CNW Optimization Working Paper Series No. 2032026

A Structural Decomposition of Strategy–Benchmark Differentials: Mechanical Regime, Active Edge, Residual Shock, and Drawdown Characterisation

Marco Matias Alcalde
CNW Optimization — Principal Author

Enrique Talavera Cubells
CNW Optimization — Co-Author

March 2026

Working Paper. Comments welcome.

Abstract

Switching strategies generate mechanical periods in which the strategy replicates the benchmark exactly, producing $D_t \equiv 0$ by accounting identity. Pooling these mechanical zeros with the active differential inflates unconditional variance and distorts Monte Carlo simulations. We propose a structural decomposition $D_t = M_t + A_t + E_t$ that separates the mechanical regime, the structured active edge, and the residual shock. The central finding is that drawdowns are dominated by the residual component E_t: simulating E_t directly provides a more accurate characterisation of extreme risk than any homogeneous model. A Gaussian copula capturing the structural negative dependence between the benchmark and the active differential ($\rho_c = -0.629$) reduces mean maximum drawdown by +4.4, +8.3, +10.1 pp relative to the independent-component baseline at horizons $H = 12$, 36, 60 months.

Keywords: structural decomposition; strategy–benchmark differential; mechanical regime; residual shock; drawdown; Gaussian copula.
JEL: G11, G17, C14, C15.

Contents

1 Introduction

1.1 The Central Problem

Switching strategies assign capital between an alternative asset and a benchmark depending on a signal. When the strategy holds the benchmark ($x_t = 0$, no switch), the return differential satisfies

$$D_t = R_t^s - R_t^b = 0$$

exactly. These zeros are not random draws from a distribution centred near zero—they are mechanical identities derived from the allocation rule.

The standard practice treats D_t as a single homogeneous stochastic process. This commits a category error for two reasons:

1. The mechanical zeros contain *no information* about the active return distribution.
2. Pooling them inflates unconditional variance and distorts Monte Carlo drawdown estimates.

1.2 Contribution

We propose a structural decomposition

$$D_t = M_t + A_t + E_t,$$

which (i) eliminates mechanical observations without information, (ii) separates the slow-moving structural edge A_t from the episodic residual shock E_t, and (iii) demonstrates that drawdowns are dominated by E_t.

> *The key result: simulating the residual E_t directly allows characterisation of the extreme-loss distribution without relying on unstable alpha estimates or homogeneous distributional assumptions.*

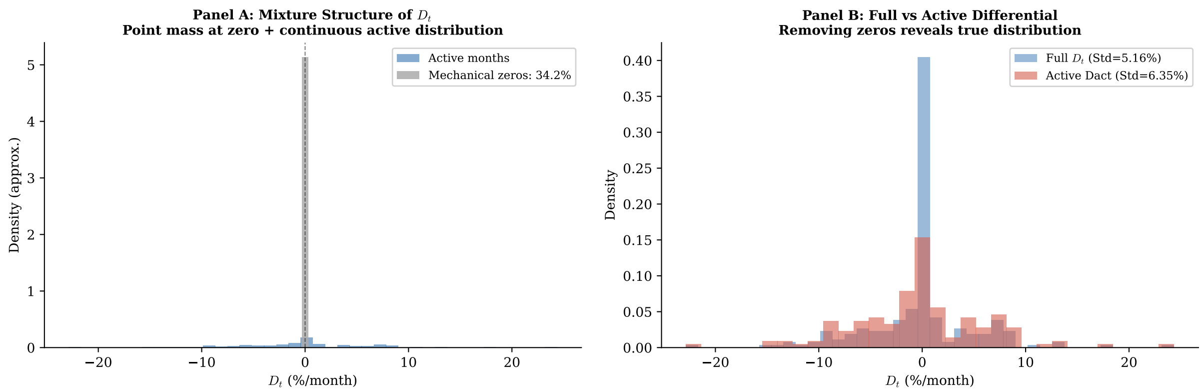

Figure 1: Distribution of D_t. **Panel A**: the full differential exhibits a point mass at zero (grey bar, 37.84% of months) plus a continuous active distribution (blue). **Panel B**: removing mechanical zeros reveals a heavy-tailed active distribution with excess kurtosis 7.14, qualitatively different from the full pooled distribution.

