Stage B: Same graphical model, dynamic factor-style selection

Same framework, different regressors. Stage B uses *exactly the same* graphical state–space specification as Stage A (same observation/state equations, same sequential Bayesian Graphical model updates, same constraints). The only change is the design matrix: instead of per-market series, we now use the six *Stage-A replications* (one per factor style) as the regressors. This lets Stage B *dynamically select the factor style mix* that best replicates the benchmark at each date.

Observation and state equations (unchanged form). We decode the benchmark on the six Stage-A replications $\widehat{\mathbf{r}}_t$:

$$r_t^{\text{Bench}} \;=\; \widehat{\mathbf{r}}_t^{\top}\boldsymbol{\beta}_t \;+\; \varepsilon_t^{(B)}, \qquad \varepsilon_t^{(B)} \sim \mathcal{N}(0, \sigma_B^2), \tag{5}$$

with time-varying coefficients following a Gaussian random walk:

$$\boldsymbol{\beta}_t \;=\; \boldsymbol{\beta}_{t-1} \;+\; \boldsymbol{\eta}_t^{(B)}, \qquad \boldsymbol{\eta}_t^{(B)} \sim \mathcal{N}(\mathbf{0}, \mathbf{W}_t^{(B)}). \tag{6}$$

We preserve interpretability with the same constraints used in Stage A:

$$\boldsymbol{\beta}_t \succeq \mathbf{0}, \qquad \mathbf{1}^{\top}\boldsymbol{\beta}_t \le 1 + \kappa \;(\kappa \text{ small}),$$

and update $\widehat{\boldsymbol{\beta}}_t$ via the *same* sequential MAP problem as in (3) (identical loss, decay, and ridge terms), now applied to the six factor-style inputs.

Replicator and attribution. The investable *Double Decoding* replicator is

$$\widehat{r}_t^{\text{DD}} \;=\; \widehat{\mathbf{r}}_t^{\top}\widehat{\boldsymbol{\beta}}_t, \tag{7}$$

where the components of $\widehat{\boldsymbol{\beta}}_t$ are *dynamic factor-style weights* (20/60/125/250/500/MKT). Horizon risk contributions under $\boldsymbol{\Sigma}_{\widehat{r}} := \operatorname{Var}(\widehat{\mathbf{r}}_t)$ are

$$\text{RC}_t^{(k)} = \frac{\widehat{\beta}_{t,k}\,\mathbf{e}_k^{\top}\,\boldsymbol{\Sigma}_{\widehat{r}}\widehat{\boldsymbol{\beta}}_t}{\widehat{\boldsymbol{\beta}}_t^{\top}\,\boldsymbol{\Sigma}_{\widehat{r}}\widehat{\boldsymbol{\beta}}_t}.$$

Turnover control (unchanged). To discourage over-reactive re-weighting we penalize ℓ_1 changes when producing implementable weights:

$$\mathcal{C}_t = c_\alpha \sum_k \|\widehat{\boldsymbol{\alpha}}_t^{(k)} - \widehat{\boldsymbol{\alpha}}_{t-1}^{(k)}\|_1 + c_\beta \|\widehat{\boldsymbol{\beta}}_t - \widehat{\boldsymbol{\beta}}_{t-1}\|_1,$$

and can report results net of a linear slippage model calibrated to TCA.

Interpretation. Stage B is a *dynamic factor-style selector*: using the same state-space machinery as Stage A, it learns the time-varying blend of the six styles that most closely reproduces the CTA benchmark, while keeping weights non-negative, stable, and directly interpretable.

Estimation and Identifiability

All models are fitted forward-only via Bayesian Graphical model updates (West and Harrison, 1997; Kim and Nelson, 1999; Carvalho et al., 2009; Koop, 2013). Decays (δ_k, δ_B) target 6–18 months of effective memory; (λ_k, λ_B) are set to match observed weight variability. Factor legs are volatility-targeted and de-meaned to remove scale/level indeterminacy; simplex constraints make Stage A identifiable; factor normalization and the budget constraint make Stage B identifiable.

Windows and Metrics

Weights use only information available up to t. We report annualized return, volatility, Sharpe ratio, Max Drawdown, and Return/MaxDD over:

- **Since Inception (SI)**: 2005-12-20 to 2025-11-05,
- **Last 10Y**: 2015-11-05 to 2025-11-05,
- **Last 5Y**: 2020-11-05 to 2025-11-05.

