Ai For Alpha

Ai-Powered Investment Replication

Strategy Spotlight:
How Can Decoding Enhance Risk Parity?

October 2025, Ai For Alpha Team

Abstract

We study a decoding framework that replicates and enhances multi-asset risk parity using only liquid instruments. The approach is implemented in two steps. **Step 1: Enhanced Returns.** We decode the benchmark to an *enhanced* target that preserves scale but improves signal-to-noise, allowing modest, data-driven allocation shifts while maintaining the benchmark's risk footprint. **Step 2: Enhanced Returns with Drawdown (DD) Mitigation.** We keep the same enhanced target, apply an asymmetric transformation that attenuates negative shocks, and activate a CTA Risk-Off overlay to mitigate drawdowns.

Overall, a Bayesian graphical decoder with shrinkage and sequential learning can reproduce risk-parity behavior and deliver better risk-adjusted outcomes through (i) enhanced-return targeting and (ii) a drawdown-mitigating overlay.

1 Introduction

Motivations

The idea of *risk-parity*—allocating portfolio risk equally across asset classes—has profoundly influenced institutional asset allocation. Its conceptual origins trace back to Bridgewater Associates' *All Weather* strategy launched in 1996, though the term itself was popularized by Qian (2005) and formalized by Maillard et al. (2010) and Roncalli (2013). Risk parity addresses the imbalance of risk contributions inherent in traditional capital-weighted or equal-weighted portfolios, emphasizing diversification in risk rather than in capital. The approach gained widespread attention after the 2008 financial crisis, when investors sought allocations that could withstand multiple macroeconomic regimes; see Choueifaty and Coignard (2008), Chaves et al. (2011), Lee (2011), and Asness et al. (2012).

While risk parity has proven robust as an asset-allocation principle, transparent and low-cost replication is feasible through liquid, rules-based implementations (e.g., S&P Risk Parity 10% Target Volatility). The objective of the paper is close tracking with potential for added value—namely, incremental excess return and improved drawdown mitigation during stressed regimes (e.g., 2018, 2020, and especially 2022). Simple linear-regression replicators offer parsimony but are often unstable and turnover-intensive, motivating more structured, model-based approaches.

We propose a *Bayesian graphical replication framework* that addresses these shortcomings by pooling information across factor models in a structured probabilistic setting. Each "sleeve" (equity, bond, credit, commodity, and a managed-futures risk-off sleeve) evolves through time as a latent state connected via cross-sleeve dependencies. This design allows the model to infer the dynamic structure of diversification, adapt exposures in real time, and exploit asymmetric factor behavior — particularly in crisis periods (see also Benhamou et al., 2024).

Structure of the Paper

The paper is organized as follows. Section 2 presents the Bayesian graphical model that captures cross-asset dependencies and latent sleeve dynamics. Section 3 states the *Baseline*. Section 4 introduces *Enhanced*. Section 5 presents *Enhanced with DD Mitigation*. Section 6 reports the results across horizons, and Section 7 concludes.

2 Methodology

Graphical Model

We model the benchmark's daily excess return r_t with a time–varying linear specification based on four observable asset returns. Let

$$\mathbf{x}_t = \left(r_{t,1},\, r_{t,2},\, r_{t,3},\, r_{t,4}\right)^\top.$$

The observation equation is

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

where $\boldsymbol{\beta}_t \in \mathbb{R}^4$ are time–varying coefficients. Coefficients evolve smoothly via 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.

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

$$\widehat{\boldsymbol{\beta}}_t \;=\; \arg\min_{\boldsymbol{\beta}\in\mathbb{R}^4} \left\{ \sum_{\tau=1}^{t} \delta^{\,t-\tau} \left(r_\tau - \mathbf{x}_\tau^\top \boldsymbol{\beta} \right)^2 \;+\; \left(\boldsymbol{\beta} - \boldsymbol{\beta}_{t-1} \right)^\top \mathbf{W}_t^{-1} \left(\boldsymbol{\beta} - \boldsymbol{\beta}_{t-1} \right) \right\}, \qquad \delta \in (0,1], \quad (3)$$

which combines an exponentially weighted prediction error with a state–smoothness penalty implied by (2). Standard forward–filtering updates implement this objective using only information available at time t (West and Harrison, 1997; Kim and Nelson, 1999; Carvalho et al., 2009; Koop and Korobilis, 2013; Benhamou et al., 2024).