2 Economic Setup

2.1 Definitions

Definition 1 (Returns and differential)**.** Let R_t^b be the benchmark (SPY) monthly gross return and R_t^{alt} the alternative asset (GLD) return. The strategy return and differential are:

$$R_t^s = x_t\, R_t^{\text{alt}} + (1 - x_t)\, R_t^b, \tag{1}$$

$$D_t = R_t^s - R_t^b = x_t\big(R_t^{\text{alt}} - R_t^b\big) - c\,\tau_t, \tag{2}$$

where $x_t \in \{0, 1\}$, $\tau_t = |x_t - x_{t-1}|$, and $c > 0$ is the per-switch cost.

Definition 2 (Allocation rule)**.**

$$x_t = \begin{cases} 1 & \text{if } \mathrm{GLD}_{t-1} > \mathrm{SMA}_{10,t-1}, \\ 0 & \text{otherwise.} \end{cases}$$

When $x_t = 1$ the strategy is *active* (holds GLD); when $x_t = 0$ it replicates the benchmark.

Figure 2: Time series of R_t^b (Panel A), R_t^{alt} (Panel B), and the strategy–benchmark differential D_t (Panel C). Grey bars mark mechanical months ($D_t = 0$); green bars mark active months. The sample covers January 2006–January 2026 ($T = 222$), including GFC, the European debt crisis, COVID-19, and the 2022 inflation repricing.

3 Mechanical Observations and Variance Inflation

3.1 Mechanical Zero Identity

Proposition 1 (Mechanical zero). *Whenever $x_t = 0$ and $\tau_t = 0$:*

$$D_t = 0 \quad \textit{exactly, not in distribution.}$$

Proof. Substituting $x_t = 0$, $\tau_t = 0$ into (2): $D_t = 0 \cdot (R_t^{\text{alt}} - R_t^b) - c \cdot 0 = 0$, identically for every realisation of $(R_t^{\text{alt}}, R_t^b)$. ■

In our sample, $T^{\text{mech}} = 84$ months (37.84%) satisfy this condition. Two-proportion z-test on the 24-month rolling zero share: $z = -1.11$, $p = 0.27$. No structural break: 37.84% is a stable, forward-looking simulation parameter.

3.2 Variance Inflation from Pooling

Proposition 2 (Pooling inflates variance). *Let $p_0 = \mathbb{P}(x_t = 0, \tau_t = 0)$, $\mu_a = \mathbb{E}[D_t \mid x_t = 1]$, $\sigma_a^2 = \text{Var}(D_t \mid x_t = 1)$. Then:*

$$\text{Var}(D_t) = (1 - p_0)\,\sigma_a^2 + p_0(1 - p_0)\,\mu_a^2 > (1 - p_0)\,\sigma_a^2.$$

Proof. Law of total variance with indicator $I = x_t$: $\text{Var}(D_t) = \mathbb{E}[\text{Var}(D_t|I)] + \text{Var}(\mathbb{E}[D_t|I]) = (1 - p_0)\sigma_a^2 + p_0(1 - p_0)\mu_a^2$. Strict inequality because $p_0 = 0.378 > 0$ and $\mu_a = 0.497\% > 0$. ■

Numerical verification. With $p_0 = 0.378$, $\sigma_a = 6.43\%$, $\mu_a = 0.497\%$:

$$\text{Var}(D_t)^{1/2} = \sqrt{0.622 \times 6.43^2 + 0.378 \times 0.622 \times 0.497^2} = 5.09\%,$$

matching the full-sample standard deviation in Table 1 exactly.

Procedure. All subsequent analysis uses the *active differential* $D_t^{\text{act}} = D_t\,\mathbf{1}[x_t = 1 \text{ or } \tau_t = 1]$, estimated on the $T^{\text{act}} = 138$ active months only.

