For Alpha

Ai-Powered Investment Replication

Strategy Spotlight: Decoding Alpha in Practice

Excess Performance of Enhanced vs. Baseline Decoders

March 2026, Ai For Alpha Team

Abstract

- **Setup.** We compare *baseline* and *enhanced* decoding across four strategy families: Hedge Funds, Long Short Equity, Global Balanced, and Risk Parity.
- **Method.** Enhanced decoding replaces the raw benchmark target with the benchmark plus a constant annualized trend objective. The clean empirical test is the portfolio *long Enhanced / short Baseline*, which removes most benchmark beta and isolates the effect of the changed inferred weights.
- **Main results.** Over the common window 2012-01-02 to 2026-03-16, *Hedge Funds* stands out as the clearest diversification result, with the **lowest correlation to U.S. equities at 15.0%** and the best return-to-drawdown profile (0.50). *Global Balanced* delivers the highest absolute excess return (182.6% cumulative, 7.3% annualized), *Long Short Equity* achieves the highest Sharpe ratio (0.99), and *Risk Parity* is positive but materially weaker.
- **Cross-strategy link.** Weekly return correlations between the four long-enhanced/short-baseline decoding strategies are moderate to high, with the strongest link between Hedge Funds and Long Short Equity (84%) and the weakest links involving Risk Parity (31–50%).
- **Interpretation.** Position analysis shows that decoding alpha is shared mainly through bonds, credit, and developed FX, while equity and U.S.-duration tilts explain most of the cross-strategy differences.

1 Introduction

Motivation

Decoding asks whether an opaque or expensive strategy can be represented by a transparent portfolio of liquid instruments. Earlier work treats this as a sequential inference problem with dynamic Bayesian graphical models and stable time-varying weights (Ohana et al., 2022). Enhanced decoding modifies the target itself: instead of decoding the raw benchmark, it decodes the benchmark plus a constant positive trend objective (Benhamou et al., 2024; Ai For Alpha Team, 2025).

- **Question.** Does enhanced decoding add value beyond baseline decoding, and where does that value come from?
- **Clean test.** We study the excess portfolio $w_t^{\text{Enh}} - w_t^{\text{Base}}$, which isolates the incremental tilts created by enhancement on the same liquid universe.
- **Scope.** We apply this comparison to four strategy families: *Hedge Funds* (HF), *Long Short Equity* (LS), *Global Balanced* (GB), and *Risk Parity* (RP).
- **Objective.** We measure whether enhancement adds value across strategy families and identify the sleeves and markets that drive the difference.
- **Preview.** A central empirical result is that *Hedge Funds* is the least equity-linked excess strategy, with the **lowest correlation to U.S. equities at 15.0%** on the common sample.

Structure of the Paper

- Section 2 presents the graphical-model framework and the liquid investment universe.
- Section 3 defines the enhanced target and the long-enhanced/short-baseline excess portfolio.
- Section 4 reports performance, cross-strategy weekly return correlations of the long-enhanced/short-baseline portfolios, position overlap, and targeted trend-shock experiments on Global Balanced.
- Section 5 summarizes the main findings and practical implications.

2 Methodology

Graphical Model

We model a target daily excess return r_t^{target} with a time-varying linear specification based on a set of p observable liquid instrument returns. Let

$$\mathbf{x}_t = \left(r_{t,1}, r_{t,2}, \ldots, r_{t,p}\right)^\top.$$

The observation equation is

$$r_t^{\text{target}} = \mathbf{x}_t^\top \boldsymbol{\beta}_t + \varepsilon_t, \qquad \varepsilon_t \sim \mathcal{N}(0, \sigma_\varepsilon^2), \tag{1}$$

where $\boldsymbol{\beta}_t \in \mathbb{R}^p$ are time-varying coefficients. Coefficients evolve smoothly through the Gaussian state equation

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

with $\mathbf{W}_t$ controlling the adaptation rate. An initial prior $\boldsymbol{\beta}_0 \sim \mathcal{N}(\mathbf{m}_0, \mathbf{C}_0)$ provides shrinkage (West and Harrison, 1997; Kim and Nelson, 1999).

Objective optimized at each t. Sequential estimation solves the one-step MAP problem

$$\widehat{\boldsymbol{\beta}}_t = \arg\min_{\boldsymbol{\beta}\in\mathbb{R}^p} \left\{ \sum_{\tau=1}^{t} \delta^{\,t-\tau} \left(r_\tau^{\text{target}} - \mathbf{x}_\tau^\top \boldsymbol{\beta}\right)^2 + \left(\boldsymbol{\beta} - \boldsymbol{\beta}_{t-1}\right)^\top \mathbf{W}_t^{-1} (\boldsymbol{\beta} - \boldsymbol{\beta}_{t-1}) \right\}, \qquad \delta \in (0,1], \quad (3)$$

which combines an exponentially weighted prediction error with a state-smoothness penalty implied by (2). Forward filtering implements this objective using only information available at time t (West and Harrison, 1997; Kim and Nelson, 1999; Ohana et al., 2022; Benhamou et al., 2024).

Figure 1 displays the target return on the *top row* and four illustrative liquid instrument returns beneath. Blue horizontal links show the temporal evolution of each observed series. At every date t, orange dotted arrows illustrate contemporaneous relationships among the liquid returns and between the target and each instrument. The decoder therefore treats the benchmark as a latent linear combination of evolving liquid sleeves rather than as a static regression.