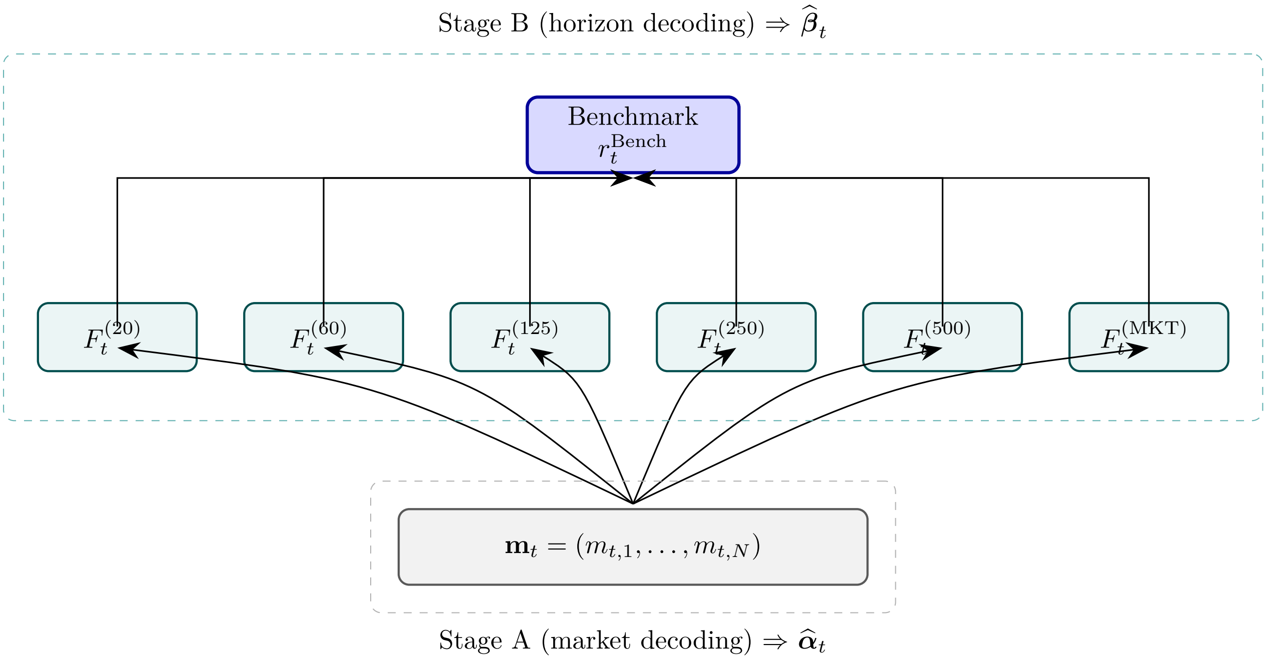

Figure 1: Two-layer *Double Decoding*: Stage A learns a market mix $\widehat{\boldsymbol{\alpha}}_t$; Stage B blends market-aligned sleeves via $\widehat{\boldsymbol{\beta}}_t$.

Stage B: Dynamic Allocation of Trend and Market Factors

What is being allocated? Stage B reuses the *same* graphical state–space framework as Stage A, but the regressors are now the six Stage–A replications (TF–20/60/125/250/500 and MKT). The filter learns time–varying, non–negative weights $\widehat{\boldsymbol{\beta}}_t$ across these *factor styles*, selecting at each date the mix that best replicates the CTA benchmark under the Stage B constraints.

Behaviour through time. Figure 2 shows the stacked evolution of the six weights. Fast sleeves (TF–20/60) expand in reversal–heavy episodes; the slow sleeves (TF–250/500) climb in persistent macro trends; the Market sleeve (MKT) remains small but non–zero, absorbing residuals when trend footprints are muted.