Figure 1 displays the benchmark on the *top row* and four illustrative asset returns beneath. Blue horizontal links show the temporal evolution of each observed series. At every date t, orange dotted arrows illustrate all contemporaneous relationships: (i) among the asset returns $\{r_{t,1}, \ldots, r_{t,4}\}$ and (ii) between the benchmark r_t and each asset. The mapping from assets to the benchmark is captured by the time–varying coefficients in (1)–(2) and optimized via (3).

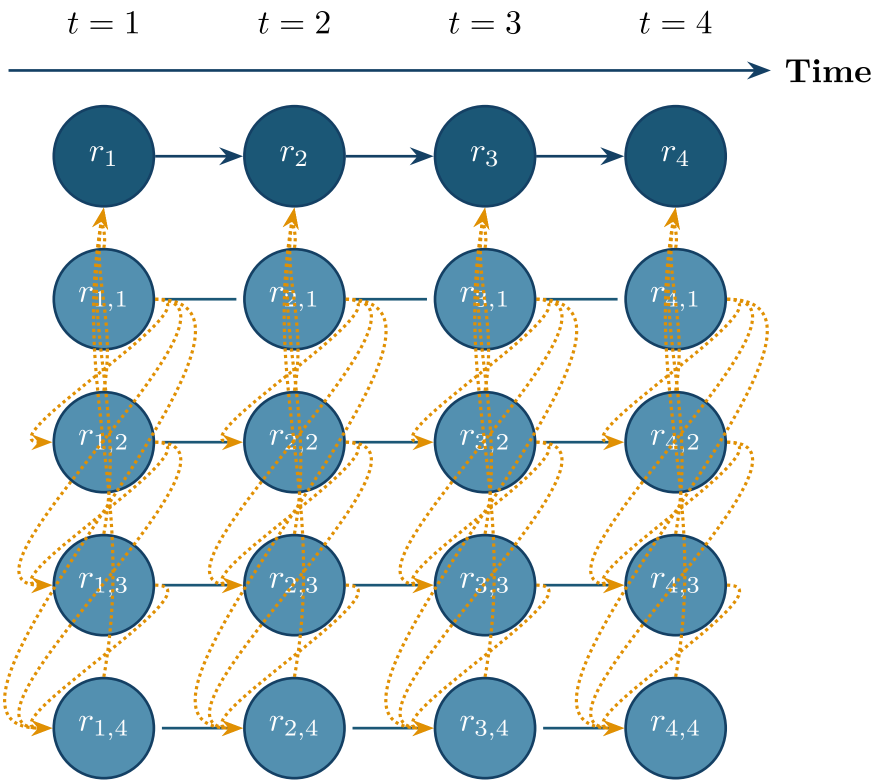

Figure 1: Graphical model with benchmark (top row) and four asset returns (rows below). Blue horizontal links show time evolution of each observed series. At each t, orange dotted arrows depict all contemporaneous relationships: among the four assets and between the benchmark r_t and each asset $r_{t,i}$. No vertical loading arrows are drawn; the regression mapping is governed by (1)–(2) and optimized via (3).

Investment Universe and Portfolio Sleeves

We base our analysis on a diversified multi-asset universe that mirrors the core structure of institutional risk-parity portfolios. The portfolio includes four traditional sleeves—Equities, Credit, Bonds, and Commodities—along with a fifth, *CTA Risk-Off*, which captures trend-following exposures designed to mitigate equity drawdowns (Bartholomew, 2025). This setup reflects the common balance between growth-oriented and defensive components found in diversified institutional portfolios (e.g., Asness et al., 2012; Ang, 2012; Roncalli, 2013). The fifth sleeve, *CTA Risk-Off*, is **optional**: it is **inactive** in **Baseline Decoding** and **Enhanced** ($w_t^{\text{Hedge}} \equiv 0$) and **active** in **Enhanced with DD Mitigation** as a dynamic overlay.

Each sleeve represents a distinct source of global risk premia. Equities and bonds form the traditional diversification core; credit indices capture exposure to corporate spreads; commodities provide an inflation-sensitive component; and the CTA Risk-Off sleeve adds a systematic hedge through time-series momentum signals (Bartholomew, 2025; Ai For Alpha, 2024). Together, these elements create a transparent and practical representation of the building blocks that underpin most institutional risk-parity strategies.

