



Long-Short Equity Fund

Portfolio Management Division

May 2026

Fund Description

MIMS – Long Short Equity Fund is a semi-automated actively-managed fund by Minerva Investment Management Society, based on a zero-net investment ‘multi-factor’ strategy. The Fund has the investment objective of achieving a positive absolute return throughout all market conditions, maintaining a constant euro and geographical exposure at each rebalancing.

Market Update

In the US, growth remains resilient, supported by consumer spending and a still-solid labor market. Inflation has reaccelerated and remains above target (CPI ~3.3%), while policy rates are around 3.50–3.75%. Equity markets continue to trade near record highs, supported by strong earnings and AI-related momentum.

In the euro area, economic activity remains weak, with growth close to stagnation (~0.1% QoQ). Inflation has reaccelerated to around ~3.0%, with the ECB keeping policy rates near ~2.00%. However, weak industrial production and energy exposure continue to weigh on the outlook.

In China, growth has stabilized at around 5.0% YoY, but structural weaknesses, particularly in real estate, continue to be a notable constraint. Furthermore, rising geopolitical tensions cast a shadow over future growth prospects.

The current conflict in the Middle East represents a material upside risk to inflation. Any prolonged disruption to energy infrastructure or shipping routes could keep oil and gas prices elevated, slowing disinflation and delaying the pace of monetary easing by central banks. The IMF expects global inflation to tick up in 2026, while the World Bank projects energy prices to surge by about 24% this year under the current shock.

Overall, global markets remain divergent: U.S. strength supported by resilient growth and record-high equities, Europe constrained by weak growth and renewed inflation pressures, and China growing but still facing structural headwinds, with geopolitical risks increasingly shaping investor positioning.



Head of Portfolio Management

Aleksandar Georgiev: aleksandar.georgiev@studbocconi.it

Portfolio Manager

Federico Indrieri: federico.indrieri@studbocconi.it

Senior Portfolio Analyst

Ivan Yepifanov: ivan.yepifanov@studbocconi.it

Portfolio Analysts

Aigerim Turuspekova: aigerim.turuspekova@studbocconi.it

Annadele Tortora: annadele.tortora@studbocconi.it

Francesco Anastasi: francesco.anastasi@studbocconi.it

Giorgio Gagliardo: giorgio.gagliardo@studbocconi.it

Hristo Petrov: hristo.petrov@studbocconi.it

Kim Lan Vo: lan.vo@studbocconi.it

Polina Drofa: polina.drofa@studbocconi.it

Szofi Varga: szofi.varga@studbocconi.it

Factor Investing Strategy

VALUE

MOMENTUM

QUALITY

LOW VOLATILITY

TAIL RISK

YIELD CURVE

3 Steps Investment Approach

Multi Factor Analysis

Fundamental metrics are identified that best proxy each of the 6 factors on which the investment style is grounded. The process involves theoretical-based frameworks as well as empirical evaluations. Cross-team expertise and Minerva IMS insights are deployed.

Screening and Normalization

Stocks are evaluated on the basis of their exposure to each single factor. The output of the process is a set of factor exposures, which is then combined with the factor weights generated by the time-series factor momentum framework to derive the final stock weights.

Strategic Asset Allocation

Portfolio allocation comes to life. Based on the ranking produced, long and short positions are taken accordingly. Usually, no intermediate rebalancing is performed. Significant changes may lead to reconsider the chosen set of factors, or their weights, thus affecting the first step of the process.



DISCLAIMER - This is an academic paper related to an academic project that does not represent any investment recommendation nor offer any solicitation to buy or sell securities or to adopt an investment strategy. The opinions expressed are subject to change. References to specific securities, asset classes and financial markets are for illustrative purposes only and are not to be intended to and should not be interpreted as recommendations. Reliance upon information in this material is at the sole risk and discretion of the reader.

© Minerva Investment Management 2026. All rights are reserved

Performance Summary

	ITD ^{1,5}	Jan	Feb	Mar	Apr	May	Jun	Jul	Aug	Sep	Oct	Nov	Dec	YTD ⁵	Vol ⁵	Sharpe ^{2,5}
2021	7.9%											+0.9%	+6.9%	+7.9%	4.6%	1.68
2022	29.1%	+9.3%	+0.1%	+0.7%	+4.9%	+2.2%	+0.6%	-3.7%	+2.2%	+0.9%	+1.0%	-1.0%	+1.5%	+19.6%	10.1%	1.61
2023	29.8%	-2.2%	+0.3%	+2.5%	+1.6%	-0.4%	+0.1%	-3.8%	+5.2%	+1.6%	+1.7%	-4.0%	-1.8%	+0.5%	7.0%	-0.64
2024	46.5%	+6.6%	+0.9%	-0.3%	+1.2%	+0.9%	-1.7%	+0.5%	+0.7%	-2.4%	+3.1%	+1.6%	+1.4%	+12.9%	6.0%	1.20
2025	58.3%	-1.1%	+0.3%	+1.0%	-1.6%	+2.1%	+2.4%	+2.2%	+0.5%	+1.6%	+0.3%	-0.85%	+0.86%	+8.01%	5.5%	0.73
2026	52.8%	-1.37%	-1.04%	-1.09%	— ⁴									-2.17%	3.1%	-1.06

Given its dynamic, opportunity-driven mandate and its wide investment universe, the portfolio does not have a reference benchmark³ upon which to calculate the alpha.


Fund Information

Key Statistics (ITD)

Key Statistics (1Y)

Total Exposure	€100K	Ann. Return⁵	10.2%	Return⁵	7.4%
Total PnL 24.4.2026	€47,194.26	Ann. Volatility⁵	7.4%	Ann. Volatility⁵	5.3%
Rebalancing timeframe	Six months	Sharpe Ratio^{1,5}	0.82	Sharpe Ratio^{1,5}	0.65
Strategy	Long-Short	Max drawdown⁵	-6.5%	Max drawdown⁵	-5.5%
Universe	S&P500 and STOXX 600	Sortino Ratio^{2,5}	1.20	Sortino Ratio^{2,5}	0.90

Performance Overview



The graph represents the performance in € between 25 April 2025 – 24 April 2026, covering the last two rebalancing. The exposure was set again at €100,000 after rebalancing

Our timeframe covers the six months from November 24th, 2025, to April 24th, 2026. Over this period, the portfolio generated an absolute return of approximately **–€1,100** on a total exposure of **€100,000**, corresponding to around **–1.1% (–2.2% annualized)**. During this semester, the portfolio initially increased, reaching a peak in early January, before experiencing a significant decline throughout February. This period was particularly challenging, with losses across both the US and European legs, with the US portfolio representing the main negative contributor, driven by weaker performance in growth-oriented equities and increased market volatility during the mid-period downturn. However, a partial recovery was observed during March and April, although the portfolio remained below its previous highs by the end of the period. By dividing more deeply, we can see that the best performers were mainly represented by long positions. Looking at the **STOXX Europe 600 leg**, the strongest performances came from **Neste Oyj (long, +65.2%)**, followed by **Prysmian SpA (long, +53.5%)** and **Mycronic AB (long, +45.2%)**. The worst performers were instead **Hemnet Group AB (long, –34.9%)**, **Dino Polska SA (long, –20.2%)**, and **Balder AB (long, –16.9%)**. Looking at the **US leg**, the best performances were represented by **Deckers Outdoor Corp (long, +23.3%)**, **Regeneron Pharmaceuticals Inc (long, +4.2%)**, and **Assurant Inc (long, +4.0%)**. The worst performers were instead **Expeditors International (long, –5.0%)**, **Johnson & Johnson (long, –5.0%)**, and **Loews Corp (long, –4.5%)**, reflecting the weaker performance of the US portfolio. In this leg, it is possible to observe that both the best and worst performances were mainly represented by long positions.

1. ITD: Inception-to-date. Inception = 21/11/2021, Date = 24/04/2026.

2. Sharpe Ratio computed with US 3-Months Treasury Bills Yields as risk-free rate proxy.

3. Sortino Ratio computed with US 3-Months Treasury Bills Yields as risk-free rate proxy.

4. The monthly performance of April 2026 as of 24/04/2026 is 1.34%


5. The fund performance is calculated based on an evolving model with a simulated NAV.



DISCLAIMER - This is an academic paper related to an academic project that does not represent any investment recommendation nor offer any solicitation to buy or sell securities or to adopt an investment strategy. The opinions expressed are subject to change. References to specific securities, asset classes and financial markets are for illustrative purposes only and are not to be intended to and should not be interpreted as recommendations. Reliance upon information in this material is at the sole risk and discretion of the reader.