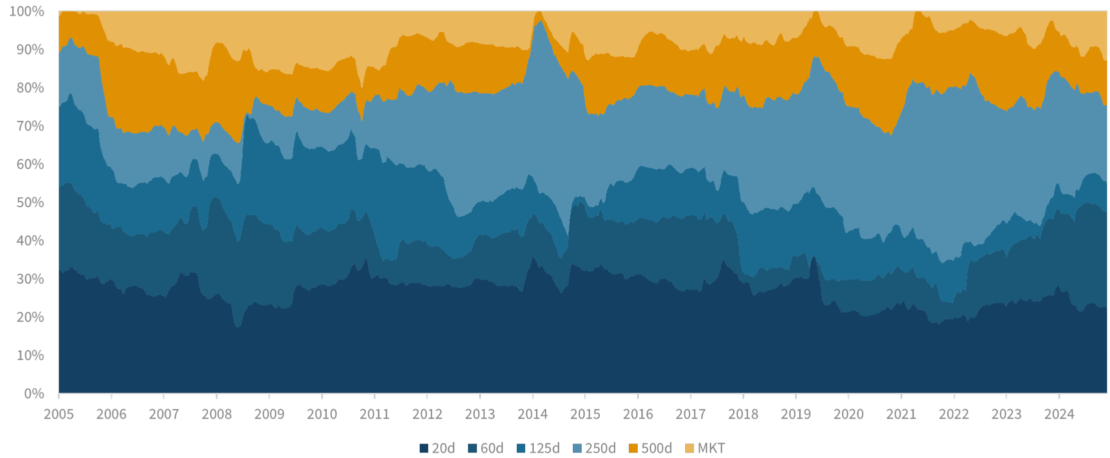

Figure 2: Stage B dynamic factor mix rebased to 100% (TF–20/60/125/250/500 and MKT shares).

Average allocations by window. Figure 3 summarizes average weights *since inception*, over the *last 10Y*, and the *last 5Y*. Two patterns stand out: (i) a structural tilt to the tails—TF–20 and TF–250/500—consistent with efficiency results; (ii) a declining mid–band share (TF–125), confirming that the 60–125d block is the weakest link for replication efficiency.

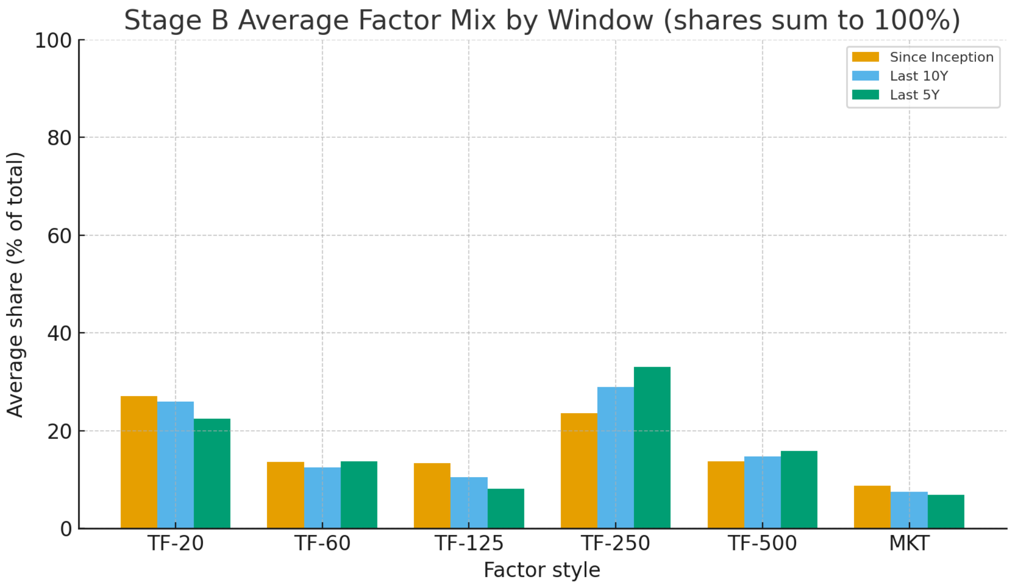

Figure 3: Average Stage B factor shares by window (Since Inception / Last 10Y / Last 5Y). Bars within each window sum to 100%.

