Factor Inertia:
Dynamic Multi-Factor Investing in the S&P 500

A dynamic factor-allocation rule that overweights factors
in an uptrend, and its economic value over 2000–2026

Javier Mestre Molins

Factor Inertia: Dynamic Multi-Factor Investing in the S&P 500

Abstract. Factor premia are not constant: their magnitude and relative structure vary with the market regime, which questions the optimality of static allocations. This report develops and evaluates a dynamic multi-factor equity strategy on the S&P 500 over 2000–2026, integrating five style factors (Value, Quality, Momentum, Growth and Low Volatility) at the individual-stock level to avoid the exposure dilution of combining single-factor indices. The contribution is *factor inertia*: an allocation rule that tilts the portfolio towards the factors with the strongest recent trend, inferring the market regime from the factors' own performance rather than from any macroeconomic forecast. The signal blends three multi-horizon exponential moving averages, converts them into weights through a *softmax* transformation under exposure caps, and rebalances semi-annually. Its parameters are fixed *a priori*, with no optimisation on the backtest. The inertia strategy beats the static equal-weight baseline on every risk-adjusted metric (CAGR 9.1% → 11.3%, Sharpe 0.49 → 0.58), with a CAPM alpha that survives the Fama–French three-factor model. The Sharpe difference is not statistically significant, yet the economic gap compounds into 69% more terminal wealth over 26 years, at roughly 20 bps of annual cost. A complementary sector analysis shows that factor predominance is industry-specific. The core contribution, however, is a reproducible, bias-free system in which the factors' own trends drive the allocation, a practical and economically meaningful route to dynamic factor investing.

Keywords: factor investing, factor inertia, dynamic factor allocation, factor timing, smart beta, S&P 500. **JEL:** G11, G12, G14.

1 The problem: factors work, but not all the time

Systematic factor investing is one of the best-documented ideas in asset management: a small set of stock characteristics (valuation, quality, price momentum, growth, volatility) explains persistent differences in returns beyond market exposure, and the industry has packaged them into low-cost *smart-beta* products. Yet putting several factors to work in a single portfolio raises two practical problems that matter to any manager.

The first is integration. Combining single-factor indices (*Top-Down*) dilutes exposures: a stock picked for cheapness may score poorly on momentum, and once indices are merged these cross-exposures cancel out (Innes, 2017). Scoring each stock on all factors at once (*Bottom-Up*) and selecting the "complete athletes" that rank high across several factors avoids this dilution.

The second, and more decisive, is cyclicality. Factor premia are not constant; their size and ranking shift with the macroeconomic regime (Hao & Watts, 2024; Kuo & Huang, 2022). The factor that leads one regime tends to lag the next, so a fixed allocation is a permanent bet on what worked in the past. This report asks a single question: *can we combine multiple factors efficiently and, on top of that, adapt their weights dynamically to the market environment?*

Our answer builds on *factor momentum*: Ehsani and Linnainmaa (2022) show that recent factor performance persists, so factors on a good run tend to keep outperforming in the short term; indeed, they argue that much of the stock-level momentum of Jegadeesh and Titman (1993) is itself a by-product of this factor-level inertia. What is new here is the practical synthesis: a transparent, parameter-light allocation rule (*factor inertia*), integrated Bottom-Up, built free of the usual backtest biases, and delivered through a reproducible platform on which every result below can be re-run.

2 Data and factor construction

The investment universe is the historical, point-in-time composition of the S&P 500 at each rebalancing date, not today's membership. This removes survivorship bias: companies later dropped for bankruptcy, merger or shrinking size are included exactly while they belonged to the index. Prices are daily and adjusted for dividends and splits; fundamentals are quarterly and enter the system only from their official release date, which removes look-ahead bias. The simulation spans 1 January 2000 to 29 April 2026 (6,626 trading days).

Following S&P Dow Jones index methodology (S&P Dow Jones Indices, 2023a, 2023b, 2023c, 2023d), each factor is a composite z-score. For a given date, every descriptor is cross-sectionally ranked and standardised ($Z_{i,d} = (R_{i,d} - \mu)/\sigma$, capped at $\pm 3\sigma$), and the factor score is the average of its descriptors' z-scores. Averaging several descriptors per factor, rather than relying on a single ratio, guards against noise and data-snooping. Standardisation is done within each GICS sector where structural sector differences would otherwise dominate (e.g. valuation ratios), and globally for price-based signals. Table 1 lists the five factors.