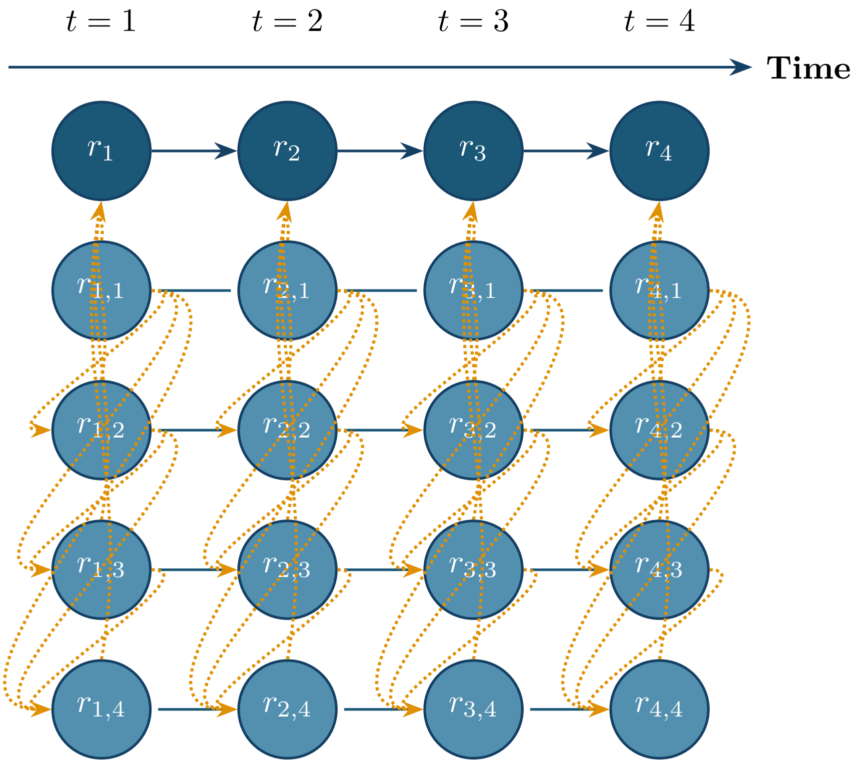

Figure 1: Graphical model with target return (top row) and four illustrative liquid instrument returns (rows below). Blue horizontal links show time evolution of each observed series. At each t, orange dotted arrows depict contemporaneous relationships among the liquid sleeves and between the target and each sleeve. The regression mapping is governed by (1)–(2) and optimized via (3).

Investment Universe and Strategy Families

We work with four decoded strategy families that all rely on a common liquid macro universe, with small family-specific variations. This shared universe makes the long-enhanced/short-baseline comparison economically meaningful: each excess strategy is implemented with the *same language of liquid instruments*, so differences arise from decoding choices rather than from a broader opportunity set.

Equity markets	U.S. broad equities; U.S. small caps; U.S. technology; Japan; Euro area; U.K.; emerging markets
Credit markets	CDX North America High Yield 5Y; iTraxx Crossover 5Y
Bond markets	U.S. 10-year; Japan 10-year; Germany 10-year; U.K. 10-year; Canada 10-year
FX markets	AUD/USD; CAD/USD; CHF/USD; EUR/USD; GBP/USD; JPY/USD (not used in Risk Parity)
Commodity markets	Gold; Brent crude; Copper; Natural Gas (Global Balanced only)

For the position analysis below, markets are also grouped into broad sleeves — Bonds, Commodities, Credit, Equities, and FX — to distinguish common macro structure from benchmark-specific expression. The Risk Parity excess strategy has no FX sleeve in the provided implementation, which is itself informative: part of the difference between RP and the other families is structural rather than merely tactical.

In addition to sharing a liquid macro universe, each family is decoded subject to the exposure discipline of the strategy it replicates. In practice, that means family-specific limits or implementation guidelines on total leverage and on major sleeves such as equities, bonds/rates, commodities, and U.S.-dollar exposure. These differ across HF, LS, GB, and RP because the replicated benchmarks differ. Within a given family, the *Baseline* and *Enhanced* decoders are evaluated under the same broad constraint framework, so the excess portfolio should mainly be read as the effect of the target change rather than as the effect of materially looser limits.

3 From Baseline Decoding to Enhanced Decoding

Baseline Decoder

Let r_t^{B} denote the raw target return of a given strategy family. The baseline decoder solves

$$r_t^{\mathrm{B}} = \mathbf{x}_t^\top \mathbf{w}_t^{\mathrm{Base}} + \varepsilon_t, \qquad \mathbf{w}_t^{\mathrm{Base}} = \mathbf{w}_{t-1}^{\mathrm{Base}} + \boldsymbol{\eta}_t, \tag{4}$$

where $\mathbf{w}_t^{\mathrm{Base}}$ are time-varying liquid exposures estimated sequentially with shrinkage. This decoder is designed to track the original target faithfully, but it still learns from the *raw* observed series and can therefore react to temporary noise.

Throughout the paper, decoder weights are *not* subject to a mechanical budget constraint such as $\sum_{i=1}^{p} w_{t,i} = 1$. Instead, both the baseline and enhanced decoders are governed by the strategy-specific exposure constraints and implementation guidelines attached to each family. All reported weight differences should therefore be read as inferred exposure changes rather than as forced budget-neutral reallocations.

Enhanced Target and Enhanced Decoder

Enhanced decoding keeps exactly the same liquid instrument universe and the same Bayesian state-space decoder as the baseline, but it changes the *objective* by decoding an improved version of the benchmark rather than the raw benchmark itself. Concretely, we define

$$\tilde{r}_t \;=\; r_t^{\mathrm{B}} + c, \qquad c \;=\; \frac{x\%}{252},$$

where c is a constant daily enhancement corresponding to a target trend of $x\%$ per annum. The term c is not a new asset and does not change the trading universe: it is simply a constant improvement objective added to the benchmark. The decoder therefore solves the same filtering problem as in the baseline, but against a target that rewards benchmark replication *plus* a steady positive drift. Economically, this creates an internal search for the most efficient risk premia: among all portfolios able to reproduce the benchmark with the same instruments, the enhanced decoder can increase or decrease exposure to sleeves and markets that deliver the benchmark *and* the added annualized drift more efficiently. Because weights are governed by the constraints associated with each strategy rather than by a hard sum-of-weights condition, any cross-market substitution that appears in the results is not mechanically imposed by budget neutrality, although the decoder may still generate such substitutions in practice. The excess portfolio, long Enhanced and short Baseline, should therefore be interpreted as the direct portfolio consequence of that trend-improvement objective.

Excess-Performance Portfolio

The clean empirical object is the excess portfolio

$$r_t^{\mathrm{XS}} = \mathbf{x}_t^\top \left(\mathbf{w}_t^{\mathrm{Enh}} - \mathbf{w}_t^{\mathrm{Base}}\right) = r_t^{\mathrm{Enh}} - r_t^{\mathrm{Base}}. \tag{5}$$

This construction is economically useful for three reasons.