Read–through. Stage B should be interpreted as a tactical horizon allocator rather than as a static six-sleeve average. The filter continuously arbitrates between reaction speed and persistence by reallocating conviction across the short end (TF–20/60), the slow block (TF–250/500), and a residual MKT sleeve, under sequential and non-negative updates on the six Stage–A replications. Figure 2 makes that rotation visible: around 2014 the mix tightens into a fast/slow barbell with a pronounced TF–250 build-up and little rebuilding of TF–60 or TF–125; in 2020 the allocation becomes more transitional, with less reliance on any single interior horizon and a broader split between fast reaction and slower confirmation; by 2022 the long end reasserts itself, with TF–250/500 forming the strategic core while TF–20 is retained as a tactical overlay for regime transition and re-entry. Figure 3 shows that this is the normal operating logic of the decoder rather than an isolated episode: the weight distribution remains structurally concentrated in the tails of the horizon curve, TF–125 is persistently subordinated to adjacent sleeves, and MKT stays small, so the middle of the curve is used opportunistically rather than as a standing core exposure.

Results

Strategy set (kept only what we define and use).

- *CTA Pure Trend Nd Decoding* ($N \in \{20, 60, 125, 250, 500\}$): single-horizon sleeves.
- *CTA Pure Trend Decoding*: Decoding with average trend factor (unique trend factor averaging all horizons) on each market.
- *CTA Adaptive Trend Decoding*: Decoding with average trend factor + Market factor on each market.
- *CTA No-MT Decoding*: Decoding with blend trend factor removing the 125d sleeve (see `https://aiforalpha.com/dist/img/Spotlight_MediumTerm.pdf`) for each market.
- *CTA Augmented ShortTerm Trend*: Market + average of short-term trend (10/20/60) for each market.
- *CTA Horizon Decoding*: our **Double CTA Decoding** (Stage A six single-factor replications blended in Stage B).

A. Since Inception (2005–12–20 to 2025–11–05)

Table 3: Since Inception metrics (required strategies only).

Strategy	Cum. Ret.	Annual Ret.	Vol	Sharpe	Max DD	Ret/MaxDD
CTA Pure Trend Decoding	628.1%	10.5%	11.1%	0.77	21.6%	0.49
CTA Pure Trend 20d Decoding	287.9%	7.1%	10.7%	0.48	18.7%	0.38
CTA Pure Trend 60d Decoding	385.5%	8.3%	11.0%	0.58	18.1%	0.46
CTA Pure Trend 125d Decoding	427.9%	8.7%	11.4%	0.60	23.7%	0.37
CTA Pure Trend 250d Decoding	499.7%	9.4%	11.3%	0.67	22.5%	0.42
CTA Pure Trend 500d Decoding	580.5%	10.1%	11.1%	0.74	14.5%	0.70
CTA Adaptive Trend Decoding	600.8%	10.3%	11.1%	0.75	18.5%	0.56
CTA No-MT Decoding	668.1%	10.8%	11.0%	0.81	20.0%	0.54
CTA Augmented ShortTerm Trend	564.3%	10.0%	10.4%	0.78	14.9%	0.67
CTA Horizon Decoding	642.7%	10.6%	10.7%	0.81	16.7%	0.64

Read-through. The best Sharpe ratio is achieved by *CTA Horizon Decoding* (0.81). The highest drawdown efficiency (Return/MaxDD) belongs to *CTA Pure Trend 500d* (0.70), which also records the lowest maximum drawdown (14.5%). *CTA Horizon Decoding* ranks among the top three strategies on this measure as well. Overall, Horizon leads on Sharpe, while the 500d sleeve remains the benchmark for drawdown resilience.

B. Last 10 Years (2015–11–05 to 2025–11–05)

Table 4: Last 10Y metrics (required strategies only).

Strategy	Cum. Ret.	Annual Ret.	Vol	Sharpe	Max DD	Ret/MaxDD
CTA Pure Trend Decoding	89.9%	6.6%	10.6%	0.41	21.6%	0.31
CTA Pure Trend 20d Decoding	56.1%	4.6%	10.0%	0.23	17.1%	0.27
CTA Pure Trend 60d Decoding	64.8%	5.1%	10.4%	0.27	16.1%	0.32
CTA Pure Trend 125d Decoding	62.6%	5.0%	10.8%	0.25	23.7%	0.21
CTA Pure Trend 250d Decoding	100.5%	7.2%	10.7%	0.46	22.5%	0.32
CTA Pure Trend 500d Decoding	111.5%	7.8%	10.5%	0.52	14.5%	0.54
CTA Adaptive Trend Decoding	120.7%	8.2%	10.8%	0.55	18.5%	0.45
CTA No-MT Decoding	97.0%	7.0%	10.4%	0.45	20.0%	0.35
CTA Augmented ShortTerm Trend	112.5%	7.8%	10.2%	0.54	14.9%	0.52
CTA Horizon Decoding	116.9%	8.0%	10.3%	0.56	16.7%	0.48