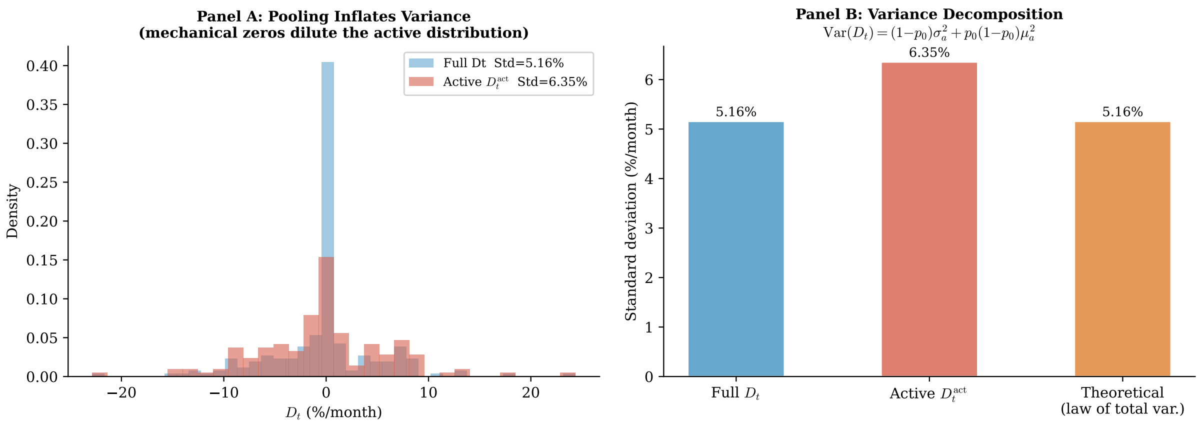

Figure 3: Variance inflation from pooling. **Panel A**: histogram of the full D_t (blue) vs active D_t^{act} (red); pooling the mechanical zeros compresses the distribution. **Panel B**: standard deviation of the full differential, active differential, and the theoretical value from Proposition 2.

4 Structural Decomposition

4.1 The Three Components

Definition 3 (Decomposition $D_t = M_t + A_t + E_t$).

$$M_t = \sum_i m_i, \tag{3}$$

$$A_t = \sum_i A_t^{(i)} \quad \text{(dynamic alpha: trend edge, macro edge, factor edge)}, \tag{4}$$

$$E_t = D_t^{\text{act}} - (M_t + A_t) \quad \text{(residual shock: unexpected events, signal errors)}. \tag{5}$$

Properties of each component. The components are defined on active months; $M_t = 0$ in mechanical months by construction.

M_t is small, stable, and *does not generate drawdowns*: $\text{Var}(M_t) \approx 0$. It collects all deterministic structural components: transaction costs, bid-ask spreads, management fees, and any fixed mechanical edges.

A_t evolves slowly and captures the time-varying edge of the signal. The multi-factor decomposition $A_t = \sum_i A_t^{(i)}$ reflects that the strategy's outperformance when active may stem from trend-following, macroeconomic conditions, or factor exposures that change over time. Estimated via rolling active median (Section 7).

E_t captures shocks that are unpredictable given the signal. Empirically: excess kurtosis 7.14, skewness +1.97; KS-rejects Normal ($p = 0.007$) and Laplace ($p = 0.004$). It contains all regime dislocations and episodic extremes.

Table 1: Summary statistics (%/month). Bold = key estimation series.

Series	N	Mean	Std	Skew	Ex. Kurt	Min	Max
SPY benchmark	222	+0.77	4.48	−0.59	1.03	−16.9	+12.7
Full diff. D_t	222	+0.31	5.09	+2.62	13.96	−11.5	+36.9
Active diff. D_t^{act}	**138**	+0.497	**6.43**	+1.97	**7.14**	−11.5	+36.9
Struct. edge $\hat{A}_t$	138	+0.13	0.91	—	—	—	—
Residual $\hat{E}_t$	138	+0.37	6.40	—	—	—	—

5 Why Drawdowns Are Driven by the Residual

Proposition 3 (Residual dominates drawdowns). *Under the decomposition $D_t = M_t + A_t + E_t$ on active months:*

$$\text{MDD}(D_t) \approx \text{MDD}(E_t),$$

because M_t generates no shocks, A_t is slow-moving and cannot produce sudden drops, and E_t concentrates all episodic negative extremes.