Performance – Top and Worst Performer


Top Performer: Neste Oyj (Long)



Neste Oyj is a leading producer of renewable fuels and refined oil products, with a focus on renewable diesel and sustainable aviation fuel, operating globally and benefiting from increasing demand for low-emission energy solutions. Over the past six months, the share price increased from approximately €16.1 to €26.6 (+65%), mainly driven by improving margins in renewable fuels and stronger earnings expectations following solid results at the end of 2025, with increased investor confidence in the long-term growth of sustainable energy further supporting the stock. This positive price development resulted in a strong contribution to the portfolio, making Neste the **best performer** within the long positions.

Source: LSEG Workspace, Total Return Index


Worst Performer: Hemnet Group AB (Long)



Hemnet Group AB operates Sweden's leading property platform, providing online real estate listings and related services. The company's performance is closely linked to housing market activity and pricing dynamics. During the semester, the total return index of the stock, expressed in euros, declined from approximately €176 to €115 (–35%), despite a slight appreciation of the Swedish krona over the period. This was driven mainly by uncertainty surrounding the company's pricing model transition, which affected short-term revenue visibility. Additionally, weaker housing market sentiment and lower transaction activity contributed to the negative performance, particularly during the mid-period downturn. As a result, Hemnet generated a negative contribution to the portfolio, making it the **worst performer**.

Source: LSEG Workspace, Total Return Index

Cumulative Performance in € (November 21, 2021 – April 24, 2026)



Source: Minerva Investment Management Society and LSEG Workspace. Past performance is not an indicator of future results.

Looking at cumulative performance since November 21st, 2021, the inception of the fund, the portfolio has reached a total cumulative performance of €47194.26. As in previous years, more than half of this substantial increase was concentrated in the first months of activity. Keeping in mind that the net invested capital in this fund is zero Euros, our benchmark is to deliver positive returns, an objective we failed to achieve this semester, but was effectively carried out in the overall life of the fund.



DISCLAIMER - This is an academic paper related to an academic project that does not represent any investment recommendation nor offer any solicitation to buy or sell securities or to adopt an investment strategy. The opinions expressed are subject to change. References to specific securities, asset classes and financial markets are for illustrative purposes only and are not to be intended to and should not be interpreted as recommendations. Reliance upon information in this material is at the sole risk and discretion of the reader.

© Minerva Investment Management 2026. All rights are reserved

Strategy

The Fund uses a multi-factor based investment style, adopting a quantitative proprietary model in order to achieve a systematic, rule-based approach to stock selection. Stocks are selected from the broad US Equity market (S&P 500 index) and the European Equity market (STOXX EUROPE 600 index).

This semester, the fund further advanced the transition initiated in the previous year by introducing a fully automated factor allocation framework. Factor and stock weights are now dynamically determined through a factor momentum methodology, and, in addition, the model was expanded with the introduction of a tail risk factor and supported by further statistical analysis

The model is based on a six-factor model, seeking to extract the risk premia coming from the size factor (Illiquidity, Amihud (2002)), value factor (HML, Fama & French (1992)), quality (QLT), momentum (MOM, Carhart (1997)), low volatility (LVOL, Baker, Bradley, and Wurgler (2011)) and yield curve factor ((BETA0, BETA1, BETA2), Nelson and Siegel (1987)). Explanations of each factor can be found under “Fund Factors”.

5 factors are constructed by ranking assets based on a specific metric (e.g., book-to-market ratio) that is presumed to reflect the factor of interest. After ranking, the assets are grouped into portfolios. At each time interval (depending on the factor), the assets are re-ranked and reassigned to the appropriate portfolios. The yield curve factor is constructed by extracting Beta coefficients from the Nelson-Siegel Model.

Following a fully automated allocation approach, factor weights are dynamically determined using a time-series factor momentum framework. After constructing the factors, we measure the recent performance of each of them and increase or reduce their allocation accordingly. This methodology is based on the idea that factors, like stocks, exhibit momentum, so recent factor performance can be used to improve portfolio construction.

The resulting factor views are then translated into stock weights by assigning positive or negative exposure to the stocks associated with each factor, depending on the factor's recent direction and adjusted for its volatility.

Once stock-level exposures are obtained, weights are aggregated across factors, so that stocks appearing in multiple factors receive the sum of their corresponding contributions.

After constructing the initial portfolio, qualitative checks are conducted to ensure sector neutrality and geographic balance. This is achieved by substituting companies from overrepresented sectors with those ranked next from underrepresented sectors. The final portfolio consists of 20 stocks in each segment (long/short, EU/US).

The portfolio is then normalized to derive the final stock allocations, while individual position weights are capped at 5% in order to avoid excessive concentration. This process results in a fully systematic and rules-based long-short portfolio construction framework.

Fund Factors

Value Factor – Long position

- This semester, we built our own Value factor instead of using the Fama-French HML series. It's based on a simple book-to-market ratio (B/M), calculated from company fundamentals. Each month, stocks are split by market capitalization (Small vs Big) and ranked by their book-to-market values.

Momentum Factor – Long position

- MOM: following the evidence provided by Jegadeesh and Titman (1993), Asness (1994), and Carhart (1997), we consider momentum defined as the cumulative return of the last 12 months excluding the most recent one divided by 11, as a buy signal. We assume, as a practical heuristic, that stocks with a positive average return over recent months will continue to perform well in the near future.

Quality Factor – Long position

- Return on Equity (ROE): we consider a high ROE as a signal of high profitability and thus a buy signal. Specifically, we are assuming that the company's profitability will stay stable in the future and will be a reliable driver of increases in prices.
- 5-year growth in ROE: to account for the growth, we assess the earnings increase over the last five years relative to the equity's book value from five years ago. This allows us to reward companies that showed an increase in profitability while smoothing earnings by considering a 5-year window.
- Debt-to-Equity (D/E): for the safety dimension of our quality factor, we consider the D/E ratio. A high D/E ratio indicates an excessive level of debt for the firm, representing a risk and also inflating ROE when earnings are positive.
- Payback Period (Total Liabilities/CFO): Measures the time a company takes to generate enough cash flow to repay its debt, with shorter periods indicating stronger stability.
- Earnings Volatility (EVOL): reflects fluctuations in quarterly ROE, with lower volatility displaying consistent performance.

Low Volatility Factor – Long position

- Standard deviation: we deem a higher standard deviation to be a sell signal, since it reveals a riskier situation where returns are less stable, and, consequently, less predictable Baker, Bradley, and Wurgler (2011).

Tail risk factor – Long position

- Based on Kelly and Jiang (2014), calculated using a pooled Hill estimator applied to the cross-section of monthly returns. It captures the thickness of the common left tail of the return distribution, where a higher λ_t implies fatter tails and greater aggregate tail risk across the market.
- Each month, stocks are ranked by their rolling tail beta — estimated via a 60-month predictive OLS regression of returns on the standardised tail risk signal $\lambda_{z(t)}$. We go long the lowest tail-beta decile and short the highest tail-beta decile.

Yield Curve Factor– Long/Short position

- The Nelson–Siegel model decomposes the yield curve into three key factors: β_0 (level), capturing parallel shifts in long-term rates; β_1 (slope), reflecting steepening or flattening driven by monetary-policy expectations; and β_2 (curvature), measuring movements in the mid-term segment of the curve. These factors reveal how stocks react to structural changes across the term structure.

Factors and Stocks weighting: Time Series Factor Momentum

In the current version of the fund, we have decided to move this semi-automated, actively managed fund closer to being a fully automated fund. One of the reasons for the semi-automated construction was the fact that, in order to construct the final portfolios, the portfolio management team was relying on the macroeconomic research and their own discretionary opinion to generate an alpha by choosing the weights of the factors, which were then used to pick stocks. In this way, by using discretionary choices, the final stocks were selected, and then the portfolios of stocks were optimized using Hierarchical Risk Parity and, before that, Mean-Variance Optimization approaches. This construction has performed well, as seen from our past performance. Nevertheless, the next natural step for improvement is the optimization of the factor weights.

The current research allows for a parsimonious yet effective way to dynamically adjust the weights of the factors to construct the factor allocations. Traditionally, we can use the momentum factor, which we are already using for our stock selection, to buy stocks that have gone up recently and sell those that have gone down. This has been widely studied and has proven to be effective in generating positive returns. In the same vein, the paper "Factor Momentum Everywhere" by Tarun Gupta and Bryan Kelly introduced a very similar idea. The key argument of this paper is that instead of applying momentum to stocks, we can apply it to factors and get similar results.

This paper shows that factors themselves exhibit momentum, which we can exploit in order to construct our long-short portfolios. Moreover, in terms of the Sharpe ratio, a portfolio strategy that buys the recent top-performing factors and sells poor-performing factors outperforms stock momentum, industry momentum, value, and other commonly studied investment factors. Moreover, the paper argues that this strategy is more stable than the stock momentum and therefore different formation windows, such as 1 month (best), 1 year, or even 5 years, all work, and both strategies are complementary. In other words, by controlling for the stock momentum, the factor momentum still exists. Finally, such a factor momentum exists globally, both in the US and EU markets, and therefore can be applied for our long-short portfolio construction. Because of these advantages, we decided to implement this strategy to improve our portfolio construction process.