Read-through (10Y). Best Sharpe: *CTA Horizon Decoding* (0.56). Best Return/MaxDD and lowest MaxDD: *CTA Pure Trend 500d* (0.54; MaxDD 14.5%). Horizon is the Sharpe leader; 500d is the risk-efficiency anchor.

C. Last 5 Years (2020–11–05 to 2025–11–05)

Table 5: Last 5Y metrics (required strategies only).

Strategy	Cum. Ret.	Annual Ret.	Vol	Sharpe	Max DD	Ret/MaxDD
CTA Pure Trend Decoding	46.4%	7.9%	11.2%	0.41	21.6%	0.37
CTA Pure Trend 20d Decoding	51.7%	8.7%	10.4%	0.51	14.6%	0.59
CTA Pure Trend 60d Decoding	36.6%	6.4%	11.0%	0.28	16.1%	0.40
CTA Pure Trend 125d Decoding	36.7%	6.5%	11.4%	0.27	23.7%	0.27
CTA Pure Trend 250d Decoding	51.6%	8.7%	11.4%	0.47	22.5%	0.38
CTA Pure Trend 500d Decoding	74.2%	11.7%	11.0%	0.75	14.1%	0.83
CTA Adaptive Trend Decoding	53.0%	8.9%	11.6%	0.47	18.5%	0.48
CTA No-MT Decoding	51.4%	8.6%	11.0%	0.48	20.0%	0.43
CTA Augmented ShortTerm Trend	66.5%	10.7%	11.0%	0.66	14.9%	0.72
CTA Horizon Decoding	64.4%	10.4%	10.8%	0.65	16.7%	0.63

Read-through (5Y). *CTA Pure Trend 500d* dominates: Annual 11.7%, Sharpe 0.75, Return/MaxDD 0.83, and the lowest MaxDD 14.1%. *CTA Horizon Decoding* remains competitive on both Sharpe and Return/MaxDD, confirming the benefit of dynamically blending the six Stage-A replications.

Correlation to the SG CTA Benchmark

Table 6: Monthly correlation to SG CTA Trend (NEIXCTAT) from inception (2005–12–20 to 2025–11–05).

Strategy	**Correlation**
CTA Horizon Decoding	0.86
CTA No-MT Decoding	0.84
CTA Adaptive Trend Decoding	0.84
CTA Pure Trend Decoding	0.84
CTA Pure Trend 250d Decoding	0.81
CTA Augmented ShortTerm Trend	0.80
CTA Pure Trend 125d Decoding	0.78
CTA Pure Trend 500d Decoding	0.77
CTA Pure Trend 60d Decoding	0.73
CTA Pure Trend 20d Decoding	0.66

Note: Correlations are computed on monthly data since inception (2005–12–20 to 2025–11–05).

Read-through. On monthly data from inception, *CTA Horizon Decoding* shows the highest benchmark fit (0.86). *CTA No-MT*, *CTA Adaptive Trend*, and the all-horizon *Pure Trend* blend cluster close behind (all 0.84). Among single sleeves, **250d** tracks best (0.81); **500d** and **125d** follow (0.77–0.78), while the fast sleeves **60d/20d** sit lower (0.73/0.66) as diversifiers.

Framework Synthesis and Utility Analysis

From single-factor decoders to an investable blend

The build is two-layer and fully sequential. Stage A runs six parallel graphical decoders—five look-back straddle trend sleeves (20/60/125/250/500 days) and one market-return sleeve (MKT)—each producing a factor-specific replication with a simplex-constrained cross-section. Stage B regresses the benchmark on these six replications to obtain a non-negative, time-varying horizon blend that is directly investable (see Section). This design keeps the replication interpretable while stabilizing risk-adjusted payoffs.

