For Alpha

Ai-Powered Investment Replication

Strategy Spotlight: CTA Horizon Style Decomposition

November 2025, Ai For Alpha Team

Abstract

We build a library of five *mono-horizon* trend-following strategies (20d, 60d, 125d, 250d, 500d) on a diversified futures universe and use them as building blocks for CTA horizon attribution.

At the index level we regress the SG CTA Trend index on these sleeves over the last five years. The index is well explained by a positive combination of three horizons (20d, 125d, 500d), with an economically negligible intercept and a total trend beta close to one: SG CTA Trend behaves like a fully invested multi-horizon trend portfolio. Crucially, the regression does *not* simply select the most correlated sleeves. The 20d strategy is the least correlated horizon to the index in isolation, yet it is retained alongside the slow 500d sleeve. This reveals genuine *replication synergies* between short- and long-term trend that are invisible in univariate correlation analysis.

At the manager level we regress seven *anonymised* SG CTA Trend constituents on the same library. Each programme is well explained by a sparse, positive mix of two or three horizons, again with small intercepts and highly significant coefficients. All managers load meaningfully on a slow sleeve (250d or 500d) and at least one fast sleeve (20d or 60d), while the mid-band (60–125d) varies across programmes. CTA performance thus appears largely driven by a small, interpretable set of horizon factors. Differences across managers are mainly differences in horizon mix and trend intensity—in particular, how they combine fast and slow sleeves—rather than fundamentally distinct strategies.

1 Background and Motivation

CTA replication is often framed either at the market level (bottom-up, using futures directly) or at the factor level (top-down, using generic trend and carry factors). Our earlier work on horizon structuring and medium-term crowding argues for decomposing trend premia by look-back window and de-emphasising the crowded mid band:

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

Here we keep the same mono-horizon sleeves but change perspective: they become regressors used to decompose both the SG CTA index and its constituents. The core idea is:

> *Construct investable mono-horizon strategies and regress each CTA—and the index—on this common horizon library.*

This yields a transparent "horizon fingerprint" for each programme. Managers are anonymised (CTA 1–CTA 7), but all belong to the SG CTA Trend universe. The key questions are: *which horizons are actually used, how are fast and slow sleeves combined, and how does this map to realised performance?*

2 Mono-Horizon Framework

2.1 Investment Universe and Costs

Tests run on a liquid, diversified futures universe representing standard CTA exposures across equities, rates, FX and commodities. Implementation costs are applied consistently across all sleeves, as summarised in Table 1.

- **Transaction cost.** Round-turn execution cost (bid–ask, brokerage, exchange and clearing fees plus slippage), in basis points of notional. Orders are *netted across sleeves* at the instrument level.
- **Replication (roll) cost.** Systematic carry/roll drag when the front contract is rolled to the next maturity, based on long-run average front-to-next calendar spreads.
- **Management fee.** Flat 50 bps per annum on AUM.

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 all mono-horizon sleeves.

Implementation. All sleeves are volatility-targeted. Costs are applied to *filled* notional after sleeve netting. Rolls follow exchange calendars and incur costs only when contracts are actually rolled. Because orders are netted per instrument across sleeves before routing, reported turnover and costs reflect the implementation of the full mono-horizon library.

2.2 Look-Back Straddles as Horizon-Controlled Trend Factors

A look-back option pays on the *path* of prices, not just the close. A long look-back straddle (call+put) 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 near zero in ranges—the same logic used in breakout CTAs.

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 $S_{t,i}^{(h)}$. A simple look-back-style trend score is

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

mapped to a bounded position $\pi_{t,i}^{(h)} \in [-1, 1]$ (e.g. via tanh), with per-market factor return

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

rescaled to a common volatility target per horizon.

We use five horizons,

$$\mathcal{H} = \{20, 60, 125, 250, 500\},$$

which span fast to very slow trend.

Table 2: Look-back straddle mechanics and CTA analogues

Concept	**Intuition**	**CTA Analogue**
Path-dependent pay-off	Rewards *max/min* excursions over h.	Trend P&L depends on the path, 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.	Fast vs. slow trend horizons.
Convexity	Positive skew from large moves.	CTAs' crisis convexity ("crisis alpha").