Equity markets	U.S. broad equities; U.S. technology; Japan; Euro area; U.K.; emerging markets
Credit (short CDS)	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
Commodity markets	Gold; Brent crude; Copper
CTA Risk-Off	Diversified time-series momentum sleeve across 26 liquid futures (equity indices, rates, FX, and commodities); positions netted against other sleeves for implementation efficiency

The addition of the CTA Risk-Off sleeve extends the traditional risk-parity framework by introducing an adaptive, crisis-responsive component that reflects the asymmetric nature of trend-following strategies documented in Moskowitz et al. (2012). In practice, this sleeve trades a diversified set of futures contracts using volatility scaling and cross-asset netting to maintain efficient execution and balanced risk. This expanded universe allows the graphical model to capture both cyclical and defensive dynamics within a unified, data-driven structure.

CTA Risk-Off Sleeve

To complement the four traditional risk-parity sleeves, we optionally introduce a dedicated *CTA Risk-Off* component that captures systematic trend-following behavior across major global futures markets while defending against equity drawdowns (Bartholomew, 2025; Ai For Alpha, 2024). The sleeve follows a time-series momentum approach across 26 highly liquid contracts spanning equities, sovereign rates, foreign exchange, and commodities. Following the framework of Greyserman and Kaminski (2014), directional exposures are inferred from recent price trends and scaled by realized volatility to maintain stable risk contributions over time.

Signals are generated using multiple lookback horizons and volatility-adjusted position sizing, subject to contract-level exposure limits. Positions switch between long and short depending on the sign of the prevailing trend and its conditional correlation with global equity returns, allowing the sleeve to strengthen its defensive posture during equity market stress. The resulting positions are aggregated into a single sleeve return, $r_t^{\text{CTA-RO}}$, which enters the portfolio allocation as an additional latent factor.

Operationally, the CTA Risk-Off sleeve is integrated with the broader portfolio architecture. When overlapping exposures occur with other sleeves—such as S&P 500, Bund, WTI, or Copper futures—positions are netted before execution to ensure a unified order per instrument. This approach maintains both computational and economic consistency across sleeves while reducing transaction costs and slippage, consistent with best practices in cross-asset implementation (e.g., Harvey et al., 2016; Koijen et al., 2018).

From an economic perspective, the CTA Risk-Off sleeve functions as an adaptive hedge that systematically loads negatively on equity-like and other procyclical risk exposures during adverse market conditions. Historically, it exhibits negative correlation with equities, credit, and cyclical commodities (such as oil and copper), while remaining positively aligned with duration-sensitive assets. Within the proposed graphical model, this sleeve acts as a dynamic shock absorber: its activation improves crisis resilience by providing convexity to portfolio returns without the need for discretionary overlays. This mechanism formalizes the crisis-alpha properties of managed futures within a probabilistic Bayesian framework, linking macro-hedging intuition to measurable latent exposures.

3 Baseline Decoding

We begin by analyzing a baseline decoding of the risk-parity allocation that *excludes* the CTA Risk-Off sleeve while keeping the Target Benchmark unchanged. In this specification, the sleeve set reduces to $\{\mathrm{Eq}, \mathrm{Cr}, \mathrm{Bd}, \mathrm{FX}, \mathrm{Cmd}\}$—representing equities, credit, bonds, FX and commodities—and follows the observation model

$$r_t = \sum_i w_{t,i}\, r_{t,i} + \varepsilon_t, \qquad \mathbf{w}_t = \mathbf{w}_{t-1} + \boldsymbol{\eta}_t, \tag{4}$$

where $\varepsilon_t \sim \mathcal{N}(0, \sigma_\varepsilon^2)$ and $\boldsymbol{\eta}_t$ represents weight innovations drawn from a Gaussian transition process. The model is estimated sequentially, conditioning only on information available at time t as introduced in Benhamou et al. (2024); Ohana et al. (2022).