Cobb–Douglas utility (since inception)

We summarize the trade-off between benchmark fidelity and drawdown efficiency with a Cobb–Douglas utility

$$U(C, R; \alpha) = C^{\alpha} R^{1-\alpha},$$

where C is the *monthly* correlation to SG CTA from inception and R is *since-inception* Return/MaxDD (Tables 3, 6). Unless specified, $\alpha = 0.5$.

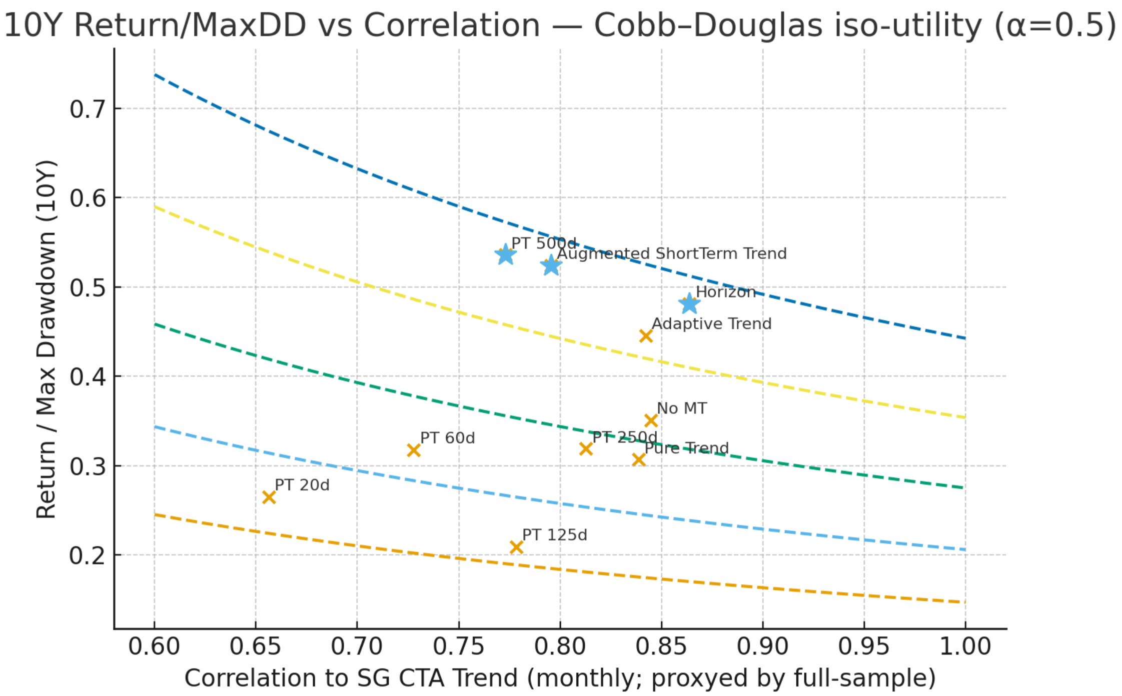

Figure 4: Since inception: Return/MaxDD vs. monthly correlation to SG CTA with Cobb–Douglas iso-utility curves ($\alpha = 0.5$).

Read-through. Figure 4 positions **CTA Horizon Decoding**, **Pure Trend 500d**, and **Augmented Short-Term** on the top iso-utility ridge at $\alpha = 0.5$, echoing the tables: Horizon leads on Sharpe, 500d on drawdown efficiency, while STT balances both with strong correlation.