- It uses the *same* implementation universe on both sides of the trade, so the excess return is attributable to decoding rather than to additional assets.
- Common benchmark beta is largely differenced out, making the track record a cleaner measure of incremental alpha.
- Position analysis is directly interpretable: every non-zero excess weight is the incremental exposure chosen by the enhanced decoder relative to the baseline decoder, and these differences need not sum to zero.

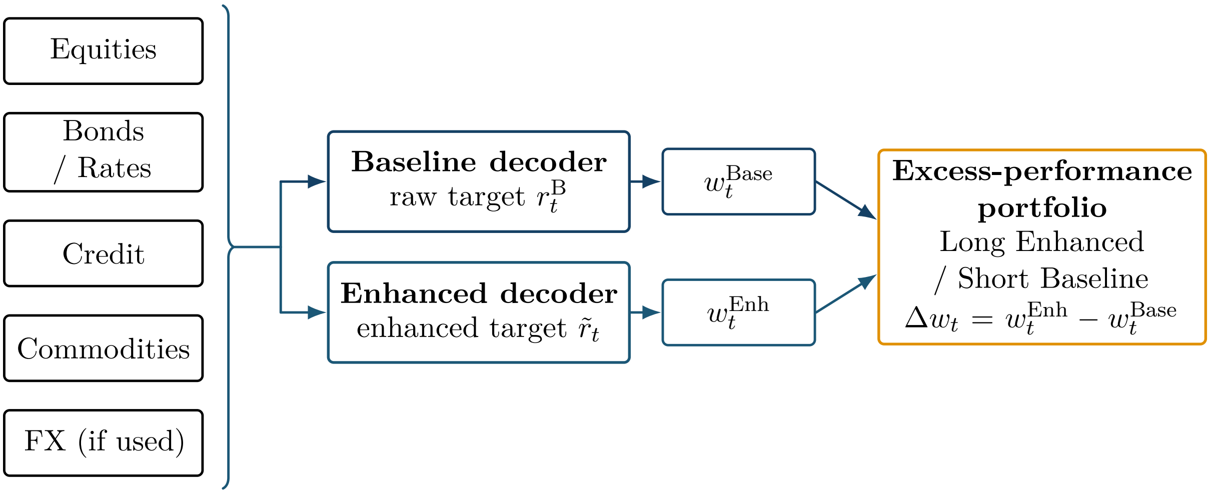

Figure 2: Same liquid sleeves, two targets. The excess-performance portfolio isolates the incremental allocation tilts generated by enhanced decoding.

4 Results

Comparative performance uses the common return window 2012-01-02 to 2026-03-16. Cross-strategy weekly return correlations of the long-enhanced/short-baseline portfolios, as well as each strategy's correlation to U.S. equities, are measured on that same common sample. Position overlap uses daily weights from 2012-01-02 to 2026-03-17. Hedge Funds, Long Short Equity, and Risk Parity have longer standalone histories, but the common window preserves an apples-to-apples comparison. One headline result on this aligned sample is that **Hedge Funds** has the **lowest correlation to U.S. equities at 15.0%**.

Track Records and Full-Sample Performance

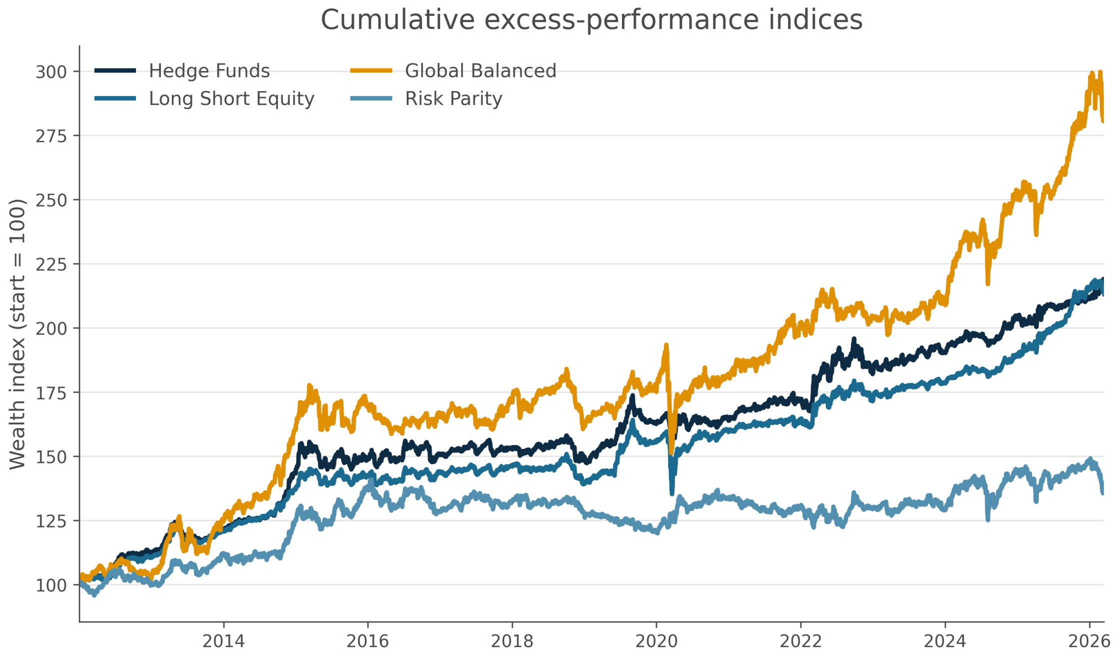

Figure 3: Common-sample cumulative excess-performance indices, normalized to 100 at the first shared observation. Each series is the P&L of the long-enhanced/short-baseline portfolio for the corresponding strategy family.

Table 1: Full-sample performance (common comparison window).