2.3 Mono-Horizon Trend Sleeves

For each horizon $h \in \mathcal{H}$ we aggregate per-market returns into a portfolio sleeve

$$F_t^{(h)} = \sum_{i=1}^{N} w_{i,t}^{(h)} \, g_{t,i}^{(h)},$$

where $w_{i,t}^{(h)}$ are simple risk-parity weights rescaled so each sleeve targets the same annualised volatility. The sleeves

$$\{F^{(20)}, F^{(60)}, F^{(125)}, F^{(250)}, F^{(500)}\}$$

constitute our mono-horizon library; they are the only regressors used in the index and manager decompositions.

3 SG CTA Trend Index: Horizon Style

3.1 Static Horizon Regression of SG CTA Trend

We model the SG CTA Trend index as a linear combination of mono-horizon sleeves:

$$r_t^{\mathrm{SG}} \;=\; \alpha_{\mathrm{SG}} + \sum_{h\in\mathcal{H}} \beta_h^{\mathrm{SG}}\, F_t^{(h)} + \varepsilon_t^{\mathrm{SG}}, \tag{1}$$

estimated by OLS over the last five years of daily excess returns. Non-significant horizons are removed by applying backward elimination: at each step we drop the horizon with the highest p-value above 5%, re-estimate, and iterate until all remaining horizons are significant at 5%. The final specification retains three sleeves.

Table 3: SG CTA Trend index: horizon decomposition (last 5Y).

Horizon	**Coef.**	**Std.Err.**	t	$P > \lvert t \rvert$	**2.5%**	**97.5%**
const	-0.0002	0.0005	-0.41	0.685	-0.0012	0.0008
20d	0.3297	0.0457	7.22	< 0.001	0.2398	0.4197
125d	0.3802	0.0560	6.79	< 0.001	0.2700	0.4905
500d	0.3465	0.0485	7.14	< 0.001	0.2509	0.4421

Key points:

- The intercept is economically small and statistically insignificant: once horizon styles are included, there is little residual "alpha".
- The index loads significantly on three horizons: fast (20d), mid (125d) and slow (500d), all with t-statistics above 6.7.
- The sum of coefficients is close to one (about 1.06), so SG CTA Trend behaves like a fully invested multi-horizon trend portfolio.
- The mid/slow block (125d+500d) accounts for roughly two-thirds of the exposure; 20d contributes the remaining third. The index is structurally slow/mid-trend with a fast overlay.
- **Importantly, the regression does not just pick the most correlated sleeves.** As seen below, 125d and 250d are the best single-horizon trackers by correlation, yet the regression keeps 20d and 500d and drops 250d. The fit relies on interactions between horizons rather than on any single "closest" sleeve.

The fitted value from (1),

$$\widehat{r}_t^{\mathrm{H}} \;=\; \widehat{\alpha}_{\mathrm{SG}} + \sum_{h} \widehat{\beta}_h^{\mathrm{SG}}\, F_t^{(h)},$$

is our regression-based multi-horizon representation of the index (CTA Horizon Decoding), used here mainly as a style descriptor.

3.2 Mono-Horizon Correlation to the CTA Index

Table 4 shows the correlation of each mono-horizon "Pure Trend Nd" sleeve to the official CTA index (NEIXCTAT), together with the correlation structure across sleeves.

Table 4: Correlation matrix of mono-horizon sleeves and CTA index (monthly, in %).
PT 20d/60d/125d/250d/500d = CTA Pure Trend Nd Decoding; CTA Idx = NEIXCTAT Index.

	PT 20d	PT 60d	PT 125d	PT 250d	PT 500d	CTA Idx
PT 20d	100%	83%	66%	55%	46%	66%
PT 60d	83%	100%	88%	73%	61%	77%
PT 125d	66%	88%	100%	90%	78%	82%
PT 250d	55%	73%	90%	100%	93%	82%
PT 500d	46%	61%	78%	93%	100%	78%
CTA Idx	66%	77%	82%	82%	78%	100%

Two observations are central:

- **20d is the least correlated horizon to the index.** Its correlation to NEIXCTAT is 66%, well below the 125d and 250d sleeves (both 82%). In isolation it is a weak tracker.
- **Yet 20d is systematically selected in regressions.** The index regression retains 20d and 500d and drops 250d, and several manager regressions below also keep 20d. The least correlated sleeve is consistently part of the optimal mix, while better individual trackers can be excluded. This highlights strong *complementarity* between horizons: fast and slow sleeves bring distinct information and combine more effectively than any sleeve on its own.