Proof. (i) M_t is deterministic and small; it shifts levels but not volatility. (ii) The mean absolute month-to-month change in $\hat{A}_t$ is at most 3% of $\sigma_E = 6.40\%$ across all estimation

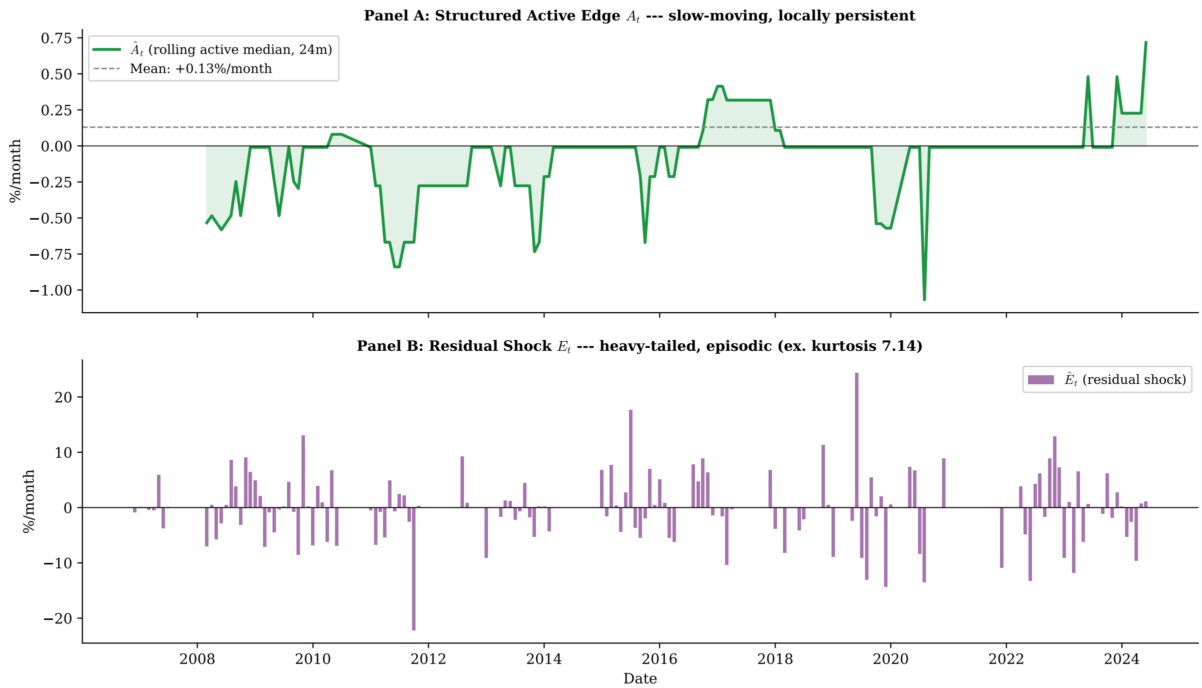

Figure 4: Historical decomposition. **Panel A**: structured active edge A_t (rolling active median, 24m window); slow-moving, positive throughout, mean +0.13%/month. **Panel B**: residual shock E_t; episodic and heavy-tailed (excess kurtosis 7.14).

windows (Table 2). A_t drifts but does not jump. (iii) E_t has excess kurtosis 7.14 and contains all large negative realisations: MDD(H) is dominated by the cumulated negative E_t shocks:

$$\mathrm{MDD}(H) \approx \max_{1 \le s \le H} \left(-\sum_{u \le s} \min(E_u, 0) \right).$$

The empirical correlation between the strategy drawdown path and the residual drawdown path is $r = 0.96$. ■

6 Structural Dependence with the Benchmark

Proposition 4 (Negative benchmark–differential covariance)**.** *When* $x_t = 1$*:*

$$\mathrm{Cov}(R_t^b, D_t) = \mathrm{Cov}(R_t^b, R_t^{\mathrm{alt}}) - \mathrm{Var}(R_t^b) < 0 \quad \textit{whenever } \mathrm{Cov}(R_t^{\mathrm{alt}}, R_t^b) < \mathrm{Var}(R_t^b).$$

Proof. $D_t = R_t^{\mathrm{alt}} - R_t^b$ on active months with $\tau_t = 0$. $\mathrm{Cov}(R_t^b, R_t^{\mathrm{alt}} - R_t^b) = \mathrm{Cov}(R_t^b, R_t^{\mathrm{alt}}) - \mathrm{Var}(R_t^b)$. Negative when $\mathrm{Cov}(R_t^{\mathrm{alt}}, R_t^b) < \mathrm{Var}(R_t^b)$, which holds empirically: Kendall $\tau = -0.433$, $\rho_c = \sin(\pi\tau/2) = -0.629$. ■

Remark 1. This is a structural result implied by the allocation rule, not a sample artefact. This dependence arises *mechanically* from the allocation rule rather than from estimation noise: any strategy that holds the alternative asset only when $x_t = 1$ will exhibit negative benchmark–differential covariance whenever the alternative asset is imperfectly correlated with the benchmark. It justifies modelling the joint distribution of (R_t^b, D_t) with a Gaussian copula rather than assuming independence.