Strategy	Cumulative Return	Annual Return	Vol	Sharpe Ratio	Max DD	Return/Max DD
Hedge Funds	117.7%	5.6%	6.2%	0.91	11.2%	0.50
Long Short Equity	113.2%	5.4%	5.5%	0.99	17.6%	0.31
Global Balanced	182.6%	7.3%	8.8%	0.85	21.8%	0.34
Risk Parity	36.6%	2.2%	7.6%	0.33	14.9%	0.15

Three patterns stand out. First, **Hedge Funds** is the paper's clearest diversification result: its absolute return is close to Long Short Equity, but its drawdown is much shallower, which produces the best return-to-drawdown ratio of the four, while its correlation to U.S. equities is only **15.0%**, the lowest in the paper. Second, **Global Balanced** dominates in absolute compounding: a 100-base index grows to roughly 283 by the end of the sample, supported by the highest annualized return. Third, **Long Short Equity** delivers the cleanest risk-adjusted profile, with the highest Sharpe ratio and materially lower volatility than Global Balanced. **Risk Parity** remains positive over the full sample but is clearly the weakest excess strategy; it compounds less, carries a lower Sharpe ratio, and underperforms sharply in the most recent year. One plausible explanation is structural: the RP implementation is subject to more positive-only asset constraints, which narrows the set of feasible long/short excess bets and may therefore limit risk-adjusted upside relative to the other families.

The main drawdowns of HF, LS, and GB all cluster around the 2020 stress episode, whereas RP experiences a more protracted drawdown profile. This already hints that the enhanced-decoding edge is expressed differently in RP than in the other three families.

Subperiods

Table 2: 10Y: 2016-03-16 to 2026-03-16.

Strategy	Cumulative Return	Annual Return	Vol	Sharpe Ratio	Max DD	Return/Max DD
Hedge Funds	44.4%	3.7%	6.3%	0.60	11.2%	0.33
Long Short Equity	50.7%	4.1%	5.8%	0.73	17.6%	0.23
Global Balanced	71.8%	5.4%	8.6%	0.65	21.8%	0.25
Risk Parity	3.2%	0.3%	7.6%	0.08	13.0%	0.02

Table 3: 5Y: 2021-03-16 to 2026-03-16.

Strategy	Cumulative Return	Annual Return	Vol	Sharpe Ratio	Max DD	Return/Max DD
Hedge Funds	29.2%	5.1%	6.1%	0.85	7.1%	0.73
Long Short Equity	31.6%	5.5%	4.8%	1.13	4.5%	1.23
Global Balanced	51.0%	8.3%	8.2%	1.01	10.4%	0.80
Risk Parity	4.6%	0.9%	8.4%	0.15	12.1%	0.08

Table 4: 3Y: 2023-03-16 to 2026-03-16.

Strategy	Cumulative Return	Annual Return	Vol	Sharpe Ratio	Max DD	Return/Max DD
Hedge Funds	18.2%	5.6%	4.4%	1.27	2.8%	2.02
Long Short Equity	23.3%	7.0%	4.2%	1.64	2.5%	2.82
Global Balanced	42.3%	12.0%	8.5%	1.38	10.4%	1.16
Risk Parity	6.3%	2.1%	8.8%	0.28	12.1%	0.17

Table 5: 1Y: 2025-03-17 to 2026-03-16.

Strategy	Cumulative Return	Annual Return	Vol	Sharpe Ratio	Max DD	Return/Max DD
Hedge Funds	7.9%	7.7%	5.3%	1.44	2.5%	3.07
Long Short Equity	11.2%	10.8%	5.6%	1.85	2.5%	4.35
Global Balanced	12.8%	12.3%	10.0%	1.22	6.7%	1.84
Risk Parity	-3.1%	-3.1%	8.3%	-0.33	9.1%	-0.34

The subperiod picture is remarkably consistent. **Hedge Funds** remains the steadiest diversification result, consistent with its role as the **least equity-linked** family on the full-sample benchmark with only **15.0%** correlation to U.S. equities. Over the last decade, **Global Balanced** remains the strongest source of raw excess return, while **Long Short Equity** is usually the most efficient implementation on a risk-adjusted basis. Over 5Y, 3Y, and 1Y, LS has the best Sharpe ratio and the best return-to-drawdown ratio, whereas GB retains the highest absolute upside. **Risk Parity** does not keep pace and turns negative over the most recent year.

Cross-Strategy Weekly Return Correlations

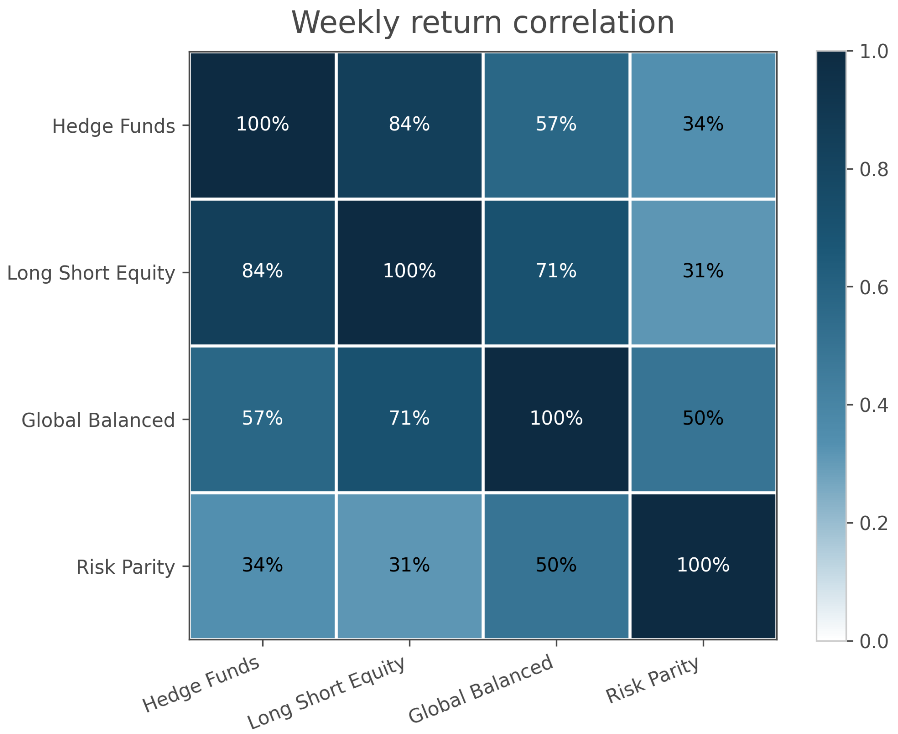

Figure 4: Weekly correlation matrix of excess-performance returns across the four long-enhanced/short-baseline decoding strategies over the common sample.