To implement this strategy, we first construct the factors as before, then we measure the recent performance of each factor, and afterwards increase or reduce the allocation to each factor based on the previous performance. In the paper, such a construction is called time-series factor momentum (TSFM). This strategy looks at each factor individually and decides to go long or short based on its own history. The authors propose another strategy called Cross-Sectional Factor Momentum (CSFM), which looks at each factor's performance relative to other factors. However, the TSFM specification is preferred to the CSFM due to better performance.

In order to implement this idea, we have developed a practical implementation of TSFM at the stock level. We first input the factor weights, then we convert the factor view into the stock weights by using the following formula:

$$w_{i,j,t+1} = f_{i,j,t+1} \left(\frac{1}{N} \right) * (-1)^{1_s}$$

Where w_i is the stock weight, f_j is the factor weight from the TSFM, N is the number of stocks, and 1_s is the indicator of the short direction that takes the values of 0 and 1. Specifically, the factor weight from the TSFM is obtained using the following formula:

$$f_{i,j,t+1} = s_{i,j,t} * f_{i,t+1}$$

and $s_{i,j,t}$ is a scaling term which is calculated as:

$$s_{i,j,t} = \min\left(\max\left(\frac{1}{\sigma_{i,j,t}} \sum_{\tau=1}^j f_{i,t-\tau+1}, -2\right), 2\right)$$

This scaling term uses the formation period from $t-j$ to t , and if formation returns are positive, it buys the factor, while if they are negative, it sells the factor. The scaling term also uses the $\sigma_{i,j,t}$, which is the annualized factor volatility, over the previous three years.

Intuitively, if the factor is strong, its long stocks get positive weight, and its short stocks get negative weight. We then aggregate across factors, and if the stock appears in multiple factors, its weights are summed. By normalizing the portfolios, we thus get the final stock allocations. In the same way as the original paper, we cap the weight of each stock at 5% in order to avoid concentration. Hence, at the end, we obtain the final portfolios and thus allocate by leveraging the time-series factor momentum strategy.

US: Factor Weights

Value	Tail Risk	Momentum	Quality	Low Volatility	BETA 0	BETA 1	BETA 2
+23.95%	+19.30%	+2.08%	-0.87%	-15.34%	-4.55%	-27.27%	-6.63%

EU: Factor Weights

Value	Tail Risk	Momentum	Quality	Low Volatility	BETA 0	BETA 1	BETA 2
+20.15%	-5.53%	+20.39%	-18.29%	-1.78%	+20.39%	-6.68%	+6.80%




DISCLAIMER - This is an academic paper related to an academic project that does not represent any investment recommendation nor offer any solicitation to buy or sell securities or to adopt an investment strategy. The opinions expressed are subject to change. References to specific securities, asset classes and financial markets are for illustrative purposes only and are not to be intended to and should not be interpreted as recommendations. Reliance upon information in this material is at the sole risk and discretion of the reader.

Correlation Heatmaps

Pairwise correlations among the eight style factors remain generally low, with only a few values showing extreme correlations up to ± 0.85 .

The strongest positive relationship is observed between Beta 0 and Beta 1 in the US leg ($\rho = +0.76$). Value and Tail Risk in the European leg are the most negatively correlated ($\rho = -0.84$).



Principal component analysis – Theoretical background

Why PCA? The fund uses factors (value, momentum, quality, etc.) to explain stock returns. When factors are highly correlated, they carry redundant information, inflating coefficient variances and making individual factor contributions hard to isolate. PCA is used here not for dimensionality reduction, but as a diagnostic tool to detect this multicollinearity.

What PCA does: PCA finds a new set of axes – Principal Components (PCs) – that are uncorrelated linear combinations of the original factors, ordered so that each one captures as much remaining variance as possible. This is done via the eigendecomposition of the covariance matrix of the standardised data. Each eigenvalue measures how much variance a PC explains; its eigenvector (the loadings) shows which original factors drive it. Standardising the data beforehand is essential to prevent factors with larger numerical scales from dominating.

Condition indices and VDP

Condition Indices & VDP: The Condition Index of a PC is $\sqrt{\lambda_1/\lambda_2}$ – the ratio of the largest eigenvalue to that component's. A high CI on a late PC signals near-collinearity. To identify which factors are involved, the Variance Decomposition Proportions (VDP) matrix breaks down each factor's coefficient variance across PCs. If two or more factors concentrate their variance on the same high-CI component, they are the collinear pair (Belsley, Kuh & Welsch, 1980).

Variance Inflation Factor

The Variance Inflation Factor for each factor is $1 / (1 - R^2)$, where R^2 comes from regressing that factor on all the others. A high R^2 means the factor is largely predictable from its peers – i.e., redundant. VIF above 10 is the standard concern threshold. It serves as a simpler, factor-level complement to the CI/VDP analysis.

Diagnostics data

All the data tables referenced are on the following page.

STOXX 600

Variance structure: PC1 alone explains 48.5% of total variance, and the first two components together account for 68.3%. A large share of EU factor co-movement is driven by just one or two underlying dimensions, consistent with a macro environment where the European rate cycle pulls several factors simultaneously — something the correlation matrix makes visually clear.

What PC1 captures: PC1 loads negatively on BETA0_EU, VAL_EU, and TAIL_EU, and positively on QLT_EU, MOM_EU, and LVOL_EU. This maps neatly onto the correlation heatmap: VAL_EU and QLT_EU are strongly negatively correlated (-0.84), and TAIL_EU is negatively correlated with both QLT_EU (-0.58) and MOM_EU (-0.40). PC1 is essentially capturing a macro sensitivity axis: on one side, factors tied to yield curve level and value; on the other, quality, momentum, and low volatility. The positioning of TAIL_EU on the yield-curve/value side is notable — since the factor goes long low-tail-beta stocks and short high-tail-beta ones, its alignment here likely reflects a temporary overlap between low-tail-sensitivity names and value-oriented stocks in the current European environment rather than a stable structural relationship.

What PC2 captures: PC2 is dominated by BETA2_EU (0.69) and BETA1_EU (0.44), reflecting the strong positive correlation between the two yield curve shape factors ($\rho = 0.44$). This is a yield curve geometry dimension largely orthogonal to the equity-style factors.

The VAL–QLT relationship: The most striking pairwise correlation in the heatmap is VAL_EU vs. QLT_EU at -0.84 . In the current European environment, cheap (high book-to-market) stocks tend to be lower-quality names, and vice versa — a well-known tension in factor investing that becomes particularly pronounced during periods of macro stress. Despite this strong correlation, the VIF for both factors remains below 5, confirming that the relationship, while economically meaningful, does not rise to the level of harmful multicollinearity.

Multicollinearity — clean bill of health: The overall condition number is 5.27, and all condition indices are LOW. The VDP matrix shows that VAL_EU and BETA0_EU concentrate most of their variance on the smallest PCs, but those components carry negligible CIs, meaning they occupy idiosyncratic dimensions rather than collinear ones. The mild co-loading of MOM_EU and TAIL_EU on PC4 is similarly harmless, given that component's CI of 2.51.

S&P 500

Variance structure: PC1 explains 58.6% of total variance, and the first two components together account for 74.8%. The US factor set is strongly dominated by a single common theme, consistent with a market environment where Fed policy expectations have been the overriding driver pulling most factors simultaneously.

What PC1 captures: PC1 loads positively on BETA0_US, TAIL_US, and VAL_US, and negatively on MOM_US, QLT_US, and LVOL_US — capturing a macro sensitivity axis broadly similar to the EU. The correlation heatmap makes this tangible: TAIL_US is strongly positively correlated with BETA0_US (0.74) and VAL_US (0.67), while being strongly negatively correlated with QLT_US (-0.63) and LVOL_US (-0.83). As with the EU leg, the positioning of TAIL_US alongside value and yield curve level — rather than the defensive cluster where its construction would suggest it belongs — likely reflects a period-specific overlap between low-tail-sensitivity stocks and value names in the US market, and should be monitored for stability in future periods.

What PC2 captures: PC2 is dominated by BETA1_US (0.50) and BETA2_US (0.50), reflecting a yield curve shape dimension independent of the equity-style factors, supported by the strong BETA0/BETA1 correlation of 0.76.

The VAL–QLT–LVOL cluster: The most structurally notable feature of the heatmap is the tight triangle of VAL_US, QLT_US, and LVOL_US. VAL and QLT are strongly negatively correlated (-0.75), VAL and LVOL similarly so (-0.83), while QLT and MOM move together (0.64). This cluster is the primary driver of the elevated VIF readings.