Fast trend thus appears as a structural overlay: not the primary standalone tracker, but a persistent component that improves the joint payoff with slow trend.

4 Manager-Level Horizon Decomposition

4.1 Sample and Specification

We select seven trend-following programmes that are, or have been, constituents of the SG CTA Trend index, anonymised as **CTA 1** to **CTA 7**. For each manager we consider the last five years of daily excess returns, using the same futures universe, volatility targeting and cost assumptions as in the index study.

For manager j we start from the full model

$$r_{j,t} = \alpha_j + \sum_{h \in \{20,60,125,250,500\}} \beta_{j,h} F_t^{(h)} + \varepsilon_{j,t}, \tag{2}$$

and then apply backward elimination: at each step we drop the horizon with the highest p-value above 5%, re-estimate, and iterate until all remaining horizons are significant at 5%. The final sparse regression is the manager's *period decomposition*. All coefficients are estimated by OLS; p-values below 0.001 are reported as < 0.001.

4.2 Period Decompositions (Last 5Y)

Table 5: Period decomposition for CTA 1 (5-year window).

Horizon	**Coef.**	**Std.Err.**	t	$P > \|t\|$	**2.5%**	**97.5%**
const	0.0008	0.0007	1.26	0.208	-0.0005	0.0021
60d	0.2741	0.0521	5.26	< 0.001	0.1715	0.3766
500d	0.4648	0.0488	9.51	< 0.001	0.3686	0.5610

Table 6: Period decomposition for CTA 2 (5-year window).

Horizon	**Coef.**	**Std.Err.**	t	$P > \|t\|$	**2.5%**	**97.5%**
const	-0.0012	0.0006	-2.00	0.047	-0.0024	-0.0000
60d	0.4422	0.0814	5.43	< 0.001	0.2819	0.6024
125d	0.2976	0.0981	3.03	0.003	0.1045	0.4908
500d	0.4135	0.0587	7.04	< 0.001	0.2979	0.5291

Table 7: Period decomposition for CTA 3 (5-year window).

Horizon	**Coef.**	**Std.Err.**	t	$P > \|t\|$	**2.5%**	**97.5%**
const	-0.0008	0.0008	-1.13	0.261	-0.0023	0.0006
20d	0.4289	0.0653	6.57	< 0.001	0.3004	0.5573
125d	0.4555	0.0829	5.49	< 0.001	0.2924	0.6185
500d	0.3243	0.0717	4.52	< 0.001	0.1832	0.4654

Table 8: Period decomposition for CTA 4 (5-year window).

Horizon	**Coef.**	**Std.Err.**	t	$P > \|t\|$	**2.5%**	**97.5%**
const	-0.0015	0.0008	-1.89	0.060	-0.0031	0.0001
20d	0.4692	0.0693	6.77	< 0.001	0.3327	0.6057
125d	0.3910	0.0849	4.61	< 0.001	0.2238	0.5582
500d	0.3453	0.0736	4.69	< 0.001	0.2003	0.4903

Table 9: Period decomposition for CTA 5 (5-year window).

Horizon	**Coef.**	**Std.Err.**	t	$P > \|t\|$	**2.5%**	**97.5%**
const	-0.0003	0.0006	-0.56	0.574	-0.0015	0.0008
20d	0.2965	0.0783	3.78	< 0.001	0.1422	0.4507
60d	0.3436	0.0868	3.96	< 0.001	0.1726	0.5146
250d	0.2320	0.0498	4.66	< 0.001	0.1339	0.3302

Table 10: Period decomposition for CTA 6 (5-year window).