Weekly return correlations between the four long-enhanced/short-baseline decoding strategies are moderate to high, but the strategies are far from redundant. The strongest link is **HF vs LS at 84%**, which indicates that enhanced decoding extracts a very similar alpha signal from those two alternative-equity-oriented families. That commonality should not be mistaken for the same equity footprint: **Hedge Funds** later proves to be the **least equity-linked** family, with only **15.0%** correlation to U.S. equities. **LS vs GB at 71%** and **HF vs GB at 57%** still show material commonality, but now with more family-specific dispersion. **Risk Parity** is the clear outlier: its cross-strategy weekly return correlation ranges only from **31% to 50%** versus the other three strategies. In other words, enhanced decoding adds value in RP as well, but through a different macro channel than in HF, LS, and GB.

Table 6: Correlation of weekly excess returns to U.S. equities (S&P 500 Total Return) over the common sample. Hedge Funds is the lowest at 15.0%.

Strategy	Correlation to U.S. equities
Hedge Funds	**15.0%**
Long Short Equity	41.6%
Global Balanced	50.6%
Risk Parity	30.2%

Correlation to U.S. equities adds a useful benchmark lens. **Hedge Funds** is the result to emphasize: its excess strategy shows only a **15.0%** correlation to U.S. equities, the **lowest** of the four families. This places HF well below **Risk Parity** at **30.2%**, **Long Short Equity** at **41.6%**, and **Global Balanced** at **50.6%**. That ranking matters because it shows that the hedge-fund enhancement signal is the least tied to broad equity beta and is instead carried more through a diversified mix of bonds, credit, FX, and selected equity shorts. **Long Short Equity** remains notable for a different reason: it still delivers the highest Sharpe ratio while keeping equity correlation materially below **Global Balanced**. These numbers should be read together with the strategy-specific implementation limits discussed above: baseline and enhanced are compared within the same broad leverage, equity, bond, commodity, and dollar framework for each family, so the correlation differences are best understood as coming from a different use of a similar feasible set rather than from a materially different constraint set.

Average Excess Positions by Category

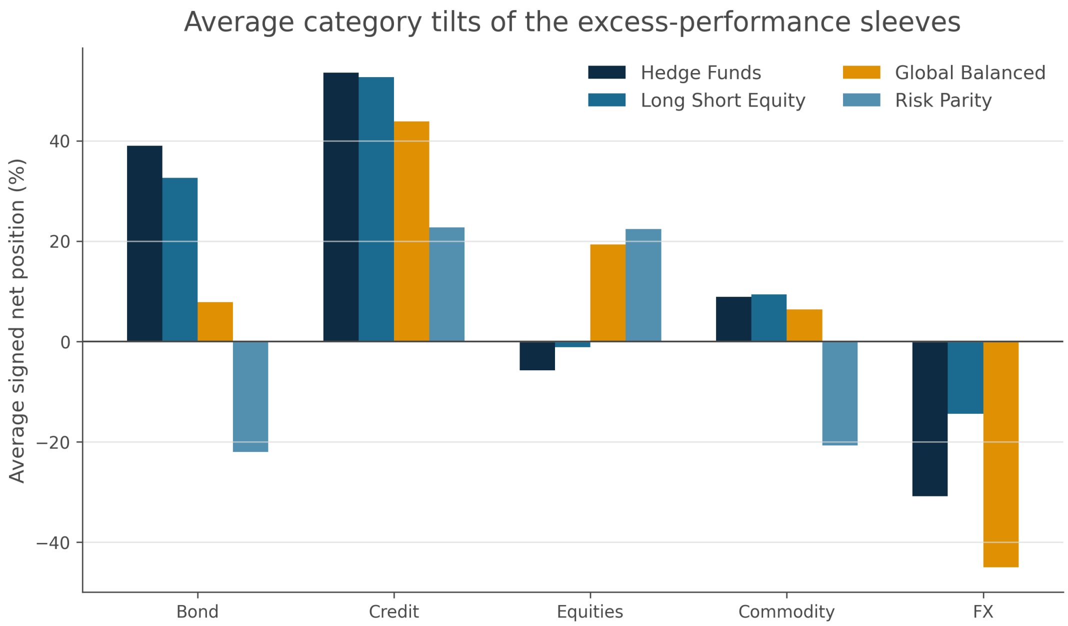

Figure 5: Average signed excess positions by category over the common position sample. Positive values indicate markets that the enhanced decoder tends to overweight relative to the baseline decoder.

The category averages are economically structured rather than arbitrary residuals. **HF** and **LS** are structurally *long bonds and credit*, mildly long commodities, lightly short equities, and short FX. **GB** remains long credit as well, but adds a more directional *long-equity / short-FX* expression. **RP** is fundamentally different: its average excess trade is *short bonds and commodities* and *long equities and credit*. For **HF** in particular, this bond/credit-heavy and lightly short-equity mix is consistent with the paper's headline diversification result: the **lowest correlation to U.S. equities at 15.0%**. More generally, excess tilts are often inversely related to baseline positioning: sleeves with elevated baseline weights tend to be reduced by enhancement, while lightly represented sleeves tend to be increased. This already indicates that enhanced decoding modifies the risk-parity book through a different macro balance than in the other three families.