Multicollinearity — moderate flags worth monitoring. The overall condition number of 7.97 remains below 10, and all condition indices are LOW, so no structural collinearity is present. However, BETA0_US (VIF = 9.42) and TAIL_US (VIF = 5.41) both register as moderate concern. The VDP matrix clarifies why: BETA0_US concentrates 95.7% of its coefficient variance on PC8, and TAIL_US concentrates 55.7% on PC7 — both late, low-eigenvalue components approaching but not breaching the critical CI threshold. This means each factor is being meaningfully explained by a combination of the others, inflating their VIFs without triggering a formal collinearity diagnosis. It is not an acute problem, but worth keeping under observation as the sample window extends.

STOXX 600 Data

EU Factor PCA — VIF Analysis						EU Factor PCA — Loadings Matrix								
Factor	R²	VIF	Tolerance	Concern	Interpretation	Each entry shows how strongly a factor loads onto a given Principal Component. Larger absolute values indicate greater influence.								
BETA0_EU	0,72	3,55	0,28	LOW	Highest R² — shared with yield curve cluster	Factor	PC1	PC2	PC3	PC4	PC5	PC6	PC7	PC8
BETA1_EU	0,44	1,77	0,56	LOW	Well-isolated slope factor	BETA0_EU	-0,40	0,35	-0,18	0,28	-0,15	-0,11	0,70	-0,30
BETA2_EU	0,59	2,45	0,41	LOW	Moderate overlap with BETA0/BETA1	BETA1_EU	-0,25	0,44	0,61	-0,21	0,35	0,43	0,07	0,13
MOM_EU	0,47	1,89	0,53	LOW	Largely independent	BETA2_EU	-0,15	0,69	-0,25	-0,03	0,01	-0,42	-0,48	0,16
QLT_EU	0,77	4,38	0,23	LOW	Highest VIF — driven by VAL/QLT tension	MOM_EU	0,35	0,26	0,15	0,70	-0,36	0,37	-0,17	0,07
TAIL_EU	0,47	1,88	0,53	LOW	Largely independent	QLT_EU	0,46	0,10	0,21	-0,03	0,02	-0,42	0,46	0,59
VAL_EU	0,77	4,35	0,23	LOW	High VIF — driven by VAL/QLT tension	TAIL_EU	-0,34	-0,29	0,22	0,61	0,48	-0,34	-0,13	0,12
LVOL_EU	0,45	1,81	0,55	LOW	Largely independent	VAL_EU	-0,43	-0,15	-0,37	0,02	-0,17	0,35	0,05	0,71
						LVOL_EU	0,35	0,15	-0,53	0,11	0,68	0,29	0,13	-0,02

EU Factor PCA — Variance Decomposition Proportions (VDP)									EU Factor PCA — Explained Variance				EU Factor PCA — Condition Indices				
Each row sums to 1. High proportion on a high-CI component (paired with another factor) signals collinearity. All EU CIs are LOW — no collinearity flagged.									Overall Condition Number: 5.2664 — Multicollinearity Concern: LOW (threshold: <10 LOW, 10–30 MODERATE, >30 HIGH)								
Factor	PC1	PC2	PC3	PC4	PC5	PC6	PC7	PC8	Component	Var. Ratio	Cumul. Var.	Incremental	Component	Eigenvalue	Condition Index	Warning	Threshold
BETA0_EU	0,01	0,02	0,01	0,04	0,01	0,01	0,72	0,18	PC1	48,47%	48,47%	48,47%	PC1	3,92	1,00	LOW	<10
BETA1_EU	0,01	0,07	0,24	0,04	0,16	0,40	0,02	0,07	PC2	19,85%	68,31%	19,85%	PC2	1,60	1,56	LOW	<10
BETA2_EU	0,00	0,12	0,03	0,00	0,00	0,28	0,49	0,08	PC3	10,98%	79,30%	10,98%	PC3	0,89	2,10	LOW	<10
MOM_EU	0,02	0,02	0,01	0,42	0,15	0,28	0,08	0,02	PC4	7,68%	86,97%	7,68%	PC4	0,62	2,51	LOW	<10
QLT_EU	0,01	0,00	0,01	0,00	0,00	0,16	0,26	0,56	PC5	5,64%	92,61%	5,64%	PC5	0,46	2,93	LOW	<10
TAIL_EU	0,02	0,03	0,03	0,32	0,27	0,23	0,05	0,05	PC6	3,24%	95,85%	3,24%	PC6	0,26	3,87	LOW	<10
VAL_EU	0,01	0,00	0,04	0,00	0,01	0,11	0,00	0,83	PC7	2,40%	98,25%	2,40%	PC7	0,19	4,49	LOW	<10
LVOL_EU	0,02	0,01	0,18	0,01	0,56	0,17	0,05	0,00	PC8	1,75%	100,00%	1,75%	PC8	0,14	5,27	LOW	<10

S&P 500 Data

US Factor PCA — VIF Analysis						US Factor PCA — Loadings Matrix								
VIF = 1/(1-R²). Threshold: <5 LOW 5–10 MODERATE >10 HIGH. BETA0_US and TAIL_US show MODERATE concern — monitor but not a formal collinearity diagnosis.						Each entry shows how strongly a factor loads on to a given Principal Component. Larger absolute values indicate greater influence.								
Factor	R²	VIF	Tolerance	Concern	Interpretation	Factor	PC1	PC2	PC3	PC4	PC5	PC6	PC7	PC8
BETA0_US	0,89	9,42	0,11	MODERATE	Highest R²/VIF — 95.7% VDP on PC8 (CI=7.97);	BETA0_US	0,41	0,32	0,02	-0,07	-0,02	0,22	0,08	0,82
BETA1_US	0,71	3,50	0,29	LOW	Some overlap with yield curve cluster	BETA1_US	0,29	0,50	-0,40	-0,43	-0,36	0,10	-0,16	-0,39
BETA2_US	0,61	2,55	0,39	LOW	Moderate overlap with BETA0/BETA1 slope	BETA2_US	0,25	0,50	0,61	0,34	-0,06	-0,31	0,19	-0,25
MOM_US	0,54	2,17	0,46	LOW	76.3% VDP on PC5 — largely idiosyncratic	MOM_US	-0,32	0,46	-0,06	-0,28	0,76	-0,15	-0,07	0,02
QLT_US	0,67	3,04	0,33	LOW	52.6% VDP on PC6 VAL/QLT tension	QLT_US	-0,36	0,35	-0,10	0,57	-0,07	0,58	-0,25	-0,04
TAIL_US	0,82	5,41	0,18	MODERATE	MODERATE — 55.7% VDP on PC7 (CI=6.47);	TAIL_US	0,41	-0,10	-0,27	0,20	0,41	0,35	0,58	-0,29
VAL_US	0,78	4,52	0,22	LOW	59.9% VDP on PC7 —VAL/QLT/LVOL cluster	VAL_US	0,38	-0,21	0,45	-0,19	0,28	0,40	-0,55	-0,18
LVOL_US	0,74	3,78	0,26	LOW	52.7% VDP on PC7 —VAL/QLT/LVOL cluster	LVOL_US	-0,38	0,04	0,42	-0,46	-0,19	0,45	0,47	-0,04

US Factor PCA — Variance Decomposition Proportions (VDP)									US Factor PCA — Explained Variance				US Factor PCA — Condition Indices				
Each row sums to 1. BETA0_US: 95.7% on PC8 (CI=7.97) and TAIL_US: 55.7% on PC7 (CI=6.47) explain the MODERATE VIF flags. All CIs remain LOW — no formal collinearity diagnosis.									Overall Condition Number: 7.9651 — Multicollinearity Concern: LOW (threshold: <10 LOW, 10–30 MODERATE, >30 HIGH) Note: PC8 CI approaching 10 — monitor								
Factor	PC1	PC2	PC3	PC4	PC5	PC6	PC7	PC8	Component	Var. Ratio	Cumul. Var.	Incremental	Component	Eigenvalue	Condition Index	Warning	Threshold
BETA0_US	0,00	0,01	0,00	0,00	0,00	0,02	0,01	0,96	PC1	58,63%	58,63%	58,63%	PC1	4,740535	1,000000	LOW	<10
BETA1_US	0,01	0,06	0,06	0,12	0,11	0,01	0,06	0,58	PC2	16,16%	74,79%	16,16%	PC2	1,306760	1,904653	LOW	<10
BETA2_US	0,01	0,08	0,18	0,10	0,00	0,18	0,13	0,33	PC3	10,06%	84,85%	10,06%	PC3	0,813831	2,413495	LOW	<10
MOM_US	0,01	0,08	0,00	0,08	0,76	0,05	0,02	0,00	PC4	5,81%	90,66%	5,81%	PC4	0,469885	3,176273	LOW	<10
QLT_US	0,01	0,03	0,00	0,23	0,00	0,53	0,19	0,01	PC5	4,34%	95,00%	4,34%	PC5	0,351028	3,674877	LOW	<10
TAIL_US	0,01	0,00	0,02	0,02	0,09	0,10	0,56	0,21	PC6	2,67%	97,68%	2,67%	PC6	0,216058	4,684124	LOW	<10
VAL_US	0,01	0,01	0,06	0,02	0,05	0,17	0,60	0,10	PC7	1,40%	99,08%	1,40%	PC7	0,113203	6,471205	LOW	<10
LVOL_US	0,01	0,00	0,06	0,12	0,03	0,25	0,53	0,01	PC8	0,92%	100,00%	0,92%	PC8	0,074721	7,965109	LOW	<10