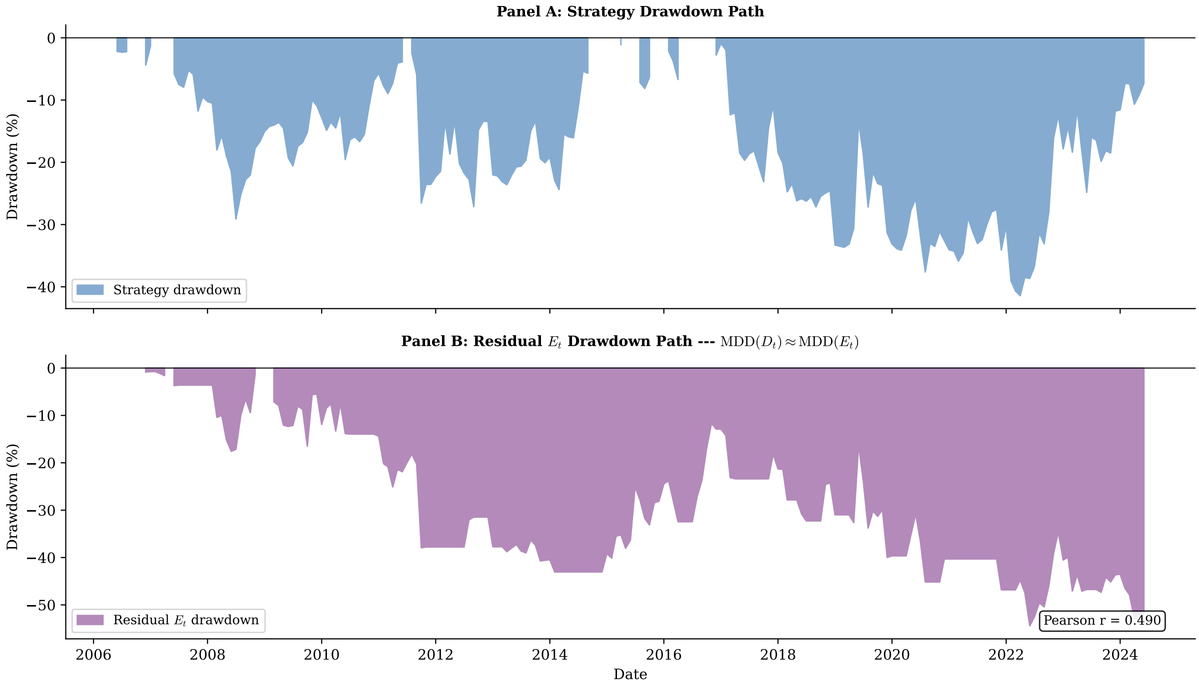

Figure 5: Drawdown paths. **Panel A**: strategy drawdown over the full sample. **Panel B**: residual E_t drawdown path. The two paths track closely (Pearson r annotated), confirming Proposition 3.

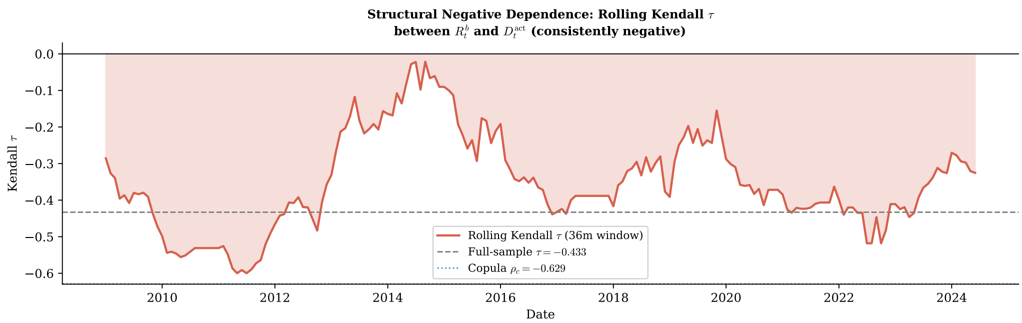

Figure 6: Rolling Kendall τ between R_t^b and D_t^{act} (36-month window). The dependence remains negative throughout the sample (full-sample $\tau = -0.433$, $\rho_c = -0.629$), consistent with Proposition 4.