Table 1 reports the average inferred weights by asset class for the baseline decoding model. Consistent with standard risk-parity intuition (e.g., Maillard et al., 2010; Roncalli, 2013), the model allocates disproportionately more capital to low-volatility assets—particularly government bonds—while assigning smaller shares to riskier sleeves such as equities, commodities, and credit. This allocation pattern reflects the leverage typically required to equalize marginal risk contributions across asset classes under a fixed 10% volatility target.

Table 1: Baseline: Average portfolio weight by asset class.

Asset class	Average weight
Bonds	168%
Equities	39%
Commodities	26%
Credit	6%
FX vs. USD	-3%

The estimated sleeve exposures highlight how the decoding model internalizes the covariance structure among asset returns without requiring explicit volatility targeting or ex ante constraints. In particular, the large allocation to bonds offsets equity risk during periods of elevated market volatility, while commodities and credit provide marginal diversification benefits. This pattern confirms that the model's inferred weights are consistent with the theoretical structure of risk budgeting and the empirical behavior of balanced portfolios documented in Asness et al. (2012) and Ang (2012).

Historical Allocation

Figure 2 presents the evolution of decoded portfolio allocations over time. The series reflects raw capital weights, inclusive of leverage, across the four baseline sleeves. Consistent with the risk-parity framework, bond exposures dominate the capital mix to balance risk contributions, while equity, commodity, and credit allocations fluctuate with changes in relative volatility and cross-asset correlations.

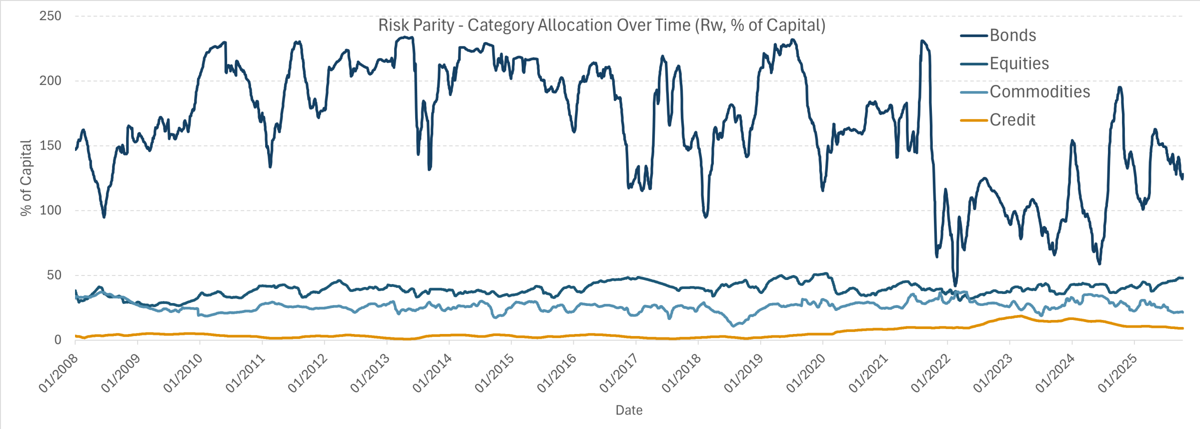

Figure 2: Historical decoded allocations across asset classes (% of capital, including leverage).

4 Methodology — Enhanced Risk Parity

Goal. Replicate the benchmark's risk structure while allowing measured allocation shifts that improve efficiency—without any explicit hedge sleeve.

Decoder. We model the benchmark excess return r_t^{B} as a noisy combination of sleeve instrument returns with time-varying weights updated in a state-space (random-walk) fashion. Estimation is forward-only with shrinkage.

Enhanced target — construction and rationale. We replace the raw benchmark return r_t^{B} with a variance-preserving, denoised target $\tilde{r}_t$ that suppresses transient noise but keeps scale. This lets the decoder learn persistent sleeve structure with less measurement error, yielding stabler weights and higher expected Sharpe without drifting away from the benchmark's risk footprint.

Why this improves the decoder. In the observation equation $r_t^{\mathrm{B}} = X_t^{\top} w_t + \varepsilon_t$, replacing the raw target r_t^{B} with the variance-preserving, denoised $\tilde{r}_t$ lowers effective measurement noise. In a linear-Gaussian state-space, this raises the Kalman gain and reduces the posterior variance of w_t, yielding steadier exposures and lower turnover at the same tracking fidelity. Conceptually, it is the target-side analogue of covariance shrinkage and hierarchical priors: a small, controlled bias that buys a large variance reduction and less overfitting.[1]

Why a portfolio close to the baseline can achieve higher expected return.