DISCLAIMER - This is an academic paper related to an academic project that does not represent any investment recommendation nor offer any solicitation to buy or sell securities or to adopt an investment strategy. The opinions expressed are subject to change. References to specific securities, asset classes and financial markets are for illustrative purposes only and are not to be intended to and should not be interpreted as recommendations. Reliance upon information in this material is at the sole risk and discretion of the reader.

Macro-Yield Curve Modeling (Nelson-Siegel Betas)

Movements of the yield curve, whether parallel shifts (Level), steepening (Slope), or changes in curvature, summarize both the stance of monetary policy and expectations on future growth and inflation. Firms with different profiles respond differently to these changes due to various characteristics, such as leverage and discount-rate sensitivity. By quantifying such sensitivities, we can construct a portfolio that systematically earns returns when those macro shocks occur by translating the dynamics of interest rates into cross-sectional equity payoffs

The Nelson–Siegel model provides a smooth parametric representation of the yield curve using only three parameters: β_0 (Level): parallel shifts of the curve; β_1 (Slope): short vs. long-term spread; β_2 (Curvature): changes in the belly of the curve. This interpolation method fills the gaps by generating a continuous and coherent curve across all maturities. It solves the problem of comparing yield dynamics consistently over time, turning scattered bond data into an analytically tractable structure.



Methodology


1. Estimate monthly Nelson–Siegel betas (β_0 , β_1 , β_2) for US and EU yield curves using the functional form:

$$\gamma(\tau) = \beta_0 + \frac{1 - e^{-\lambda\tau}}{\lambda\tau} + \beta_2 \left(\frac{1 - e^{-\lambda\tau}}{-\lambda\tau} \right)$$

2. Compute monthly shocks $\Delta\beta_0$, $\Delta\beta_1$, $\Delta\beta_2$, defined as changes in each factor between consecutive months, capturing yield-curve movements.

3. For each stock, it is computed a rolling 60-month OLS of excess returns on the three factors at the same time, in order to reduce the risk of incorporating variations attributable to others, and hence reducing multicollinearity.


Each coefficient (γ) measures how much a stock reacts to a one-unit shock in that yield-curve component and thanks to these results the long and short legs for each factor can be constructed by taking respectively the deciles with the most positively and negatively sensitive stocks for each beta at each point in time: $r_{i,t}^E = d_i + \gamma_i^{(k)} \Delta\beta_{k,t} + \varepsilon_{i,t}$, $k \in \{0,1,2\}$



This approach ensures that the factor's returns are driven by yield-curve movements. By design, long and short legs have equal dollar exposure, making the strategy approximately market-neutral. When the Level factor rises (yields increase), the portfolio gains because it is long in stocks positively exposed to rate shocks and short in those less sensitive. Conversely, during negative yield shocks, shorting the factor generates positive performance. The portfolio was subject to yearly rebalancing, as in classic Fama-French factor construction.

The final cumulative performance of the Level Factor confirms the macro-consistency of the strategy. The 2020 drawdown reflects the collapse in global yields during the Covid crisis, while the recovery mirrors the normalization of rates through 2025.

The NS Yield Curve Factor captures the equity market's response to yield-curve shifts, turning interest-rate movements into tradable signals.



DISCLAIMER - This is an academic paper related to an academic project that does not represent any investment recommendation nor offer any solicitation to buy or sell securities or to adopt an investment strategy. The opinions expressed are subject to change. References to specific securities, asset classes and financial markets are for illustrative purposes only and are not to be intended to and should not be interpreted as recommendations. Reliance upon information in this material is at the sole risk and discretion of the reader.

Factors – Quality Factor

Building on the findings of Quality Minus Junk (Asness, Frazzini & Pedersen, 2017), we include a strengthened Quality factor that integrates five components (Earnings Volatility (EVOL), ROE, ROE-growth, Debt-to-Equity, and Safety) and that allocates weights to emphasize some selected aspects.

The most recent addition remains EVOL, computed as the standard deviation of quarterly ROE over the most recent twenty observations, providing a straightforward insight into the predictability of a firm's earnings and profitability (Figure 1).

Firms scoring in the lowest decile of EVOL generate extremely stable profits that are, hence, usually less subject to industry-wide and macroeconomic downturns. Conversely, elevated EVOL indicates an erratic earnings pattern, and in such cases, the stock is systematically penalized in our z-score composite.

By rewarding low-volatility earners, the updated Quality factor can also capture the natural shift toward predictable investments that becomes prominent during macro stress and times of uncertainty.

The safety metric is computed as the ratio of Operating Cash Flow (OCF) to total liabilities (TL), capturing a firm's ability to generate sufficient internal cash to meet its financial obligations. A higher ratio suggests that the company can comfortably cover its liabilities through ongoing operations, signaling lower default risk and stronger balance sheet stability.

The approach for this semester retains these five pillars and weights them as follows: ROE 24 %, ROE-growth 22 %, EVOL 22 %, Debt-to-Equity 16 %, and Safety 16%.

This distribution reflects a broad industry consensus: a survey of leading asset managers (BlackRock, Vanguard, and Fidelity among others) shows that profitability and earnings stability consistently outrank raw leverage in contemporary quality screens (Figure 2).

Furthermore, during times of economic uncertainty (as opposed to monetary instability), profitability becomes an even more valuable component within the quality factor. Firms that consistently generate strong profits are more likely to maintain operational resilience, reinvest in growth, and weather downturns without resorting to external financing or cost-cutting that might impair long-term value. High profitability often signals durable competitive advantages, pricing power, and efficient management — traits that offer stability in unpredictable macroeconomic environments.

While leverage and balance sheet strength are crucial, especially during periods of monetary tightening, economic uncertainty calls for a focus on firms that can internally sustain themselves through robust income generation. As such, assigning a greater weight to profitability within the quality composite can enhance a portfolio's defensive profile and return consistency across cycles.

To reflect this priority and to align closer with the industry standard, we have hence slightly adjusted the equal-weighted allocation of the quality subcomponents, redistributing weights in favor of those more aligned with profitability, as shown earlier.

$$\text{Earnings Volatility} = \sqrt{\frac{1}{20} \sum_{i=1}^{20} (ROE_i - \overline{ROE})^2}$$

FIGURE 1: Earnings Volatility Formula

Source:

EVOL (Earnings Volatility) is the standard deviation of ROE calculated over a 20 quarters (5-year) basis. It aims to capture the stability and predictability of profitability over time.




FIGURE 2: Usage of quality sub-factors amongst asset management firms

Source:

Personal research conducted on the construction factors of various quality funds of leading asset management firms.

The figure depicts the frequency in which each quality measure is used to evaluate the stocks in the various analyzed funds.

The firms whose funds were chosen are: BlackRock; AQR; Vanguard; Fidelity; J.P. Morgan, Capital Group, Invesco, Wisdom Tree, Northern Trust, Amundi.

One quality fund for each firm was chosen as a sample.



DISCLAIMER - This is an academic paper related to an academic project that does not represent any investment recommendation nor offer any solicitation to buy or sell securities or to adopt an investment strategy. The opinions expressed are subject to change. References to specific securities, asset classes and financial markets are for illustrative purposes only and are not to be intended to and should not be interpreted as recommendations. Reliance upon information in this material is at the sole risk and discretion of the reader.

Factors – Tail risk factor

Tail risk measures how exposed a stock is to rare, severe market crashes. Kelly and Jiang (2014) show that this risk is time-varying and common across assets: during stress, the left tail of returns becomes fatter for all stocks simultaneously. Stocks that co-move most strongly with this aggregate tail risk, those with high tail betas, should earn a premium to compensate investors for bearing that exposure.

Empirically, however, the opposite holds: high-tail-beta stocks systematically underperform low-tail-beta stocks. This anomaly reflects demand-driven overpricing of crash-sensitive assets, where investors overpay for lottery-like payoffs or perceived crash hedges. The factor captures this persistent mispricing through a monthly long-short decile strategy.