Daily Position Correlations by Category

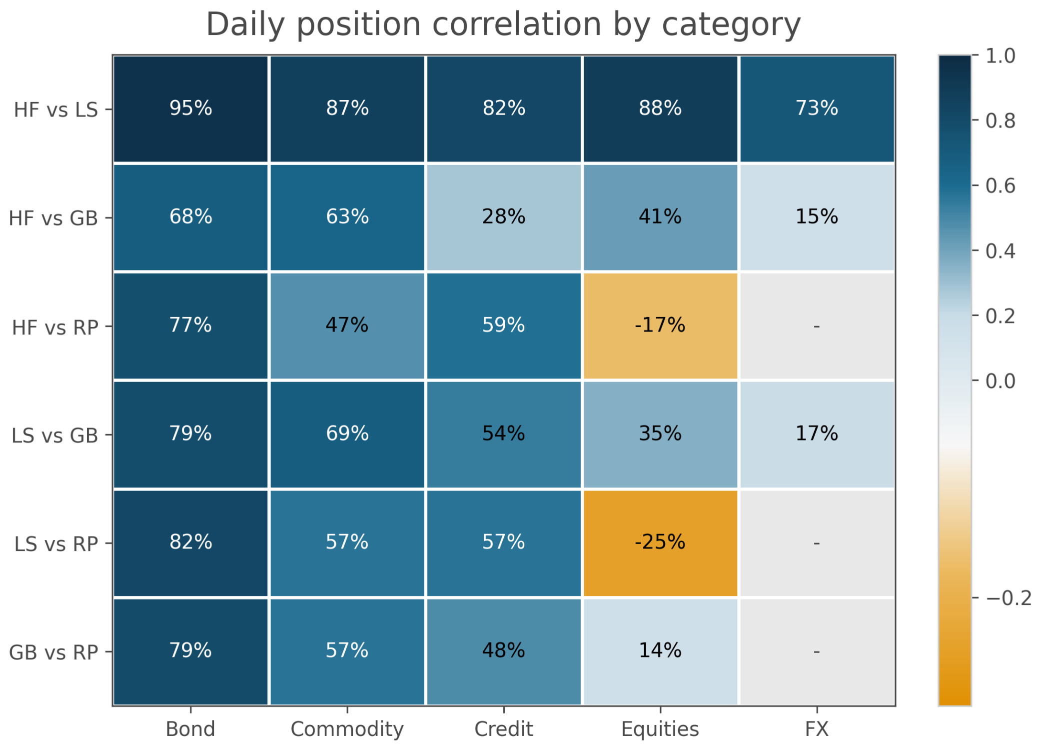

Figure 6: Pairwise daily position correlation by category, computed on excess weights.

Category-level overlap is strongest between **HF and LS**, which align across every sleeve, especially bonds (95%), commodities (87%), credit (82%), and equities (88%). More broadly, **bond exposures are the common macro backbone** of the excess portfolios: even the least similar pairs still show roughly 69–82% correlation in the bond sleeve. The main divergence sits in **equities**. HF versus RP is **-17%** and LS versus RP is **-25%**, meaning that the RP excess strategy typically expresses the opposite equity exposure change from HF and LS. FX overlap is weakest between GB and the HF/LS pair (14–17%), which is consistent with GB's more persistent short-FX stance.

Position Correlations by Specific Market

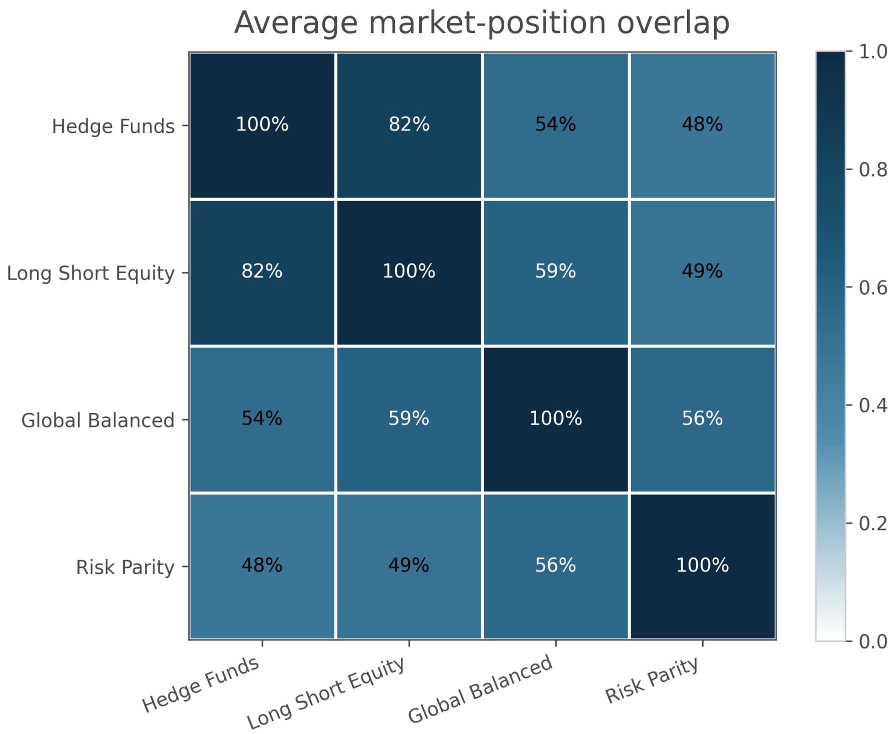

Figure 7: Average market-position overlap, measured as the average pairwise daily position correlation across specific markets shared by each strategy pair.

Across all common instruments, **HF and LS** show the highest average market-position overlap (about 82%). **GB** has moderate overlap with HF and LS (roughly 54–59%), while **RP** is the least similar to HF and LS (48–49%) but somewhat closer to GB (56%). This already suggests that RP's excess performance is driven less by alternative-beta tilts and more by distinct macro exposure changes.

Targeted Trend-Shock Experiments on Global Balanced

To make the mechanism more explicit, we run two targeted experiments on the **Global Balanced** enhanced decoder. In each case, one market receives a constant +10% per annum trend bump over calendar year 2016,

$$r'_t = r_t + \frac{10\%}{252},$$

for trading days from 2016-01-04 to 2016-12-30, while the rest of the universe is left unchanged. We then compare the resulting enhanced-model weights with the original enhanced weights. The two perturbations are chosen to test *risk-premium substitution across similar assets*: when one asset is given an artificial trend premium, the decoder does not simply add directional exposure, but reallocates part of the excess position against a close substitute and effectively arbitrages trends within the same sleeve. The two examples are **EUR/USD versus GBP/USD** in developed FX, and **Russell 2000 versus the S&P 500 sleeve** in U.S. equities.