Table 1: The five style factors and their descriptors.

Factor	Descriptors	Rationale
Value	Book/Price, Earnings/Price, Sales/Price	Cheap stocks earn more
Quality	ROE, Accruals ratio, Financial leverage	Sound fundamentals
Momentum	12-1 month return / volatility	Winners keep winning
Growth	3Y sales growth, ΔEPS/Price, mom. signal	Fundamental dynamism
Low Volatility	Inverse realised volatility	Low-risk stocks over-earn

Portfolios are equal-weighted and rebalanced semi-annually, which keeps turnover and costs low while still tracking the factor signals. Each single-factor portfolio holds the top 100 stocks; the combined multi-factor portfolio holds the top 25. Realistic transaction costs are charged on traded volume only (6 bps proportional plus a $1 fixed fee), deducted from capital before rebalancing.

3 The factor inertia rule

Each stock i receives a combined multi-factor score $S_i = \sum_{f=1}^{5} w_f \, Z_{i,f}$, and the 25 highest-scoring stocks form the portfolio. The only difference between the static baseline and the dynamic strategy is the weight vector w_f:

- **Equal-weight (baseline, H_0):** $w_f = 1/5$ for all factors, fixed forever.
- **Factor inertia:** w_f tilts each rebalance towards the factors in a sustained uptrend.

The rule's input is each factor's own track record: the equity curve EC_f of its single-factor portfolio. It reads the *trend* of that curve and turns it into weights in three transparent steps.

1. Reading the trend (moving averages). For each factor we measure how far its own performance curve EC_f sits above or below its *exponential moving average* (EMA) at three horizons, roughly one month, one quarter and one year:

$$I_f^{(h)} = \frac{EC_f - \mathrm{EMA}_f^{(h)}}{\mathrm{EMA}_f^{(h)}}, \qquad h \in \{21\text{d},\, 63\text{d},\, 252\text{d}\}. \tag{1}$$

The short horizon catches fast regime shifts, the long one confirms persistent trends, and the medium one is the central signal. The three are blended into a single *inertia* reading per factor, with the medium horizon weighted double:

$$\bar{I}_f = 0.25\, I_f^{21\text{d}} + 0.50\, I_f^{63\text{d}} + 0.25\, I_f^{252\text{d}}. \tag{2}$$

A positive $\bar{I}_f$ means the factor trades above its own trend (bullish), a negative value below it (bearish). An EMA is used rather than a simple average because it reacts to turns with less lag.

2. Turning trend into weights (softmax). The five inertia readings, once standardised into z-scores across the five factors (z_f), so their scale is comparable, are converted into portfolio weights with a *softmax* transformation,

$$\tilde{w}_f(t) = \frac{e^{z_f(t)/T}}{\sum_{j=1}^{5} e^{z_j(t)/T}}, \tag{3}$$

so that factors with stronger inertia receive exponentially more weight. A single temperature parameter T governs how aggressive that tilt is: as $T \to \infty$ the weights collapse back to the static $1/N$, while lower values sharpen the response. We set $T = 0.8$, a moderately assertive tilt.

3. Diversification caps. Finally, the weights are constrained to the range 0–50%. The floor lets the rule switch off a persistently weak factor; the ceiling prevents any single factor from dominating the book. The full formulation of all three steps, with the parameter values, is given in Appendix A.

Crucially, all parameters (horizons, blend weights, temperature, caps) are fixed *a priori* from standard practice and the momentum literature, *not* optimised on the backtest. This removes the risk of data snooping and is what lets us read the economic improvement as a genuine signal rather than an artefact of curve-fitting.

4 Results

4.1 The factors capture risk premia, not anomalies

Each single-factor portfolio earns a CAPM alpha that is statistically significant at the 1% level (between 4.2% and 6.0% per year; joint GRS test $F = 3.02$, $p = 0.010$ (Gibbons et al., 1989); Newey–West standard errors). However, once we control for the Fama–French five-factor model (Fama & French, 2015), every alpha disappears. This is what we would expect: the portfolios load on the documented size, value, profitability and investment premia, so their excess return reflects known premia rather than an unexploited anomaly. This validates the engine and sets an honest baseline for what follows. Figure 1 shows the five single-factor portfolios compounding steadily over the period, with a clear risk–return trade-off: Value leads on return but with the deepest drawdowns, while Low Volatility rides visibly smoother. Full single-factor metrics and regression tables are reported in Appendix B.