- **Estimation efficiency and noise filtering.** Denoising the target and using smooth, cross-sleeve priors reduce the impact of transient noise, so updates load more on persistent premia and less on randomness. The portfolio remains highly correlated to the benchmark but aligns more with systematic drivers that typically carry higher long-run return per unit of risk.
- **Implicit risk-budget optimization.** With time-varying covariances, the latent-state structure reweights sleeves to keep behaviour close to the benchmark while improving marginal risk–return (e.g., leaning slightly toward the sleeve with the best conditional Sharpe in the current regime), keeping tracking error modest.

[1]See, for example, the shrinkage and denoising approaches of Ledoit and Wolf (2004), and the dynamic learning frameworks discussed by Carvalho et al. (2009) and Koop and Korobilis (2013).

- **Cross-sectional information sharing.** By pooling information across sleeves, the decoder infers common factors; when the benchmark allocates mechanically, small probabilistic tilts (e.g., a bit more equity when bond volatility falls) raise expected return without materially changing total portfolio risk.

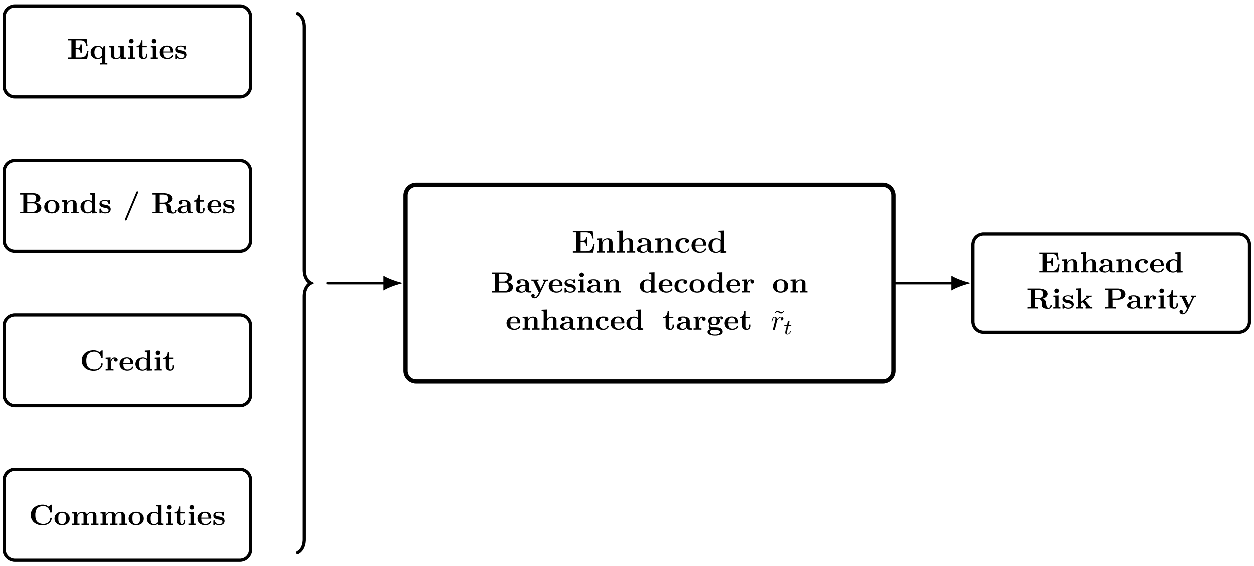

Figure 3: Flow: sleeves (grouped) → Augmented Decoding → Enhanced Risk Parity.

5 Methodology — Enhanced Risk Parity with DD Mitigation

Goal. Keep the upside gained in the no-hedge step, add downside convexity.

Asymmetric mapping. We use the same enhanced target $\tilde{r}_t$ but apply an asymmetric transform that attenuates negative shocks. This changes the loss the decoder optimizes—penalizing downside tracking errors more than upside ones.