The factor is constructed in three steps applied to the STOXX Europe 600 and S&P 500:

- **Pooled Hill Estimator:** at each month t , the 5th-percentile left-tail threshold u_t is computed from all stock returns over a rolling 12-month pooled window. The Hill estimator λ_t measures the thickness of the common left tail as the average log-exceedance ratio across all returns below u_t .

- **Rolling Tail Betas:** the standardized signal $\lambda_{z(t)}$ is regressed on each stock's next-period return over a rolling 60-month OLS window. The slope β_i captures each stock's sensitivity to changes in aggregate tail risk: $r_{i,t+1} = \alpha_i + \beta_i \cdot \lambda_{z(t)} + \varepsilon$.

- **Long-Short Portfolio:** stocks are ranked monthly by β_i . The strategy goes long the bottom decile (lowest tail beta) and short the top decile (highest tail beta), forming an equal-weighted, zero-investment portfolio rebalanced monthly.

The factor delivers a positive unconditional premium over the full sample period, with a Sharpe ratio of 0.36 and low correlation to the market ($\rho \approx -0.05$), confirming its diversifying role within the multi-factor framework. Its premium is strongly concentrated in high-volatility regimes (Sharpe = 0.84 vs -0.11 in calm markets), consistent with the theoretical prediction that tail-risk pricing is most acute when investors are most fearful.



Factors – Value Factor

In earlier versions of the model, the fund relied on the traditional Fama–French High-Minus-Low (HML) factor to capture the value premium. HML sorts stocks by book-to-market (B/M) and takes a long position in “cheap” firms (high B/M) and a short position in “expensive” firms (low B/M). However, over the last decade, the classical value premium has weakened in several markets, especially during long periods of low interest rates, high sector concentration, and the rising importance of intangible-heavy companies. Under these conditions, the standard HML series tends to reflect the structure of the broad CRSP universe rather than the characteristics of the indices in which the fund actually invests.

To address this issue, we construct our Value factor directly from the S&P 500 and STOXX Europe 600. This ensures that the factor reflects only the stocks we can trade and avoids exposures that do not exist in our investment universe. Our factor follows the standard academic B/M definition but aligns it with the fund’s own investment universe, risk framework, and portfolio-construction process. The factor is built monthly using the formula:

$$Book\ Equity_i = BVPS_i \times Shares\ Outstanding_i$$

$$Value\ Signal_i = \frac{B}{M}_i = \frac{Book\ Equity_i}{Market\ Cap_i}$$

To remove the interaction between Value and Size, the cross-section is split by median market capitalization into Small and Big buckets. The median-split procedure ensures the Value factor captures fundamental cheapness and not exposure to small-cap risk. Within each bucket, we identify the top 20% high-B/M stocks (Value) and bottom 20% low-B/M stocks (Growth). These long and short selections are then combined into one size-neutral HML-style portfolio. Signals are lagged by one month, and positions are held for twelve months in overlapping cohorts. It reduces noise from reporting delays and makes the factor more realistic to implement.

Book-to-Market remains a meaningful indicator of valuation and captures a risk premium linked to economic uncertainty, earnings sensitivity, and valuation mean-reversion. High-B/M firms usually trade at lower expectations and can be more exposed to shifts in macro conditions, interest-rate movements, and inflation trends. Recent market behavior highlights these differences across regions.

- In the US, the cumulative performance of the factor has been weak, which reflects the strong dominance of mega-cap growth stocks and the prolonged underperformance of cyclical sectors.
- In Europe, where the sector composition is more value-oriented (financials, industrials, energy), the factor displayed a meaningfully stronger recovery during 2021–2026.

Our empirical results confirm these macro patterns.

Cumulative Value Factors



DISCLAIMER - This is an academic paper related to an academic project that does not represent any investment recommendation nor offer any solicitation to buy or sell securities or to adopt an investment strategy. The opinions expressed are subject to change. References to specific securities, asset classes and financial markets are for illustrative purposes only and are not to be intended to and should not be interpreted as recommendations. Reliance upon information in this material is at the sole risk and discretion of the reader.

© Minerva Investment Management 2026. All rights are reserved

New Fund Positioning


S&P 500

Intel Corp	3.74%	Erie Indemnity Co	-3.59%
Tyson Foods Inc	3.55%	PTC Inc	-3.38%
Align Technology Inc	3.07%	Fair Isaac Corp	-3.12%
Newmont Corporation	2.80%	Super Micro Computer Inc	-2.89%
Micron Technology Inc	2.71%	Trade Desk Inc	-2.89%
Seagate Technology Holdings PLC	2.70%	EMCOR Group Inc	-2.67%
Western Digital Corp	2.64%	McKesson Corp	-2.67%
Edwards Lifesciences Corp	2.47%	Norwegian Cruise Line Holdings Ltd	-2.63%
International Flavors & Fragrances Inc	2.45%	Uber Technologies Inc	-2.61%
Estee Lauder Companies Inc	2.40%	Gartner Inc	-2.56%
Stanley Black & Decker Inc	2.34%	Hubbell Inc	-2.48%
Expedia Group Inc	2.32%	CDW Corp	-2.45%
Revvity Inc	2.26%	Rockwell Automation Inc	-2.41%
Ciena Corp	2.19%	Phillips 66	-2.36%
Kenvue Inc	2.19%	Old Dominion Freight Line Inc	-2.04%
Lumentum Holdings Inc	2.19%	Occidental Petroleum Corp	-1.71%
EchoStar Corp	2.19%	Boeing Co	-1.98%
Fox Corp	1.55%	Interactive Brokers Group Inc	-1.71%
Paramount Skydance Corp	2.11%	Humana Inc	-1.94%
Warner Bros Discovery Inc	2.11%	Eaton Corporation PLC	-1.90%

STOXX 600

thyssenkrupp AG	3.61%	Nemetschek SE	-4.42%
Saab AB	3.45%	Sonova Holding AG	-3.56%
Commerzbank AG	3.37%	Reply SpA	-3.55%
Rheinmetall AG	2.71%	Partners Group Holding AG	-3.19%
Asseco Poland SA	2.70%	Brunello Cucinelli SpA	-2.90%
KGHM Polska Miedz SA	2.65%	Ferrari NV	-2.87%
ABN Amro Bank NV	2.57%	Thule Group AB	-2.57%
Standard Chartered PLC	2.57%	Givaudan SA	-2.41%
Elia Group SA	2.56%	Adidas AG	-2.32%
Bayer AG	2.38%	Amplifon SpA	-2.32%
Boliden AB	2.38%	CTS Eventim AG & Co KgaA	-2.32%
Orlen SA	2.38%	DSM-Firmenich AG	-1.86%
Raiffeisen Bank International AG	2.38%	Straumann Holding AG	-2.32%
Repsol SA	1.69%	Symrise AG	-1.97%
Italgas SpA	2.33%	Rational AG	-2.19%
Leonardo SpA	2.33%	Evonik Industries AG	-1.60%
Bilfinger Se	2.01%	P/F Bakkafrøst	-1.60%
Telecom Italia SpA	2.09%	Sectra AB	-2.10%
Novo Nordisk A/S	2.24%	Sika AG	-1.97%
Mercedes-Benz Group AG	1.59%	SIG Group AG	-1.97%

Industries breakdown



Neutrality Hedges

For the European leg of the portfolio, the 10-year historical beta using monthly observations has been 0.06488. For the US leg, the raw monthly beta has been -0.04057. There is no clear linear relationship between the portfolio and the S&P 500 and STOXX EURO 600. As a result, we do not need any beta hedging to stay neutral to the underlying movements in the indices. It is important to note that beta exposures are subject to change over time, and maintaining a consistently beta-neutral position would require dynamic hedging with futures. However, this approach is not applied since portfolio allocations are rebalanced on a semiannual basis. Another risk faced by the fund is currency risk. At inception, the fund's currency risk is zero. However, currency risk may evolve over time. Currently, no currency hedging measures have been implemented.



DISCLAIMER - This is an academic paper related to an academic project that does not represent any investment recommendation nor offer any solicitation to buy or sell securities or to adopt an investment strategy. The opinions expressed are subject to change. References to specific securities, asset classes and financial markets are for illustrative purposes only and are not to be intended to and should not be interpreted as recommendations. Reliance upon information in this material is at the sole risk and discretion of the reader.