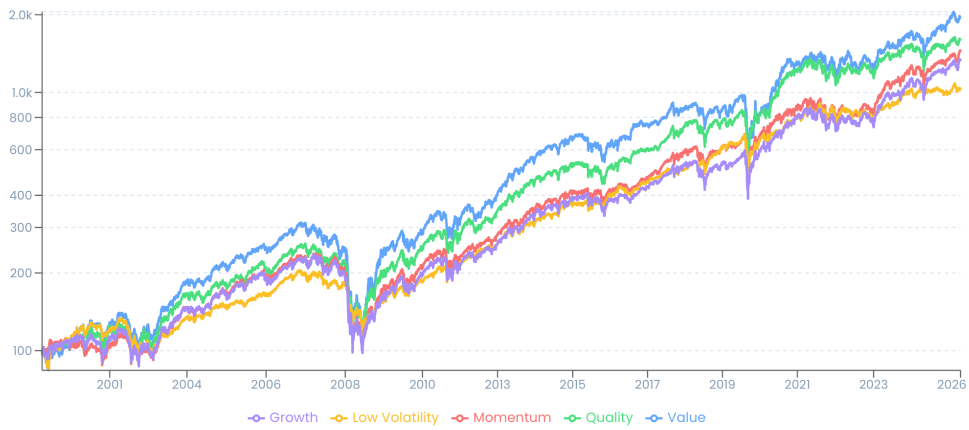

Figure 1: Cumulative return of the five single-factor portfolios (base 100, log scale), 2000–2026.

4.2 Dynamic allocation adds economic value

Table 2 compares the two combined strategies and the index. The inertia rule dominates the static baseline on every risk-adjusted metric, and both beat the S&P 500 comfortably, confirming that factor selection adds value and that inertia amplifies it.

Table 2: Performance of the combined strategies vs. the S&P 500 (total return), 2000–2026. Best value of each metric in bold.

Metric	S&P 500 TR	Equal Weight	Factor Inertia
CAGR	8.12%	9.07%	**11.28%**
Volatility	19.33%	**16.48%**	18.12%
Sharpe ratio	0.40	0.49	**0.58**
Sortino ratio	0.57	0.63	**0.76**
Calmar ratio	0.15	0.18	**0.23**
Max. drawdown	−55.8%	−51.4%	**−50.1%**
CAPM alpha (ann.)	—	3.88%*	**6.18%****

* significant at 5%, ** at 1%.

Figure 2 shows the capital curves: the inertia line separates progressively from the equal-weight line, especially after 2013. That gap is the compounding of +2.2 percentage points of annual CAGR over 26 years.

Figure 2: Cumulative value of the combined strategies vs. the S&P 500 total return (base 100, log scale), 2000–2026.

Where does the edge come from? We regress each strategy on progressively richer factor models (Table 3). The inertia alpha survives the three-factor model (Fama & French, 1993) (4.1%, still significant) and only fades under the five-factor model, where the profitability (RMW) and investment (CMA) premia absorb it ($\beta_{\text{CMA}} = 0.43$, $p < 0.001$). The equal-weight alpha, by contrast, already loses significance at the three-factor stage. The inertia edge is genuine alpha: it comes from a dynamic tilt towards high-quality, disciplined-capital companies, not from a passive bet on small or cheap stocks.

Table 3: Annual alpha of the combined strategies across nested factor models (p-value in parentheses).

Model	Equal Weight	Factor Inertia
CAPM (market only)	3.88% (0.020)*	6.18% (0.002)**
FF3 (+ size, value)	1.96% (0.213)	4.08% (0.034)*
FF5 (+ profitab., invest.)	−0.41% (0.785)	2.02% (0.269)

The dynamic weights are the mechanism, and Figure 3 makes the contrast plain: the static baseline (panel b) is five flat 20% bands, indifferent to the environment, while the inertia allocation (panel a) "breathes" with the cycle. Value spikes in the dot-com bust and the 2022 inflation shock, Growth and Momentum dominate the low-rate 2010s, and Low Volatility takes over in stress episodes. The rule uses no macro variable; it reads these regimes purely from each factor's own inertia. This produces a markedly convex capture profile: the strategy captures 94.9% of up-market moves but only 64.4% of down-moves, against 81.6% and 60.0% for the static baseline.