Hedge sleeve. We activate a CTA-style managed-futures overlay (time-series momentum across major equity, rate, FX, and commodity futures) with volatility scaling and contract limits. The hedge tends to be negatively correlated to procyclical risk during sell-offs and slightly duration-positive. Exposures that overlap with core sleeves are netted at execution to keep a single order per instrument. The combined portfolio is our **Enhanced Risk Parity with DD Mitigation**.

Why the hedge improves outcomes.

- **Left-tail protection, better compounding.** A negative conditional beta in sell-offs reduces drawdown depth and time under water, lifting geometric returns for a similar volatility budget.
- **Conditional and risk-budgeted.** Activation is gated (trend breadth, elevated volatility, benchmark drawdown) and sleeve risk is capped; in normal regimes exposure stays small, preserving high correlation and modest tracking error.
- **Cross-asset convexity, clean implementation.** Dynamic crisis tilts—short equities/commodities, long duration—add convexity; systematic netting versus core sleeves avoids double-counting and contains turnover/costs.

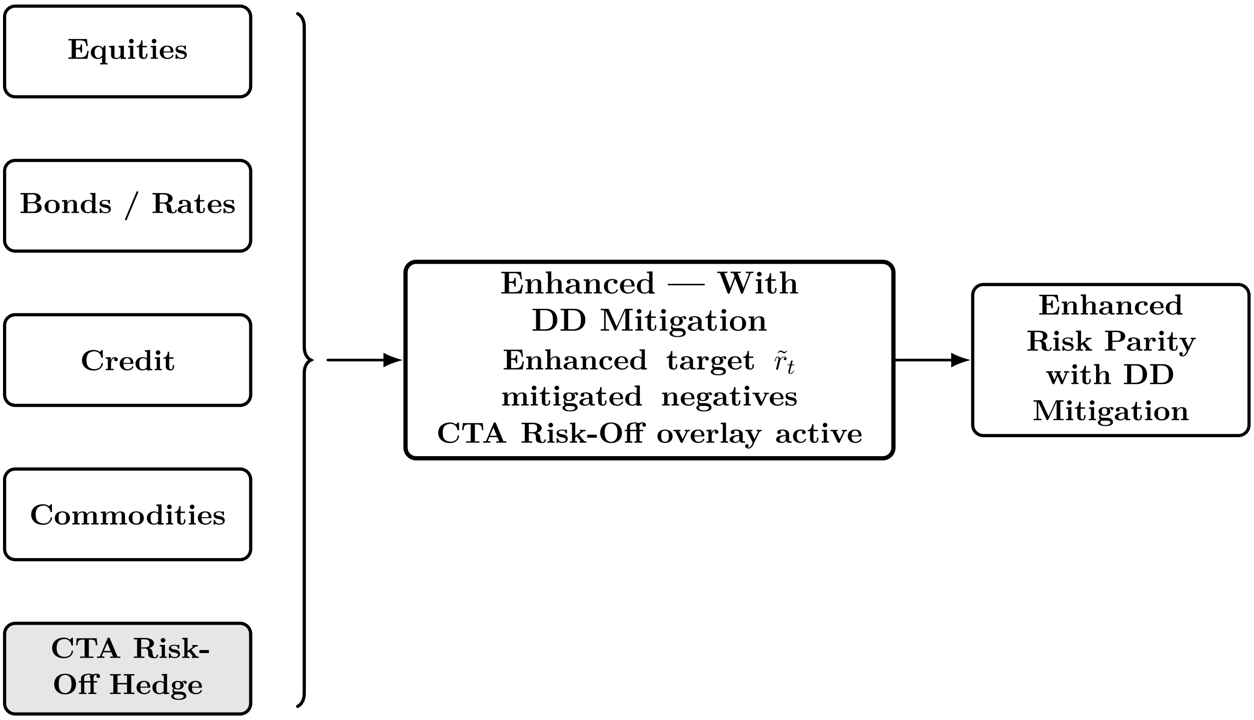

Figure 4: Flow: sleeves (grouped) → Enhanced Decoding with asymmetric returns → *Enhanced Risk Parity with DD Mitigation*. The CTA Risk-Off sleeve is enabled.

6 Results

We now evaluate two variants—*Enhanced* and *Enhanced with DD Mitigation*—against the *Risk-Parity Benchmark*. We report annualized return, volatility, Sharpe ratio, maximum drawdown, and return-to-drawdown for the full sample and for standard subperiods. All figures are net of fees and estimated implementation costs.