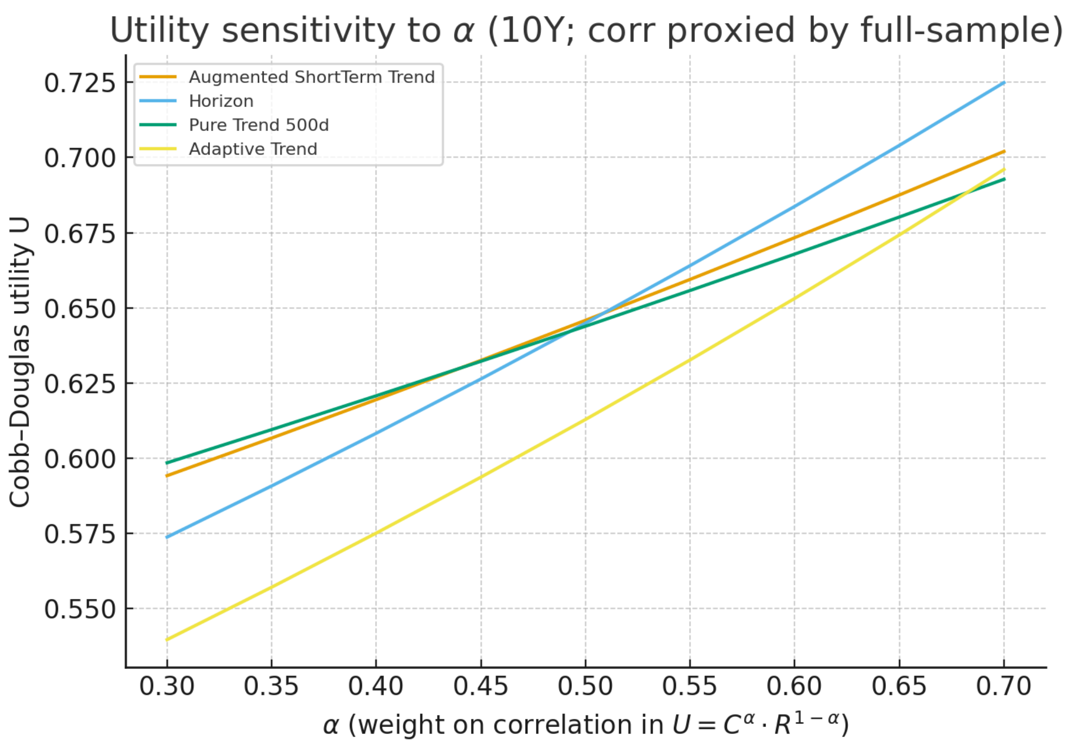

Figure 5: Since inception: utility sensitivity to α in $U = C^{\alpha} R^{1-\alpha}$ (preference for correlation vs. drawdown efficiency).

Preference map. Figure 5 shows the choice is robust to reasonable preference shifts:

- As α rises (heavier weight on fidelity), **CTA Horizon Decoding** becomes the natural top choice.
- As α falls (heavier weight on drawdown efficiency), **500d** becomes preferred.
- **Augmented Short-Term** stays competitive across the range, providing convex diversification.

Practical takeaway

For allocators balancing fidelity and efficiency equally ($\alpha = 0.5$), the **CTA Horizon Decoding** offers the most robust configuration. When benchmark tracking becomes the dominant objective (larger α), **CTA Horizon Decoding** naturally remains the preferred choice.

Discussion

The two-layer *Double CTA Decoding* delivers both benchmark fidelity and robust efficiency. Its architecture directly connects to the broader Ai For Alpha research program on interpretable replication and Bayesian graphical learning (Benhamou et al., 2025; Etienne et al., 2025; Ohana et al., 2024; Benhamou et al., 2024), where similar state-space formulations have been used to decode trend premia and redundancy across horizons. A few themes are consistent across windows and diagnostics.

Fidelity to SG CTA. Monthly correlations from inception (Table 6) place **CTA Horizon Decoding** at the top (0.86). **No-MT**, **Adaptive Trend**, and the all-horizon **Pure Trend** blend cluster just below (0.84). Among single sleeves, **250d** is the cleanest tracker (0.81) while fast sleeves (**60d/20d**) sit lower by design.

Risk-adjusted efficiency by window. Since inception (Table 3), **Horizon** and **No-MT** are essentially tied on Sharpe (both ≈ 0.81 within rounding), while **500d** is the drawdown-efficiency anchor (Return/MaxDD 0.70 with the lowest MaxDD 14.5%). Over the last 10 years (Table 4), **Horizon** leads on Sharpe (0.56) and **500d** leads on Return/MaxDD (0.54). In the last 5 years (Table 5), **500d** clearly dominates both Sharpe (0.75) and Return/MaxDD (0.83).