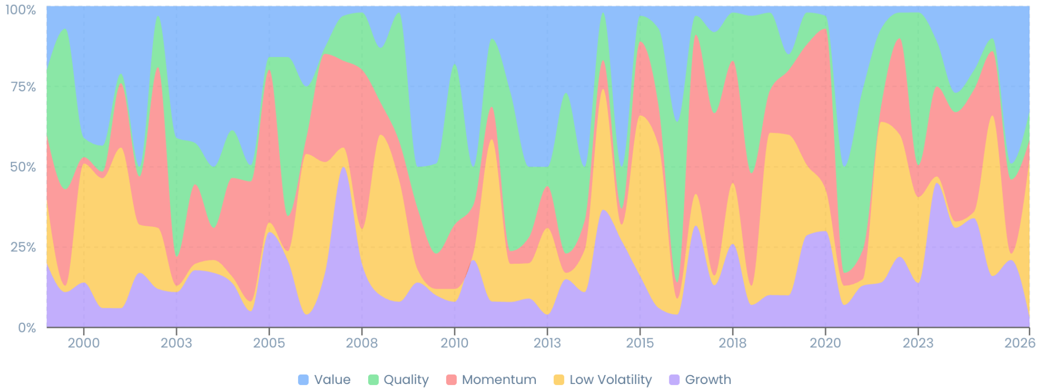

(a) Factor inertia: dynamic weights that rotate with the regime.

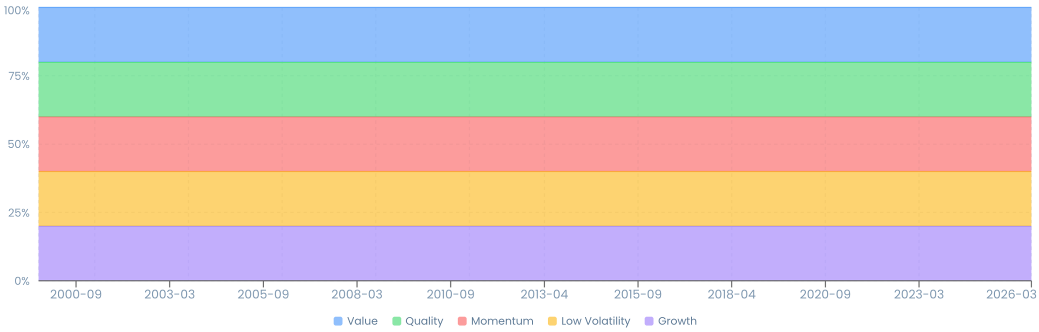

(b) Equal weight: static $1/N$ allocation, constant by construction.

Figure 3: Factor weights over time, 2000–2026. The inertia rule (a) shifts exposure towards the trending factors, inferred only from the factors' own performance; the static baseline (b) holds five constant 20% bands. This contrast is the mechanism behind the performance gap.

An honest statistical verdict. Is the Sharpe improvement statistically conclusive? No. The Jobson–Korkie test (Jobson & Korkie, 1981; Memmel, 2003) gives $p = 0.42$ and a block bootstrap $p = 0.23$; we cannot formally reject H_0. This is the expected outcome, not a failure: the two strategies are 90% correlated (same universe, same method) and the system makes only 52 independent allocation decisions in 26 years, so statistical power to detect a small Sharpe gap is intrinsically low. What is not ambiguous is the economic magnitude: compounding +2.2 pp of CAGR, the inertia strategy multiplies initial capital by 16.6 against 9.8 for the baseline, i.e. 69% more terminal wealth, achieved at ∼20 bps of annual cost. The result is economically substantial but statistically inconclusive, a pattern that is common in the factor-timing literature.