Full Sample (2008-01-04 to 2025-10-14)

Table 2: Full-sample performance.

Metric	Enhanced	Enhanced with DD Mitigation	Risk-Parity Benchmark
Annual Return	8.9%	10.5%	7.4%
Vol	9.8%	9.3%	10.7%
Sharpe Ratio	0.76	0.97	0.56
Max DD	30.9%	23.5%	33.5%
Return/maxDD	0.29	0.45	0.22

Subperiods

Table 3: 10Y: 2015-10-15 to 2025-10-14.

Metric	Enhanced	Enhanced with DD Mitigation	Benchmark
Annual Return	8.9%	9.3%	7.5%
Vol	9.7%	9.5%	10.4%
Sharpe Ratio	0.68	0.74	0.51
Max DD	22.2%	19.2%	24.4%
Return/maxDD	0.40	0.49	0.31

Table 4: 5Y: 2020-10-15 to 2025-10-14.

Metric	Enhanced	Enhanced with DD Mitigation	Benchmark
Annual Return	7.9%	8.3%	7.6%
Vol	10.0%	10.1%	11.1%
Sharpe Ratio	0.46	0.50	0.39
Max DD	22.2%	18.6%	23.8%
Return/maxDD	0.35	0.45	0.32

Table 5: 3Y: 2022-10-17 to 2025-10-14.

Metric	Enhanced	Enhanced with DD Mitigation	Benchmark
Annual Return	17.3%	16.1%	12.4%
Vol	9.9%	10.1%	10.3%
Sharpe Ratio	1.19	1.05	0.69
Max DD	11.8%	11.9%	8.8%
Return/maxDD	1.46	1.35	1.40

Table 6: 1Y: 2024-10-15 to 2025-10-14.

Metric	Enhanced	Enhanced with DD Mitigation	Benchmark
Cumulative Return	13.3%	12.9%	11.2%
Annual Return	13.4%	12.9%	11.2%
Vol	9.5%	9.6%	8.5%
Sharpe Ratio	0.89	0.83	0.75
Max DD	11.8%	11.9%	7.9%
Return/maxDD	1.13	1.08	1.42

Monthly Correlations

Table 7: Monthly correlation matrix.

	Enhanced	Enhanced with DD Mitigation	Benchmark
Enhanced	100%	97%	92%
Enhanced with DD Mitigation	97%	100%	88%
Benchmark	92%	88%	100%

Commentary. All three series are highly aligned (every entry ≥ 0.88), which is what we want from a faithful replication. Two useful takeaways:

1. **97% between Enhanced and Enhanced with DD Mitigation:** the hedge doesn't reinvent the portfolio—it overlays convexity while keeping the same core exposures.

2. **88–92% versus the Benchmark:** the enhanced target introduces enough deviation to earn excess Sharpe and reduce drawdowns, yet stays close enough that risk parity's character is preserved.

Drawdown preservation. The CTA overlay is doing what it should: it reduces the *depth* of major equity crises sell-offs while keeping the core risk-parity footprint. MaxDD drops from 30.9% (No Hedge) to 23.5% (With DD Mitigation), a **24% relative reduction**. Versus the benchmark's 33.5% drawdown, the reduction is about **30%**.
Mechanism. The hedge engages mainly when trends extend through stress (rates rallies, equity sell-offs, commodity breaks). Its low/negative conditional beta to procyclical sleeves clips the left tail, so the portfolio spends less time and depth under water in crisis windows, while remaining 97% correlated with the No-Hedge variant.

7 Conclusion

The two-step enhanced decoding design improves efficiency while preserving the strategy's profile. **Enhanced** decodes to a variance-preserving, denoised target and raises risk-adjusted returns with volatility close to the benchmark. **Enhanced With DD Mitigation** retains this structure, adds a mild asymmetric mapping and a CTA risk-off overlay, and, on average, reduces drawdown depth and improves the return-to-drawdown ratio. High cross-correlations and modest tracking error indicate that gains stem from small, persistent tilts and crisis-time convexity rather than from a wholesale change in exposures.

8 Disclaimer

The analyses and results are based on historical data and simulations and do not constitute investment advice. Past performance is not indicative of future results. Model outcomes depend on assumptions, input data, and parameter choices; real-world results may differ due to market volatility, liquidity, and regime changes.