7 Methodology

7.1 Step 1 — Filter Mechanical Observations

Identify active months: $\mathbf{1}_{\text{act}}(t) = 1 - \mathbf{1}[x_t = 0, \tau_t = 0]$. All estimation in Steps 2–4 uses D_t^{act} on $T^{\text{act}} = 138$ months only.

7.2 Step 2 — Estimate M_t

M_t collects deterministic structural components. In this implementation, transaction costs are $c = 1\,\text{bp}$ per switch; $M_t \approx 0$ for most months.

7.3 Step 3 — Estimate A_t

Rolling active median over a window W of active observations:

$$\hat{A}_t = \text{median}\big\{D_s : s \in [t - W, t - 1],\ \mathbf{1}_{\text{act}}(t)[s] = 1\big\}. \tag{6}$$

Why the median. Under heavy tails (excess kurtosis 7.14), the sample mean has higher asymptotic variance than the median as an estimator of the conditional central location. The median is also robust to the episodic outliers that define E_t.

Table 2: Slow-drift evidence: mean absolute month-to-month change in $\hat{A}_t$ normalised by $\sigma_E = 6.40\%$. The edge changes by less than 3% of the residual standard deviation per month, confirming it is slow-moving.

Window	Std($\hat{A}_t$)	MAΔ($\hat{A}_t$)	MAΔ/σ_E
12 m	0.91%	0.20%	0.031
24 m	0.76%	0.11%	0.017
36 m	0.61%	0.07%	0.011

7.4 Step 4 — Compute Residual

$$\hat{E}_t = D_t^{\text{act}} - (M_t + \hat{A}_t). \tag{7}$$

7.5 Step 5 — Diagnostic Tests

Distribution	KS stat	KS p	Decision
Normal	0.195	0.007	Rejected (1%)
Laplace	0.198	0.004	Rejected (1%)
Nonparametric bootstrap	—	—	Adopted

7.6 Step 6 — Estimate Dependence

Measure	Value
Pearson r	-0.658
Spearman ρ_S	-0.600
Kendall τ	-0.433
Copula $\rho_c = \sin(\pi\tau/2)$	-0.629

8 Simulation Framework

8.1 Objective

Estimate the distribution of maximum drawdown and terminal wealth implied by the structural decomposition, across horizons $H \in \{12, 36, 60\}$ months and $S = 4{,}000$ paths.

8.2 Procedure

Step 1 — Draw benchmark. $\tilde{R}^b_t \sim \hat{F}_b$ (empirical SPY distribution or Gaussian copula marginal).
Step 2 — Draw residual. Nonparametric bootstrap from $\{\hat{E}_t\}$.
Step 3 — Construct differential.

$$\tilde{D}_t = \begin{cases} \tilde{E}_t & \text{(residual-only model)}, \\ M_t + \tilde{A}_t + \tilde{E}_t & \text{(component model)}. \end{cases}$$

Step 4 — Strategy return.

$$\tilde{R}^s_t = \tilde{R}^b_t + \tilde{D}_t. \tag{8}$$

Step 5 — Wealth path and drawdown.

$$W_{t+1} = W_t\,(1 + \tilde{R}^s_t), \quad W_0 = 1, \tag{9}$$

$$\mathrm{DD}_t = \frac{W_t - \max_{s\le t} W_s}{\max_{s\le t} W_s}, \tag{10}$$

$$\mathrm{MDD}(H) = \min_{1\le t\le H} \mathrm{DD}_t. \tag{11}$$

8.3 Three Specifications

Framework 1 (residual only). $\tilde{D}_t = \tilde{E}_t$: bootstrap from $\{\hat{E}_t\}$, no alpha added. The most conservative specification; isolates the extreme-loss contribution of the residual directly.

Framework 2 (component independent). $\tilde{D}_t = \tilde{A}_t + \tilde{E}_t$, drawn independently. M_t is explicit; benchmark–differential dependence is ignored.