4.3 A complementary result: factor predominance by sector

Running the same engine independently inside each of the eleven GICS sectors adds a secondary but interpretable finding: no factor wins everywhere. Value leads in cyclical, balance-sheet-driven sectors (Financials, Energy, Industrials) and in Health Care; Quality where cash flows are stable and multiples high (Technology, Consumer Staples); Momentum and Growth in price-trend-driven sectors; and Low Volatility where cutting risk is the main lever. The pattern is robust, agreeing under a pure-return ranking in 9 of the 11 sectors, and it offers a practical side-benefit: a map of which factor to prioritise in each industry, useful even to an investor who never adopts the dynamic rule. Figure 4 maps every factor–sector combination.

Sector	Value	Quality	Momentum	Low Volatility	Growth	Equal Weighted	Inertia Weighted
Communication Services	0.26	0.24	0.27	0.29	0.27	0.26	0.29
Consumer Discretionary	0.41	0.52	0.42	0.57	0.41	0.51	0.50
Consumer Staples	0.50	0.54	0.46	0.50	0.42	0.52	0.55
Energy	0.48	0.42	0.40	0.38	0.41	0.45	0.42
Financials	0.48	0.41	0.41	0.41	0.30	0.44	0.42
Health Care	0.63	0.53	0.51	0.58	0.43	0.63	0.49
Industrials	0.61	0.50	0.57	0.52	0.50	0.58	0.56
Information Technology	0.38	0.42	0.39	0.34	0.40	0.37	0.40
Materials	0.38	0.41	0.48	0.35	0.47	0.40	0.42
Real Estate	0.36	0.35	0.37	0.33	0.37	0.35	0.36
Utilities	0.50	0.53	0.57	0.51	0.50	0.55	0.57

Figure 4: Sharpe ratio by factor and GICS sector, 2000–2026. Green: best risk-adjusted performance in the row; red: worst. The last two columns are the combined equal-weight and inertia portfolios within each sector.

5 Why it matters in practice

For a portfolio manager, the study offers a few concrete, usable takeaways.

How to build the book. Integrating Bottom-Up, selecting the "complete athletes" that rank high across several factors at once, produces cleaner and more intentional exposures than stacking single-factor sleeves, whose cross-exposures cancel out. This is a construction choice a manager can adopt directly, independently of any timing view.

How to allocate across factors. The inertia rule offers a way to move factor weights through the cycle *without* a forecasting apparatus: it uses no macroeconomic inputs and no regime model, only each factor's own recent trend, refreshed on a simple semi-annual schedule with a handful of transparent parameters. Its most valuable practical property is the *shape* of its risk: the strategy captures 94.9% of up-market moves but only 64.4% of the downside (against 81.6% and 60.0% for the static book). That convex profile, more participation in rallies and more protection in drawdowns, is what a risk budget is meant to buy.

The approach is, moreover, implementable and governable. Rebalancing 25 positions twice a year gives annual turnover of 1.56 and total costs of roughly 20 bps, which barely dent the edge and sit comfortably within an institutional book. Because every parameter is fixed *a priori* and every result is fully reproducible, the system is auditable rather than a black box tuned to the sample, a property that matters to allocators and risk committees as much as raw performance. And while the Sharpe gain is not statistically decisive, its economic size, 69% more terminal wealth over the horizon, is the kind of compounding difference that dominates a long-horizon mandate.

6 Limitations and conclusion

Three limitations frame the results honestly. First, there is no formal out-of-sample split; this is mitigated because the rule decides in real time, using only the information available up to each date, and because its parameters are set *a priori*, but confirming the edge on other markets and windows remains the key pending test. Second, the universe is confined to US large caps. Third, statistical power is low, for the structural reasons already noted, so the Sharpe improvement, though consistent in sign, does not reach significance.

The conclusion has two sides, and both should be stated plainly. Factor inertia is an allocation signal that is economically valuable and operationally viable: it beats the static baseline on every risk-adjusted metric, its edge is genuine alpha that survives the three-factor model, and it compounds into 69% more terminal wealth, all at negligible cost and with no macro inputs. At the same time, its superiority is not statistically confirmed against the static alternative. Alongside a robust, industry-specific map of factor predominance, the report delivers a reproducible, bias-free multi-factor system and shows that inferring the regime from the factors themselves is a promising, practical path to dynamic factor allocation. The priority ahead is to test whether that edge generalises beyond the S&P 500 of 2000–2026, the true out-of-sample proof.

References

Ehsani, S., & Linnainmaa, J. T. (2022). Factor momentum and the momentum factor. *The Journal of Finance*, *77*(3), 1877–1919.