Horizon	**Coef.**	**Std.Err.**	t	$P > \|t\|$	**2.5%**	**97.5%**
const	-0.0014	0.0007	-1.83	0.069	-0.0028	0.0001
20d	0.3795	0.0649	5.85	< 0.001	0.2518	0.5072
125d	0.3292	0.0794	4.14	< 0.001	0.1728	0.4856
500d	0.4962	0.0689	7.20	< 0.001	0.3605	0.6319

Table 11: Period decomposition for CTA 7 (5-year window).

Horizon	**Coef.**	**Std.Err.**	t	$P > \|t\|$	**2.5%**	**97.5%**
const	0.0001	0.0004	0.13	0.896	-0.0008	0.0009
20d	0.2405	0.0376	6.40	< 0.001	0.1665	0.3144
125d	0.2897	0.0460	6.30	< 0.001	0.1992	0.3803
500d	0.2504	0.0399	6.28	< 0.001	0.1718	0.3289

4.3 Cross-Sectional Patterns

Table 12: Horizon shares (in %) for the index SG CTA Trend and the 7 CTAs (5Y regressions on mono-horizon trend factors, coefficients rebased to 100%).

	20d	**60d**	**125d**	**250d**	**500d**
Index (SG CTA Trend)	31%	0.00	36%	0.00	33%
CTA 1	0.00	37%	0.00	0.00	63%
CTA 2	0.00	38%	26%	0.00	36%
CTA 3	35%	0.00	38%	0.00	27%
CTA 4	39%	0.00	32%	0.00	29%
CTA 5	34%	39%	0.00	27%	0.00
CTA 6	31%	0.00	27%	0.00	41%
CTA 7	31%	0.00	37%	0.00	32%

Four messages summarise the cross-section:

(i) **Trend sleeves explain most of the action.** Horizon coefficients are positive and strongly significant across all managers. Intercepts are small; only one manager shows a borderline significant negative intercept, consistent with fees and implementation drag. The mono-horizon library captures most of each CTA's time-series variation.

(ii) **Common structure: fast + slow, flexible mid band.** Every manager loads materially on at least one short horizon (20d or 60d) and one long horizon (250d or 500d). The slow sleeve is particularly stable: six out of seven managers tilt to 500d (the remaining one uses 250d). Exposure to the 60–125d band is more heterogeneous and a key differentiator.

(iii) **Total trend beta near (but not fixed at) one.** The sum of horizon coefficients per manager ranges from about 0.75 to 1.20. Some managers (CTA 3, CTA 4, CTA 6) resemble fully invested multi-horizon portfolios; others (CTA 1, CTA 5, CTA 7) run slightly lower trend beta. Differences across CTAs are primarily differences in horizon mix and intensity.

(iv) **Regression favours horizon synergies over single best sleeves.** As for the index, regressions do not systematically pick the single most correlated horizon. Several CTAs combine the relatively low-correlation 20d sleeve with slow sleeves instead of loading only on mid-band horizons. Short and long horizons act as complements, diversifying each other's misspecification and improving overall tracking.

4.4 Performance and Horizon Mix (Last 5Y)

We now link horizon fingerprints to realised risk-adjusted performance. Table 13 reports Sharpe ratios and Return/MaxDD for the seven CTAs and for the SG CTA Trend index over the last five years.

Table 13: Risk-adjusted performance of CTA managers and SG CTA Trend index (last 5Y).

	Sharpe Ratio	Return/Max DD
CTA 1	0.75	0.84
CTA 2	-0.09	0.06
CTA 3	0.01	0.13
CTA 4	-0.14	0.04
CTA 5	0.10	0.16
CTA 6	-0.04	0.07
CTA 7	0.38	0.37
SG CTA Trend	0.38	0.35

Combined with Table 12, three patterns emerge.

(a) Slow-trend–dominated CTA 1 as risk-efficiency outlier. CTA 1 delivers the best risk-adjusted performance, with Sharpe 0.75 and Return/MaxDD 0.84, far above the index (0.38 and 0.35). Its horizon mix is dominated by slow trend (about 63% in 500d, 37% in 60d). This echoes earlier backtests where the 500d sleeve offered the strongest drawdown efficiency. CTA 1 looks like a well-implemented "slow-plus-somewhat-fast" trend portfolio.