Because the experiment is intentionally local, performance is summarized by the cumulative contribution of the local spread trade on the shocked market and its closest substitute, rather than by a full-portfolio counterfactual. This is enough to distinguish three effects: *instantaneous sign reaction*, *partial substitution toward or away from the nearest peer*, and *progressive mean reversion once the temporary bump is no longer refreshed.*

Table 7: Summary of targeted 2016 trend shocks on Global Balanced.

Experiment	Peak shocked-asset tilt	2016 peer hedge ratio	2016 local pair contribution	Peak local contribution	Half-life of shocked tilt
EUR/USD bump vs. GBP/USD	+75.3 pp (2017-08-17)	21.8%	0.15%	8.8% (2018-01-25)	2018-12-14
Russell 2000 bump vs. S&P 500 sleeve	+11.2 pp (2017-01-11)	52.7%	1.8%	2.0% (2017-01-04)	2017-11-24

The hedge ratio is the average fraction of the shocked tilt that is matched by an opposite move in the nearest peer during 2016, computed as $-\Delta w^{\mathrm{peer}}/\Delta w^{\mathrm{shock}}$ on days with a nonzero shocked tilt. It is a descriptive co-movement statistic, not the consequence of a budget constraint. The two rows already show a clear difference in transmission speed. The FX perturbation is more persistent and reaches its largest effect only *after* the bump window has ended, whereas the Russell perturbation peaks almost immediately after year-end and then fades much faster.

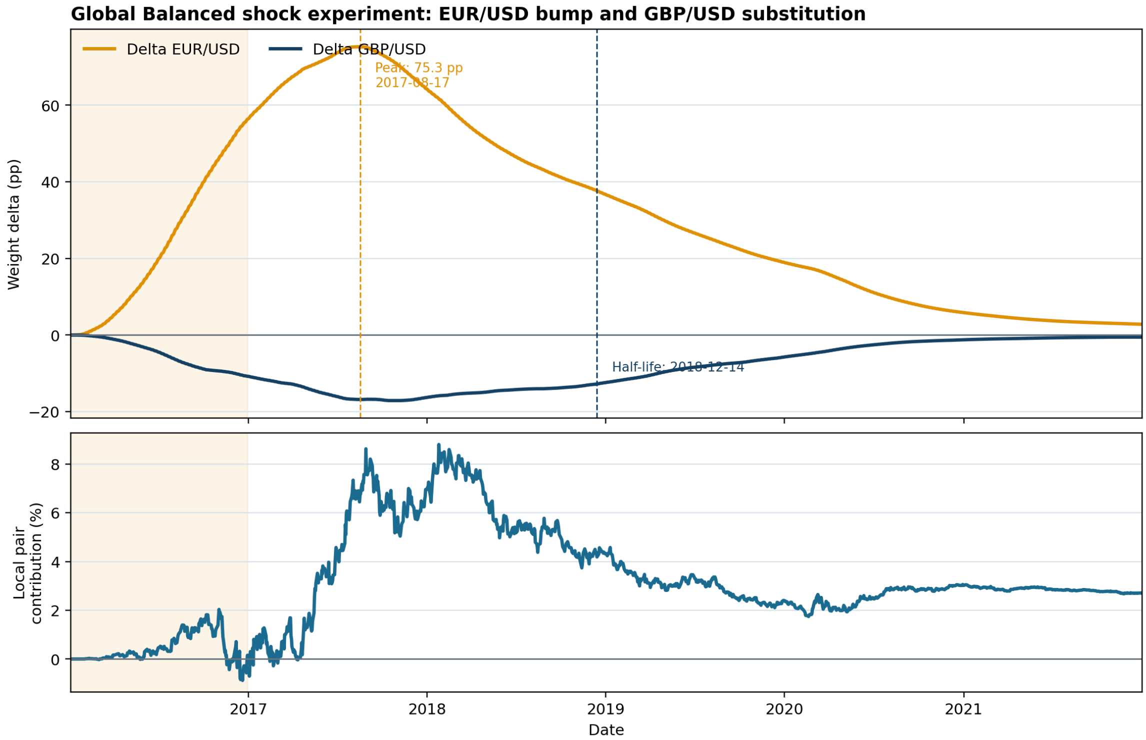

Figure 8: Global Balanced response to a constant +10% p.a. EUR/USD trend bump in 2016. The shaded band marks the bump window. Top panel: weight differences between the bumped enhanced model and the original enhanced model, in percentage points. Bottom panel: cumulative contribution of the local EUR/USD versus GBP/USD spread trade.

The **EUR/USD experiment** shows an immediate reaction in *sign*, but not in full magnitude, because the constant daily bump must pass through the state-space filter and therefore builds up over time. By 2016-02-01, the enhanced model is already **0.33 pp less short EUR/USD** and **0.07 pp more short GBP/USD**. By 2016-12-30, those numbers reach **+56.3 pp** on EUR/USD and **-10.8 pp** on GBP/USD, and the EUR tilt continues to build until **2017-08-17**, where it peaks at **+75.3 pp**. The key economic point is that the model does *not* interpret the higher EUR trend as a generic risk-on signal. Instead, it treats the shock as a relative-value opportunity inside developed FX: part of the added EUR premium is substituted against GBP/USD and other nearby markets, so the decoder is effectively arbitraging trends across similar assets rather than simply levering up one directional bet. Over 2016, **GBP/USD offsets 21.8%** of the shocked EUR tilt on average, while JPY and Euro-credit also move in the opposite direction. After the bump year, the model progressively comes back toward the original enhanced weights: the shocked EUR tilt is halved by **2018-12-14** and is close to zero by 2021. The local EUR/GBP contribution is almost flat through 2016 (**+0.15%**), then rises sharply to **+8.8%** by 2018-01-25 as the post-shock overweight continues to monetize before gradually decaying.