Fama, E. F., & French, K. R. (1993). Common risk factors in the returns on stocks and bonds. *Journal of Financial Economics*, *33*(1), 3–56.

Fama, E. F., & French, K. R. (2015). A five-factor asset pricing model. *Journal of Financial Economics*, *116*(1), 1–22.

Gibbons, M. R., Ross, S. A., & Shanken, J. (1989). A test of the efficiency of a given portfolio. *Econometrica*, *57*(5), 1121–1152.

Hao, W. B., & Watts, R. (2024). *A historical perspective on factor index performance across macroeconomic cycles* (Research Paper). S&P Dow Jones Indices.

Innes, A. (2017). *The merits and methods of multi-factor investing* (Research Paper). S&P Dow Jones Indices.

Jegadeesh, N., & Titman, S. (1993). Returns to buying winners and selling losers: Implications for stock market efficiency. *The Journal of Finance*, *48*(1), 65–91.

Jobson, J. D., & Korkie, B. M. (1981). Performance hypothesis testing with the Sharpe and Treynor measures. *The Journal of Finance*, *36*(4), 889–908.

Kuo, Y.-S., & Huang, J.-T. (2022). Factor-based investing in market cycles: Fama–French five-factor model of market interest rate and market sentiment. *Journal of Risk and Financial Management*, *15*(10), 460.

Memmel, C. (2003). Performance hypothesis testing with the Sharpe ratio. *Finance Letters*, *1*(1), 21–23.

S&P Dow Jones Indices. (2023a). *S&P Enhanced Value Indices Methodology* (Index Methodology). S&P Global.

S&P Dow Jones Indices. (2023b). *S&P Growth Indices Methodology* (Index Methodology). S&P Global.

S&P Dow Jones Indices. (2023c). *S&P Momentum Indices Methodology* (Index Methodology). S&P Global.

S&P Dow Jones Indices. (2023d). *S&P Quality Indices Methodology* (Index Methodology). S&P Global.

Appendices

A The factor inertia algorithm in full

Section 3 defines the per-horizon inertia $I_f^{(h)}$, its blend $\bar{I}_f$ and the softmax weighting. This appendix completes the formulation with the EMA recursion and the exposure-clamping step. Each EMA is computed recursively, weighting recent observations more heavily:

$$\mathrm{EMA}_f^{(h)}(t) = \alpha_h \, EC_f(t) + (1-\alpha_h)\, \mathrm{EMA}_f^{(h)}(t-1), \qquad \alpha_h = \frac{2}{h+1}, \quad h \in \{21, 63, 252\}. \tag{4}$$

The blended signals $\bar{I}_f$ are standardised across the five factors (z_f) and mapped to target weights with a softmax at temperature T, then clamped to $[w_{\min}, w_{\max}]$ and re-normalised iteratively so that $\sum_f w_f = 1$:

$$\tilde{w}_f(t) = \frac{e^{z_f(t)/T}}{\sum_{j=1}^{5} e^{z_j(t)/T}}, \qquad w_{\min} \leq w_f \leq w_{\max}. \tag{5}$$

While no EMA has enough history (simulation start), weights default to $1/N$; each horizon enters the blend as soon as it matures, and the weights renormalise automatically. All parameters are fixed *a priori* (Table 4).

Table 4: Parameters of the dynamic weight-allocation algorithm.

Parameter	**Value**	**Role**
EMA short (h_S)	21 days	Fast reversal / incipient regime
EMA medium (h_M)	63 days	Central trend signal
EMA long (h_L)	252 days	Persistent trend
Blend weights (ω)	0.25 / 0.50 / 0.25	Short / medium / long
Softmax temperature (T)	0.8	Moderately assertive tilt
Min. weight ($w_{\min}$)	0%	Allows switching a factor off
Max. weight ($w_{\max}$)	50%	Prevents concentration

B Single-factor validation

Tables 5–8 report the performance, CAPM regressions, nested-model alpha progression and return correlations of the five single-factor portfolios (top 100 stocks each, semi-annual rebalancing, 2000–2026). All alphas use Newey–West HAC standard errors and survive a Benjamini–Hochberg correction for multiple testing.

Table 5: Performance metrics of the single-factor portfolios.