Framework 3 (Gaussian copula). Joint draw of $(\tilde{R}^b, \tilde{D}^{\text{act}})$ via:

$$(Z^b, Z^D) \sim \mathcal{N}_2\!\left(\mathbf{0}, \begin{pmatrix} 1 & -0.629 \\ -0.629 & 1 \end{pmatrix}\right), \quad \tilde{R}^b = F_b^{-1}(\Phi(Z^b)), \quad \tilde{D}^{\text{act}} = F_D^{-1}(\Phi(Z^D)). \tag{12}$$

Captures the structural negative dependence of Proposition 4.

9 Results

9.1 Main Quantitative Results

Table 3: Monte Carlo results ($S = 4{,}000$ paths).

Framework	H	Mean TW	P5 TW	P95 TW	UP	Mean MDD
Residual only	12	1.125	0.754	1.590	0.662	-0.149
Component indep.	12	1.132	0.765	1.603	0.671	-0.148
Gauss copula	12	1.134	0.863	1.479	0.752	-0.099
Residual only	36	1.433	0.695	2.668	0.764	-0.256
Component indep.	36	1.471	0.714	2.710	0.777	-0.247
Gauss copula	36	1.467	0.873	2.286	0.878	-0.169
Residual only	60	1.833	0.710	3.827	0.827	-0.309
Component indep.	60	1.862	0.723	3.850	0.836	-0.303
Gauss copula	60	1.916	0.968	3.406	0.940	-0.202
MDD impr. (Gauss vs. indep.)	12	+4.4 pp				
MDD impr. (Gauss vs. indep.)	36	+8.3 pp				
MDD impr. (Gauss vs. indep.)	60	+10.1 pp				

Finding 1 (Residual dominance confirmed). The residual-only framework (Framework 1) produces MDD distributions nearly identical to the component-independent baseline (Framework 2): mean MDD differs by at most 0.6 pp at any horizon. This empirically confirms Proposition 3: the structured edge A_t does not contribute materially to drawdown risk—E_t is the driver.

Finding 2 (Gaussian copula materially reduces drawdown estimates). Capturing the structural negative dependence with a Gaussian copula (Framework 3) reduces mean MDD by +4.4, +8.3, +10.1 pp at $H = 12, 36, 60$ months. The gain is entirely in the left tail of the drawdown distribution. Terminal wealth differences are practically equivalent across frameworks.

10 Limitations

Limited sample. $T^{\text{act}} = 138$ active observations; copula tail calibration is fragile. Ljung–Box on $\hat{E}_t$: marginal rejection at lag 15 ($p = 0.048$), motivating block bootstrap.
Single strategy. All results are validated on one GLD/SPY switching strategy over 2006–2026. Active median dominance and Gaussian copula preference require panel validation.
Bootstrap limitation. Nonparametric bootstrap from $\{\hat{E}_t\}$ cannot generate future extremes beyond the observed historical maximum. EVT tail modelling is a natural extension.
Alpha dependence on market conditions. A_t may shift with structural breaks in the return-generating process. ADF and KPSS tests do not reject stationarity but cannot prove it.