References

Ai For Alpha (2024). Turning trend following into an actual risk-off strategy: A novel approach to enhancing portfolio resilience. Technical report, Ai For Alpha. 7 pages. Available at `https://aiforalpha.com/dist/img/Spotlight_RiskOff.pdf`.

Ang, A. (2012). *Asset Management: A Systematic Approach to Factor Investing*. Oxford University Press, New York.

Ascheberg, M., Laeven, L., and Tüzemen, D. (2023). Macro networks and the propagation of shocks. *Journal of Financial Economics*, 149(2):387–411.

Asness, C. S., Frazzini, A., and Pedersen, L. H. (2012). Leverage aversion and risk parity. *Financial Analysts Journal*, 68(1):47–59.

Bartholomew, H. (2025). Novel risk-off cta strategy passes tariff test. Risk.net. 07 May 2025; accessed October 19, 2025.

Benhamou, E., Ohana, J.-J., and Guez, B. (2024). Generative ai: Crafting portfolios tailored to investor preferences. Technical Report SSRN Working Paper No. 4809740, Université Paris-Dauphine and AI For Alpha and EB AI Advisory. 9 pages. Posted: 15 Apr 2024.

Carvalho, C. M., Lopes, H. F., Polson, N. G., and Taddy, M. (2009). Dynamic modeling of latent factors: A bayesian approach. *The Annals of Applied Statistics*, 3(2):611–639.

Chaves, D. B., Hsu, J. C., Li, F., and Shakernia, O. (2011). Risk parity portfolio vs. other asset allocation heuristic portfolios. *The Journal of Investing*, 20(1):108–118.

Choueifaty, Y. and Coignard, Y. (2008). Toward maximum diversification. *The Journal of Portfolio Management*, 35(1):40–51.

Greyserman, A. and Kaminski, K. M. (2014). *Trend Following with Managed Futures: The Search for Crisis Alpha*. John Wiley & Sons, Hoboken, NJ.

Harvey, C. R., Liu, Y., and Zhu, H. (2016). Replicating anomalies. *The Review of Financial Studies*, 29(6):104–147. Originally circulated as NBER Working Paper No. 23377 (2017).

Hurst, B., Ooi, Y. H., and Pedersen, L. H. (2017). A century of evidence on trend-following investing. *The Journal of Portfolio Management*, 44(1):15–29.

Kim, C.-J. and Nelson, C. R. (1999). State-space models with regime switching: Classical and gibbs-sampling approaches with applications. *Journal of Econometrics*, 60(1–2):1–22.

Koijen, R. S., Moskowitz, T. J., Pedersen, L. H., and Vrugt, E. B. (2018). The cross-section of risk and return. *Journal of Financial Economics*, 127(2):358–379.

Koop, G. and Korobilis, D. (2013). *Bayesian Econometrics*. Palgrave Macmillan, London.

Ledoit, O. and Wolf, M. (2004). Honey, i shrunk the sample covariance matrix. *Journal of Portfolio Management*, 30(4):110–119.

Lee, W. Y. (2011). Risk-based asset allocation: A new answer to an old question? *The Journal of Portfolio Management*, 37(4):11–28.

Maillard, S., Roncalli, T., and Teïletche, J. (2010). On the properties of equally-weighted risk contributions portfolios. *The Journal of Portfolio Management*, 36(4):60–70.

Moskowitz, T. J., Ooi, Y. H., and Pedersen, L. H. (2012). Time series momentum. *Journal of Financial Economics*, 104(2):228–250.

Ohana, J.-J., Benhamou, E., Saltiel, D., and Guez, B. (2022). Deep decoding of strategies. Technical Report Université Paris-Dauphine Research Paper No. 4128693, Université Paris-Dauphine, PSL Research University. Last revised: May 5, 2024. Available at SSRN: `https://ssrn.com/abstract=4128693`.

Qian, E. (2005). Risk parity portfolios: Efficient portfolios through true diversification. *PanAgora Asset Management White Paper*. Available at SSRN: `https://ssrn.com/abstract=1895570`.

Roncalli, T. (2013). *Introduction to Risk Parity and Budgeting*. Chapman and Hall/CRC, Boca Raton, FL.

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