Utility view (since inception). The Cobb–Douglas analysis balances fidelity and efficiency via $U = C^{\alpha}R^{1-\alpha}$. At $\alpha = 0.5$, Figure 4 places **Horizon**, **500d**, and **Augmented Short-Term** on the top iso-utility ridge, echoing the tables. Preference sensitivity in Figure 5 shows a stable handoff: for lower α (heavier weight on drawdown efficiency) **500d** is preferred; as α rises (heavier weight on correlation), **Horizon** becomes the natural top choice. **Augmented Short-Term** stays competitive across the range and is a useful convex complement.

Role of the mid-band. The mid-term block (60–125d) is consistently the weakest contributor to efficiency: **125d** and often **60d** sit lower on Sharpe and Return/MaxDD across windows. This supports a structural *de-weight* of the mid-band in replication blends.

Interpretability and implementation. Stage A's six decoders reveal where each time-scale lives in the cross-section; Stage B's non-negative blend turns those into a transparent, investable tracker. The outcome is a high-fidelity, low-complexity tracker that avoids over-reliance on any single sleeve while preserving clear, additive attribution by horizon.

Conclusion

Double CTA Decoding replicates SG CTA with high fidelity and robust efficiency while keeping the allocation fully interpretable.

- **Core replicator. CTA Horizon Decoding** offers the most robust configuration when balancing fidelity and efficiency equally ($\alpha = 0.5$), sits on the top utility ridge (Figure 4), and achieves the highest benchmark fit (Table 6).
- **Drawdown anchor.** The **500d** sleeve is the most efficient drawdown hedge across windows (Tables 3, 4, 5) and becomes the preferred sleeve as preferences tilt toward drawdown efficiency (lower α in Figure 5).
- **Convex complement. Augmented Short-Term** provides a persistent, convex contribution and remains competitive on the utility frontier; it is a natural complement to Horizon+500d.
- **Mid-band hygiene.** Consistent with the evidence, **de-weight** the 60–125d block; removing or lightening the medium sleeve improves efficiency without harming correlation.

Practical takeaway. For allocators balancing fidelity and efficiency ($\alpha = 0.5$), use **CTA Horizon Decoding** as the core replicator, supported by a **500d** backbone and an **Augmented Short-Term** sleeve. As tracking becomes paramount (larger α), **Horizon** remains the preferred choice; as drawdown efficiency dominates (smaller α), tilt toward **500d**. As tracking stringency rises (larger α), **Horizon** remains the default core; as drawdown priority rises (smaller α), tilt toward **500d** and keep an **Augmented Short-Term** sleeve as a convex complement.

Disclaimer

This document is for research discussion only. Backtested results are hypothetical and subject to model risk; they do not reflect trading costs, capacity limits, or operational frictions beyond those stated. Past performance is not indicative of future results.

References

Benhamou, E., Ohana, J.-J., Etienne, A., Guez, B., Setrouk, E., and Jacquot, T. (2025). Re-evaluating short- and long-term trend factors in cta replication: A bayesian graphical approach. *arXiv preprint arXiv:2507.15876*.

Benhamou, E., Ohana, J.-J., and Guez, B. (2024). Grip: Graphical models revealing insights for portfolio replication – a learning approach. *SSRN Electronic Journal*, (4780148).

Carvalho, C. M., Polson, N. G., and Scott, J. G. (2009). Handling sparsity via the horseshoe. *Journal of Machine Learning Research*, 5(Nov):73–80. See also related dynamic shrinkage work by the authors in Bayesian time series.

Etienne, A., Ohana, J.-J., Benhamou, E., Guez, B., Setrouk, E., and Jacquot, T. (2025). Revisiting the structure of trend premia: When diversification hides redundancy. *arXiv preprint arXiv:2510.23150*.

Kim, C.-J. and Nelson, C. R. (1999). *State-Space Models with Regime Switching: Classical and Gibbs-Sampling Approaches with Applications*. MIT Press, Cambridge, MA.

Koop, G. (2013). *Bayesian Econometrics*. Wiley, Chichester, 2nd edition.

Ohana, J.-J., Benhamou, E., Saltiel, D., and Guez, B. (2024). Deep decoding of strategies. *SSRN Electronic Journal*, (4128693).

West, M. and Harrison, J. (1997). *Bayesian Forecasting and Dynamic Models*. Springer, New York, 2nd edition.