Metric	**Value**	**Quality**	**Momentum**	**Low Vol**	**Growth**
CAGR	11.93%	11.09%	10.68%	9.29%	10.31%
Volatility	21.66%	19.55%	19.27%	15.01%	19.52%
Sharpe	0.54	0.54	0.52	0.54	0.50
Sortino	0.69	0.71	0.67	0.68	0.64
Calmar	0.18	0.20	0.18	0.21	0.18
Max. drawdown	−66.1%	−55.5%	−57.8%	−43.9%	−58.4%
Win rate	53.5%	53.4%	54.2%	53.9%	53.9%
Annual turnover	0.49	0.84	1.21	0.58	1.13

Table 6: CAPM regression of the single-factor portfolios (Newey–West HAC).

Factor	**α (ann.)**	**t(α)**	**p-value**	**β**	**R^2**
Value	5.96%	2.94	0.003	0.99	0.783
Quality	5.03%	3.50	0.001	0.94	0.854
Momentum	4.78%	3.17	0.002	0.91	0.831
Low Volatility	4.22%	2.74	0.006	0.66	0.726
Growth	4.35%	3.11	0.002	0.94	0.859

Table 7: Annual alpha across nested models (p-value in parentheses).

Factor	**α CAPM**	**α FF3**	**α FF5**
Value	5.96% (0.003)	2.01% (0.111)	0.31% (0.797)
Quality	5.03% (0.001)	2.32% (0.038)	1.08% (0.301)
Momentum	4.78% (0.002)	2.32% (0.063)	1.16% (0.344)
Low Volatility	4.22% (0.006)	2.63% (0.079)	−0.02% (0.990)
Growth	4.35% (0.002)	1.79% (0.095)	0.67% (0.520)

Table 8: Correlation matrix of the single-factor daily returns.

	Value	**Quality**	**Mom.**	**Low Vol**	**Growth**
Value	1.00	0.95	0.89	0.85	0.92
Quality	0.95	1.00	0.93	0.87	0.96
Momentum	0.89	0.93	1.00	0.88	0.98
Low Volatility	0.85	0.87	0.88	1.00	0.89
Growth	0.92	0.96	0.98	0.89	1.00

C Combined-strategy factor loadings

Table 9 reports the five-factor loadings behind the alpha progression of Section 4.2. The inertia strategy carries essentially no size tilt ($\beta_{\mathrm{SMB}} = 0.01$, n.s.) and a slightly negative value tilt ($\beta_{\mathrm{HML}} = -0.09$), which is why its alpha survives the three-factor model; its return is instead absorbed by the profitability and, above all, the investment premium ($\beta_{\mathrm{CMA}} = 0.43$), locating the edge in a dynamic tilt towards high-quality, disciplined-capital firms.

Table 9: Fama–French five-factor loadings of the combined strategies (Newey–West HAC; p-value in parentheses).

Loading	**Equal Weight**	**Factor Inertia**
β_{MKT}	0.80 (0.000)	0.84 (0.000)
β_{SMB}	−0.04 (0.028)	0.01 (0.799)
β_{HML}	0.03 (0.188)	−0.09 (0.002)
β_{RMW}	0.25 (0.000)	0.17 (0.000)
β_{CMA}	0.38 (0.000)	0.43 (0.000)

D Hypothesis tests and robustness

The central hypothesis ($\mathrm{Sharpe}_{\mathrm{inertia}} > \mathrm{Sharpe}_{1/N}$) is tested with three methods that do not assume normal returns (Table 10). None reaches the conventional 5% threshold, so H_0 cannot be rejected; the spanning joint test comes closest at $p = 0.08$. As argued in Section 4.2, this low power is expected: the two strategies are 90% correlated and the system makes only 52 independent allocation decisions over the horizon.

Table 10: Tests of the Sharpe-ratio difference (inertia vs. equal weight).

Test	**Statistic**	**p-value**
Jobson–Korkie (Memmel correction)	$z = 0.81$	0.42
Block bootstrap (1,000 resamples)	95% CI $[-0.07, 0.26]$	0.23
Spanning regression (alpha)	2.39% ann.	0.12
Spanning regression (joint)	—	0.08

All algorithm parameters (EMA horizons, blend weights, softmax temperature, exposure caps) were fixed *a priori* from standard institutional practice and the momentum literature, with no optimisation on the backtest, which removes the risk of data snooping.