Further Research – Adding a Climate Hedge Factor

Based on De Nard, Engle & Kelly (2024), “Factor-Mimicking Portfolios for Climate Risk”, Financial Analysts Journal, Vol. 80, No. 3

Paper Overview

Motivation

Climate change is a structural risk, meaning that no counterparty is willing to guarantee an insurance that pays out in case of a systemic climate event. Therefore, investors seeking protection must hedge themselves through portfolio construction. Drawing on Merton (1987) and Pastor, Stambaugh & Taylor (2021), the paper treats climate risk as a priced factor. A portfolio short this factor — appreciating when adverse climate news arrives and reprices exposed assets — functions as a hedge. The authors also distinguish **physical risk** (storms, wildfires, droughts) from **transition risk** (regulatory shifts toward a low-carbon economy). While negatively correlated at short horizons, both carry positive risk premia over time.

Climate News Indices

Climate risk is proxied by text-based weekly indices from the NYT and WSJ. The **NYT Cosine Similarity** index computes TF-IDF cosine distance between each daily edition and a curated climate dictionary, emphasising rare terminology. The **NYT Tag Index** measures the fraction of articles tagged 'global warming' or 'climate change' by LexisNexis NLP. Three WSJ indices capture general **climate risk**, **physical hazards**, and **transition risk** separately via ML keyword methods. All climate news indices trend upward over the 2001–2021 sample, spiking around the Paris Agreement and Copenhagen Conference, and remain positively correlated with each other.

Portfolio Construction & Results

The investment universe is 177 climate-focused funds from NYU Stern's V-LAB (fossil fuel-free, low-carbon, low environmental risk, and sustainable-sector ETFs and mutual funds). For each fund, the authors run a rolling 12-month OLS regression with the climate news index and five investable control factors — MARKET (VTI), HML, SMB, Stranded Asset (SPY – 0.7·KOL – 0.3·XLE), and Oil (WTI crude) — as regressors, yielding a climate β_i orthogonal to conventional risk premia. A signal vector is formed by weighting positive β_i 's and zeroing negative- β funds (as the investment universe is already “green”, short positions are eliminated). Portfolio weights are then assigned by solving $\min \mathbf{w}'\Sigma\mathbf{w} - \lambda\mathbf{w}'\mathbf{b}$, s.t. $w_i \geq 0$, $\Sigma w_i = 1$, with $\lambda \approx 0.002$. The covariance matrix Σ is estimated via the **AFM-CVC shrinkage estimator** of De Nard (2022), which outperforms sample covariance and GARCH-based alternatives since the number of funds is relatively much larger than the length of the return time series. Out-of-sample performance is evaluated via **Adaptive Least Squares** with HC3 standard errors.

Metric	CEP (Optimal)	Base Case	Equal-Weight
Alpha (ann.)	7.34% **	2.46%	-2.32%
Beta (news)	3.30 **	1.71 *	-0.40
Info Ratio	0.69	0.51	0.27
Avg Return	14.71%	9.78%	4.88%

NYT Tag Index, June 2001–January 2021. ** $p < 0.01$, * $p < 0.05$. Results hold for WSJ indices, alternative signals, and after 10 bps transaction costs (alpha falls to 6.20%, still significant at 1%). Avg Return is in annualised terms.

The CEP comprehensively dominates. The base case (simple signal-weighting with an identity covariance) delivers roughly one-third of the alpha and half the beta. The equal-weighted portfolio (without the climate signal) produces negative alpha and negative beta despite drawing entirely from a green universe, confirming that naive diversification across sustainable funds is not a substitute for proper signal construction. These results are achieved **after controlling for the same traditional factors** (MARKET, HML, SMB, Stranded Asset, Oil) used in construction, demonstrating that the climate signal is genuinely orthogonal to conventional equity risk premia, and not covered by the existing factors.

Relevance to the Fund

The paper's evidence that the CEP earns significant alpha even after controlling for MARKET, HML, SMB, Stranded Asset, and Oil (factors that mostly overlap with the fund's existing investment strategy) directly supports that a climate signal would deliver **genuinely additive return contribution**, rather than merely recycling existing exposures.

The main structural gap is the investment universe: the paper uses a dedicated green ETF pool, whereas the fund selects from the S&P 500 and STOXX Europe 600. This could be handled in two ways: (i) instead of using the investment universe of the paper, we consider the stocks already used by the fund and estimate the climate signal β , then the final portfolio weight using these. It should be noted that this way we will have to allow for short positions as well, since our investment universe is no longer exclusively “green”. The second approach could be extending the investment universe (ii), meaning the addition of the “green” assets to our portfolio. This approach would require estimating both the existing factor exposures and the new climate factor consistently across the extended set of assets.

From an operational perspective, the first approach is simpler to implement, however, it might yield a more limited impact, as the original stock universe might be less sensitive to the climate factor.



DISCLAIMER - This is an academic paper related to an academic project that does not represent any investment recommendation nor offer any solicitation to buy or sell securities or to adopt an investment strategy. The opinions expressed are subject to change. References to specific securities, asset classes and financial markets are for illustrative purposes only and are not to be intended to and should not be interpreted as recommendations. Reliance upon information in this material is at the sole risk and discretion of the reader.

Further Research – Adding a Climate Hedge Factor

Based on De Nard, Engle & Kelly (2024), “Factor-Mimicking Portfolios for Climate Risk”, Financial Analysts Journal, Vol. 80, No. 3

Implementation Roadmap

- Climate index: NYT Article Search API (weekly relative frequency of climate-tagged articles) as primary; Google Trends as a cost-free robustness check. The fact that the signals used by the original paper are all positively correlated suggests that resorting to only one or two climate news series won't lead to a significant loss of information.
- Universe: In case of choosing option (ii) mentioned before, V-LAB's 177 funds cross-referenced against Yahoo Finance availability; drop any with < 12 months of daily return history. In case of option (i), no changes to the existing universe (S&P 500 and STOXX Europe 600) are required.
- Signal & optimization: rolling 12-month OLS for beta; solving $\min \mathbf{w}'\Sigma\mathbf{w} - \lambda\mathbf{w}'\mathbf{b}$ with AFM-CVC covariance and $\lambda = 0.002$ to determine the final weights.
- Validation: 1-year quasi-OOS testing; regress CEP returns on climate index via ALS; compare alpha and beta against equal-weight and base-case benchmarks before final allocation.

Potential Caveats

- Turnover: the paper's optimal CEP averages 95% monthly turnover. Taking transaction costs at 10 bps into account, alpha falls to 6.20% (still significant), however, it is important to investigate how the less frequent rebalancing of the fund impacts the quasi-OOS performance of the additional climate factor.
- Proxy quality: an NYT API or Google Trends index will be noisier than the paper's proprietary TF-IDF pipeline, likely compressing estimated betas and reducing signal strength. However, this approach is preferred when taking implementation feasibility into consideration.
- Post-2021 regime: the paper's sample ends in January 2021. The 2022 energy crisis, the US Inflation Reduction Act, and the ECB's evolving green framework represent a structurally different environment, also from a regulatory standpoint. Therefore, the climate factor might perform differently from the expectations set by the paper, making a rigorous quasi-OOS testing even more necessary.



DISCLAIMER - This is an academic paper related to an academic project that does not represent any investment recommendation nor offer any solicitation to buy or sell securities or to adopt an investment strategy. The opinions expressed are subject to change. References to specific securities, asset classes and financial markets are for illustrative purposes only and are not to be intended to and should not be interpreted as recommendations. Reliance upon information in this material is at the sole risk and discretion of the reader.

© Minerva Investment Management 2026. All rights are reserved

Overview

Motivation

Traditional asset pricing models, such as the Fama–French five-factor model, explain stock returns using a limited number of linear factors. While these models are widely used, they may not fully capture more complex relationships between firm characteristics. Recent developments in machine learning suggest that financial data may contain nonlinear patterns that are not detected by standard models. This research investigates whether machine learning methods can help identify such patterns and whether these findings can improve traditional factor models.

Data and Research Setup

This analysis is based on a large dataset of U.S. listed companies constructed from CRSP and Compustat. The final dataset includes more than 1.2 million monthly observations and over 12,000 firms, covering the period from 2000 to 2024. Only common stocks listed on major exchanges (NYSE, AMEX, NASDAQ) are included to ensure consistency and reliability.

Firm-level characteristics are constructed using both accounting and market data. These include size, book-to-market ratio, profitability, leverage, and investment, as well as market-based indicators like momentum, volatility, and turnover. All variables are winsorized at the 1% level and standardized (z-scored) each month to ensure comparability.

The target variable in the analysis is the next-month excess return. This allows the models to use current firm characteristics to predict future performance without look-ahead bias.


Models and Methodology

The analysis applies both linear and nonlinear machine learning models to study the relationship between firm characteristics and stock returns. First, a LASSO regression model is estimated as a benchmark. LASSO is a linear model that selects only the most relevant predictors. However, the results show that all coefficients are shrunk to zero, and the model achieves almost no predictive power ($R^2 \approx 0$). This suggests that simple linear relationships are not sufficient to explain stock returns.