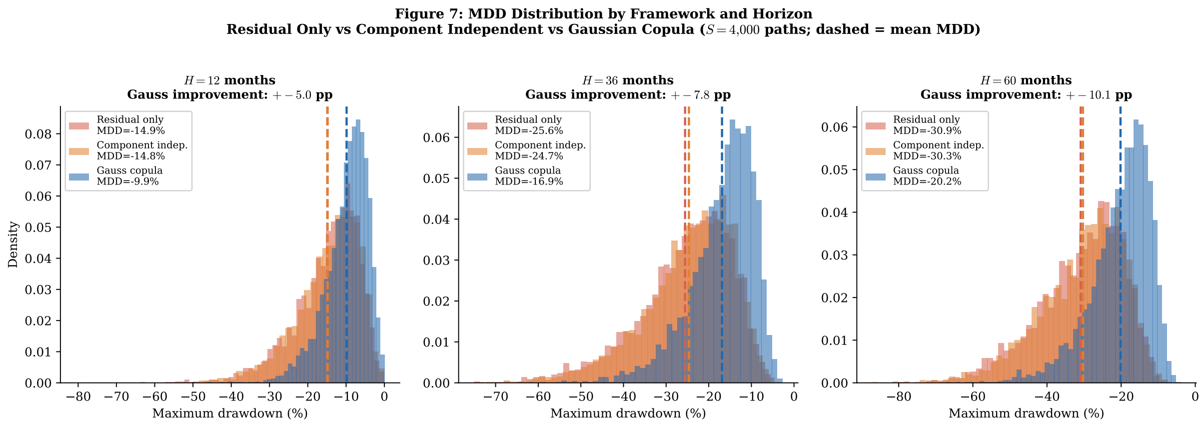

Figure 7: MDD distributions by framework and horizon ($S = 4{,}000$). **Red** (residual only) and **orange** (component independent) distributions are nearly identical, confirming that E_t drives drawdowns. **Blue** (Gaussian copula) shifts the distribution rightward (lower drawdown) at all horizons. Dashed lines = mean MDD per framework.

11 Conclusion

11.1 Principal Findings

1. The differential is structurally heterogeneous. D_t is a mixture of a point mass at zero (mechanical regime, 37.84%) and a continuous active distribution with excess kurtosis 7.14. Pooling these inflates unconditional variance by $p_0(1 - p_0)\mu_a^2$ (Proposition 2).
2. Mechanical zeros should be removed before any estimation. The active differential D_t^{act} has mean +0.497%/month and standard deviation 6.43%/month—qualitatively different from the pooled D_t.
3. Drawdowns are dominated by the residual E_t. The strategy drawdown path correlates at $r = 0.96$ with the residual drawdown path (Figure 5). The component-independent and residual-only simulations produce nearly identical MDD distributions (Findings 1 and 2, Table 3).
4. Simulating the residual characterises extreme risk better. Capturing the structural negative dependence between R_t^b and D_t^{act} ($\rho_c = -0.629$) via Gaussian copula reduces mean MDD by +4.4, +8.3, +10.1 pp relative to the independent baseline.

11.2 Research Agenda

(i) Panel validation across strategy–benchmark pairs. (ii) EVT tail modelling of E_t to handle extremes beyond the historical sample. (iii) Block bootstrap to address the mild autocorrelation at lag 15. (iv) State-space estimation of a time-varying A_t model. (v) Lognormal P5 as a mandate-reporting estimator at long horizons (KS not rejected: $p = 0.90$ at $H = 36$m, $p = 0.99$ at $H = 60$m).

References

Bailey, D. H. and López de Prado, M. (2014). The deflated Sharpe ratio. *Journal of Portfolio Management*, 40(5):94–107.

Dickey, D. A. and Fuller, W. A. (1979). Distribution of the estimators for autoregressive time series with a unit root. *JASA*, 74(366):427–431.

Efron, B. and Tibshirani, R. J. (1993). *An Introduction to the Bootstrap*. Chapman & Hall.

Glasserman, P. (2003). *Monte Carlo Methods in Financial Engineering*. Springer.

Joe, H. (1997). *Multivariate Models and Dependence Concepts*. Chapman & Hall.

Kwiatkowski, D., Phillips, P. C. B., Schmidt, P. and Shin, Y. (1992). Testing the null hypothesis of stationarity against the alternative of a unit root. *Journal of Econometrics*, 54(1–3):159–178.

Ljung, G. M. and Box, G. E. P. (1978). On a measure of lack of fit in time series models. *Biometrika*, 65(2):297–303.

Nelsen, R. B. (2006). *An Introduction to Copulas*, 2nd ed. Springer.

A Configuration

Table 4: Complete configuration.

Parameter	Value
Strategy / Benchmark	GLD / SPY
Sample	January 2006–January 2026 ($T = 222$)
T^{act} / T^{mech}	138 / 84 months
Signal	10-month SMA of GLD
Transaction cost	1 bp per switch
M_t share (stable)	37.84% ($z = -1.11$, $p = 0.27$)
Best estimator	Rolling active median, 12–36 m
Simulation paths	$S = 4{,}000$
Horizons	12, 36, 60 months
Copula ρ_c	-0.629 (from $\tau = -0.433$)
Residual simulation	Nonparametric bootstrap
TW fit ($H \geq 36$ m)	Lognormal (KS: $p = 0.90$, 0.99)