(b) CTA 7 as an index-like multi-horizon blend. CTA 7 exhibits risk-adjusted metrics very close to SG CTA Trend and a horizon mix (roughly one-third 20d, one-third 125d, one-third 500d) that closely matches the index's 31/36/33 split. It effectively delivers an efficient, slightly de-levered version of the benchmark's own fast/mid/slow structure.

(c) Underperformers concentrate fast or mid-band risk. CTAs 2, 4 and 6 all show weak or negative Sharpe and poor Return/MaxDD despite meaningful slow exposure:

- CTA 2 combines 60d, 125d and 500d with a relatively heavy mid-band component but little evidence of improved risk-efficiency.
- CTA 4 emphasises very fast trend (about 39% in 20d) on top of balanced 125d/500d exposure yet has the weakest profile.
- CTA 6 has the highest share of 500d (about 41%) but still underperforms, underscoring that a slow backbone is necessary but not sufficient; implementation quality, risk scaling and non-trend components also matter.

CTA 5, which mixes 20d, 60d and 250d without 125d/500d, sits in between: moderate performance and a more idiosyncratic horizon profile.

Overall, managers with index-like horizon mixes deliver index-like risk-adjusted returns (CTA 7), a strong tilt to slow trend can materially improve drawdown-adjusted performance (CTA 1), and aggressive fast or mid-band tilts without an overwhelming slow core are associated with weaker realised Sharpe and Return/MaxDD in this sample.

5 Allocation Takeaways

The index and manager decompositions, together with realised performance, suggest a simple allocation framework:

- **Slow trend as backbone.** The index and all managers load strongly on slow horizons (250d/500d). This block drives drawdown resilience and forms the common core of CTA behaviour. The outperformance of CTA 1 reinforces the value of a strong slow backbone.

- **Fast overlay for convexity and regime detection.** Fast sleeves (20d/60d) are used by all managers and appear in the index regression despite lower standalone correlation. They add reactivity and convexity and are valuable *marginally*, in combination with slow trend.
- **Mid band as style dial.** The 60–125d region is the main area of stylistic dispersion. It is present in the index and in several CTAs, but in very different proportions. The weaker Sharpe of mid-band–heavy managers suggests that this dial should be used cautiously rather than as a default core.
- **Focus on horizon combinations, not isolated sleeves.** The most striking result is that regressions systematically choose horizon mixes that are *not* simply the most correlated individual sleeves: the 20d horizon is always retained, even though it is the least correlated sleeve to the index. Replication quality is driven by cross-horizon diversification—how fast and slow sleeves interact—more than by any single sleeve's correlation.

For allocators, the mono-horizon library offers a concise "style basis": CTAs can be compared and combined by their fast/mid/slow mix and overall trend beta, and by how they exploit the synergy between short- and long-term trend.

6 Conclusion

Mono-horizon trend decomposition provides a compact, interpretable map of trend-following risk:

- At the **index level**, a small set of mono-horizon sleeves explains the SG CTA Trend index with a negligible intercept and a total trend beta close to one. The index is effectively a convex combination of 20d, 125d and 500d horizons, with a structural tilt to mid/slow trend and a meaningful fast overlay. The regression selects this mix because of horizon synergies, not because these are the most correlated sleeves individually.
- At the **manager level**, the same sleeves explain the systematic component of anonymised SG CTA Trend constituents over the last five years. Each manager corresponds to a different convex combination of the same horizons, with a shared slow backbone and specific fast/mid tilts. Again, regressions favour combinations of fast and slow sleeves over any single "best" horizon.
- The **fast/slow structure** is robust across managers; the mid band is optional and heterogeneous. Differences across CTAs can therefore be read as differences in horizon mix, trend intensity and the way fast and slow sleeves are combined, rather than as fundamentally different risk premia.
- The **5Y performance evidence** is consistent with this map: the best risk-efficiency belongs to the most slow-trend–dominated manager (CTA 1); managers with index-like horizon mixes deliver index-like Sharpe and drawdown profiles (CTA 7); and programmes that take large fast or mid-band bets without a dominant slow backbone tend to underperform on a Sharpe and Return/MaxDD basis.

Mono-horizon sleeves thus provide a practical common language for both CTA analysis and allocation: they reveal how much trend risk is taken at each horizon and how that choice translates into observed performance.

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.