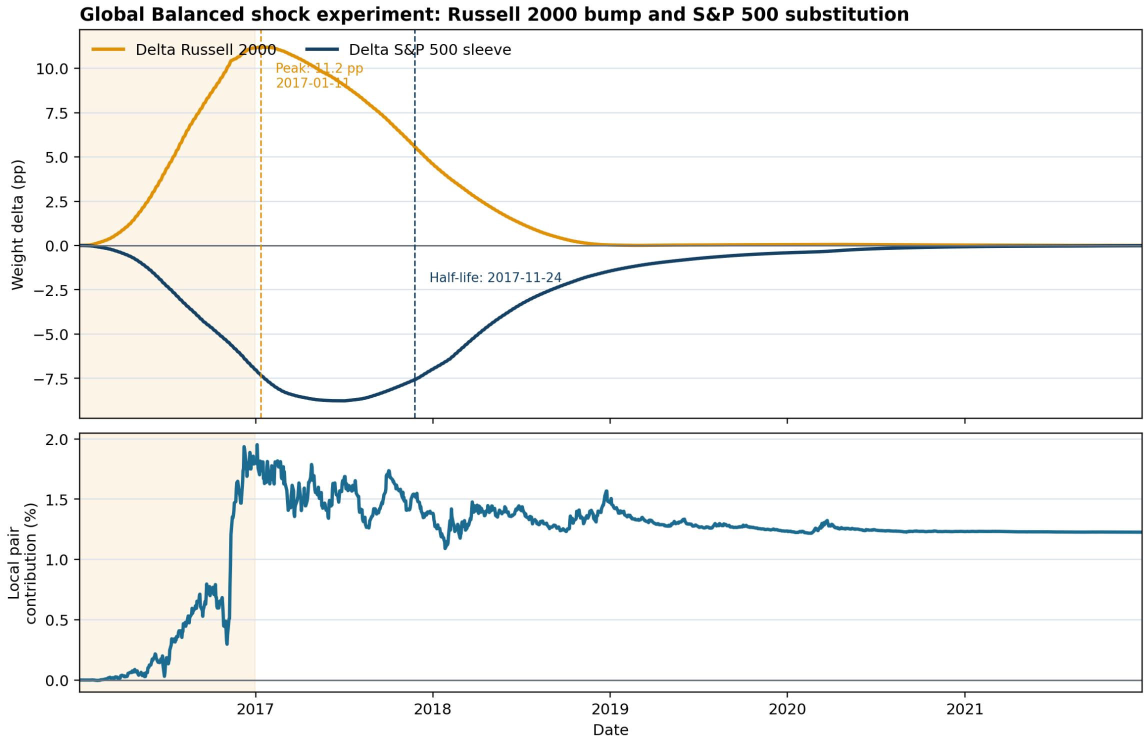

Figure 9: Global Balanced response to a constant +10% p.a. Russell 2000 trend bump in 2016. The shaded band marks the bump window. Top panel: weight differences between the bumped enhanced model and the original enhanced model, in percentage points. Bottom panel: cumulative contribution of the local Russell 2000 versus S&P 500 sleeve spread trade.

The **Russell 2000 experiment** is cleaner and faster. By 2016-02-01, the model has already started to respond within U.S. equities, with a **+0.06 pp** increase in Russell 2000 versus a **-0.03 pp** cut in the S&P 500 sleeve. By 2016-12-30, the differential reaches **+11.2 pp** on Russell 2000 and **-7.0 pp** on the S&P 500 sleeve. Roughly **52.7%** of the shocked Russell tilt is matched by an opposite move in that sleeve during 2016, which is exactly the footprint of *intra-equity risk-premium substitution*: the decoder reallocates the excess trade from one equity premium to a close alternative instead of simply adding more aggregate equity beta. The relative spread peaks on **2017-01-11** and then mean-reverts much faster than in FX, with a half-life reached by **2017-11-24**. The associated local Russell/S&P contribution reaches **+1.8%** by end-2016 and peaks at about **+2.0%** in early January 2017 before slowly fading.

Taken together, the two experiments clarify the economic intuition of the trend-shock test. A temporary increase in the trend premium of one asset tends to trigger *risk-premium substitution* into or against a close alternative asset: the decoder arbitrages trends across similar markets rather than merely adding standalone exposure to the shocked instrument. That substitution is empirical rather than imposed by a sum-of-weights constraint. The speed of normalization then depends on how specific the premium is. The Russell shock behaves like a local intra-equity opportunity and is absorbed relatively quickly, whereas the EUR shock propagates through a broader developed-FX and macro-credit complex and therefore decays much more slowly.

5 Conclusion

- **Profile.** Over the common sample, *Hedge Funds* is the paper's clearest diversification result: it provides the best return-to-drawdown profile and the **lowest correlation to U.S. equities at 15.0%**. *Global Balanced* still delivers the highest absolute excess return, and *Long Short Equity* achieves the highest Sharpe ratio.
- **Common core.** HF, LS, and GB share excess tilts mainly in bonds, credit, and developed FX, which explains why their track records remain positively correlated.
- **Risk Parity.** RP is the main exception: its excess track record is weaker and less correlated because enhancement changes a different mix of duration, commodities, and equity exposure there. Tighter positive-only constraints may also reduce the feasible long/short opportunity set.
- **Shock tests.** In Global Balanced, a trend shock triggers risk-premium substitution across similar assets: a EUR/USD bump is absorbed mainly through GBP/USD and normalizes slowly, while a Russell 2000 bump is offset mostly through the S&P 500 sleeve and mean-reverts faster.
- **Takeaway.** The portfolio **long Enhanced / short Baseline** is a useful diagnostic tool: it isolates where the enhanced target changes the inferred portfolio and makes decoding alpha interpretable by sleeve and by market. More importantly, the evidence suggests that enhancement is most worthwhile when it starts from an already efficient and balanced decoded strategy such as *Hedge Funds*, rather than from a more structurally biased starting point such as *Risk Parity*.

Disclaimer

References

Ai For Alpha Team. *Strategy Spotlight: How Can Decoding Enhance Risk Parity?* October 2025.

Eric Benhamou, Jean-Jacques Ohana, and Beatrice Guez. *Generative AI: Crafting Portfolios Tailored to Investor Preferences.* SSRN Working Paper 4780034, April 2024.

Chang-Jin Kim and Charles Nelson. *State-Space Models with Regime Switching.* MIT Press, 1999.

Jean-Jacques Ohana, Eric Benhamou, David Saltiel, and Beatrice Guez. *Deep Decoding of Strategies.* Universite Paris-Dauphine Research Paper No. 4128693, June 2022.

Mike West and Jeff Harrison. *Bayesian Forecasting and Dynamic Models.* Springer, second edition, 1997.