Second, a non-linear model (XGBoost) is implemented. This model is designed to capture more complex interactions between variables. Despite its flexibility, the predictive performance remains very low, with an out-of-sample R^2 of approximately -0.0006 . This confirms that predicting individual stock returns is extremely difficult.

Instead of focusing only on prediction, the model is used for interpretation. Using SHAP values, the most important variables are identified. The results show that leverage, earnings yield, and book-to-market are the most influential factors in explaining return variation.

Model Performance (XGBoost)	
Test R^2 :	-0.0006
Test MSE:	0.0576



SHAP values measure the average contribution of each variable to the model's predictions. Higher values indicate greater importance in explaining return variation.

From Machine Learning to Factor Construction

Instead of using the model predictions directly, the most important variables identified by SHAP are transformed into factor portfolios. For each variable, stocks are sorted into deciles, and a long–short portfolio is constructed by going long the highest decile and short the lowest.

This leads to the creation of several machine-learning-inspired factors:

- Leverage-based factor
- Earnings yield factor
- Book-to-market factor
- Interaction factor that combines leverage and earnings yield

These factors are constructed using value-weighted returns and represent economically meaningful patterns identified by the model.



DISCLAIMER - This is an academic paper related to an academic project that does not represent any investment recommendation nor offer any solicitation to buy or sell securities or to adopt an investment strategy. The opinions expressed are subject to change. References to specific securities, asset classes and financial markets are for illustrative purposes only and are not to be intended to and should not be interpreted as recommendations. Reliance upon information in this material is at the sole risk and discretion of the reader.

Further Research – Machine Learning and Factor Construction in Asset Pricing

Asset Pricing Results

The main objective is to evaluate whether these ML-based factors improve traditional asset pricing models.

A comparison is done between the standard Fama–French 5-factor model (FF5) and an augmented version that includes the new ML factors.

The results show a **consistent but moderate improvement**. When only FF5 factors are used, the explanatory power is already relatively high, with average R^2 values around 0.88–0.91 depending on the portfolio set. However, after adding the ML-based factors, the average R^2 increases to approximately 0.91–0.92.

In addition, **all 25 portfolios** in each test show an improvement in R^2 , meaning that the ML factors consistently add explanatory power. In some cases, the increase in R^2 is particularly strong, exceeding 0.10 for individual portfolios.

However, the reduction in pricing errors (alphas) is smaller and not consistent across all portfolios. Some portfolios show improvement, while others do not.

PRICING SUMMARY: FF5 vs FF5 + ML				
Asset Set	Mean R^2 (FF5)	Mean R^2 (FF5 + ML)	Mean ΔR^2	Improved Portfolios Count
Size $\times$ BM	0,8881	0,9132	0,0251	25,0000
Size $\times$ Leverage	0,9074	0,9221	0,0147	25,0000
Size $\times$ EarnYield	0,8837	0,9030	0,0193	25,0000

R^2 measures the proportion of return variation explained by the model. ΔR^2 represents the increase in explanatory power after adding machine-learning-based factors.

TOP PORTFOLIOS WITH BIGGEST R2 GAIN								
Asset Set	Portfolio (Size $\times$ Characteristic)	Alpha (FF5)	t-stat (FF5)	R^2 (FF5)	Alpha (FF5 + ML)	t-stat (FF5 + ML)	R^2 (FF5 + ML)	ΔR^2
Size $\times$ BM	Small–High	0,0031	1,3685	0,8069	0,0024	1,7457	0,9078	0,1009
Size $\times$ BM	Small–High	0,0009	0,4033	0,8128	-0,0002	-0,1804	0,9027	0,0899
Size $\times$ BM	Large–High	0,0006	0,3522	0,7715	-0,0012	-1,1246	0,8550	0,0835
Size $\times$ BM	Medium–High	0,0005	0,2509	0,8358	-0,0002	-0,1171	0,9051	0,0693
Size $\times$ Leverage	Small–High	-0,0033	-1,4369	0,8322	-0,0023	-1,3653	0,8957	0,0635
Size $\times$ BM	Large–High	0,0002	0,1303	0,8100	-0,0005	-0,3772	0,8728	0,0628
Size $\times$ EarnYield	Small–Low	-0,0009	-0,4376	0,8659	-0,0009	-0,5699	0,9158	0,0499
Size $\times$ EarnYield	Large–High	0,0010	0,5837	0,8415	-0,0004	-0,3921	0,8888	0,0473
Size $\times$ EarnYield	Large–Low	0,0007	0,4445	0,8893	0,0023	1,6224	0,9262	0,0369
Size $\times$ EarnYield	Large–Low	-0,0008	-0,3968	0,7798	0,0006	0,3031	0,8165	0,0367
Size $\times$ EarnYield	Small–Low	0,0017	1,2693	0,9103	0,0022	2,0601	0,9463	0,0360
Size $\times$ Leverage	Small–High	-0,0011	-0,7811	0,8904	-0,0005	-0,5001	0,9255	0,0351
Size $\times$ EarnYield	Large–High	0,0013	1,2559	0,8593	0,0002	0,2524	0,8938	0,0346
Size $\times$ EarnYield	Small–High	0,0025	1,6433	0,8725	0,0018	1,5410	0,9065	0,0340
Size $\times$ EarnYield	Medium–High	0,0008	0,5114	0,8711	0,0002	0,1732	0,9037	0,0326

Portfolios are formed by sorting firms by size and a given characteristic. “Small–High” refers to firms in the smallest size group with the highest values of the characteristic. ΔR^2 denotes the increase in explanatory power.

Interpretation and Conclusion

The results suggest that machine learning models are not effective for directly predicting stock returns. However, they are valuable as a tool for identifying underlying structure in financial data.

By identifying the most relevant variables and transforming them into factor portfolios, it is possible to improve the explanatory power of traditional asset pricing models. In particular, variables such as leverage and earnings yield appear to capture important sources of return variation that are not fully explained by standard factors.

Overall, these findings show that machine learning can complement, rather than replace, classical financial models. Instead of focusing on prediction alone, its main contribution lies in enhancing the structure and interpretation of factor-based investment strategies.

Limitations and Further Research

This study has several limitations. First, the machine learning models are applied to a limited set of firm characteristics and a specific sample period. Different variables or alternative datasets may lead to different results. Second, the analysis focuses on monthly data and does not consider higher-frequency information, which may contain additional signals. Third, the improvement in model performance is relatively modest, suggesting that the economic significance of the results should be interpreted with caution.

Future research could explore the use of alternative data sources, such as textual or ESG-related variables, as well as more advanced model specifications. Another direction would be to test whether machine-learning-based factors perform consistently across different markets or time periods.



DISCLAIMER - This is an academic paper related to an academic project that does not represent any investment recommendation nor offer any solicitation to buy or sell securities or to adopt an investment strategy. The opinions expressed are subject to change. References to specific securities, asset classes and financial markets are for illustrative purposes only and are not to be intended to and should not be interpreted as recommendations. Reliance upon information in this material is at the sole risk and discretion of the reader.

© Minerva Investment Management 2026. All rights are reserved

Introduction

The main objective of this section is to assess and quantify the risk embedded in the Minerva IMS Long-Short Equity Fund built by the portfolio team. We use a daily perspective on the potential extreme behaviour of a basket of assets selected by the portfolio analysts. The analysis will include three VaR and ES models (two parametric and one non-parametric) and an overview of how sentiment analysis can be considered a factor for short term investments.

As the Quantitative Research division, our focus is the estimation of the two main risk indicators:

- The daily Value at Risk (VaR): the maximum portfolio loss that occurs with $\alpha\%$ of probability over a time horizon of 1 day. For instance, if the VaR ($\alpha=5\%$) = -3.00%, it means that tomorrow there is a 5% probability of encountering a loss in the interval [-100%, -3.00%] potentially;
- The daily Expected Shortfall (ES): the expected return on the portfolio in the worst $\alpha\%$ of cases. So, it is just a mean of the returns lower than the VaR.

A simple technique to estimate these two measure is based on a historical approach: given a time series of returns of a financial security, we can easily compute the desired quantile of the historical distribution to estimate the VaR, and, after that, estimate the ES just by averaging the values below this threshold.

However, this naive approach is not well suited for our purpose: in fact, by considering our portfolio as a single financial asset, we are losing all the information that comes from all the components; moreover, with this approach we are simply focusing on the past behaviour of the fund, while our main goal is to retrieve a risk metric for the future possible trends.

In order to overcome these issues, we propose two alternative techniques that provide better risk estimates:

- Parametric approach (simple approach and time-series modelling approach),
- Bootstrapping

The first method is very well suited for understanding the main vulnerabilities in the portfolio composition, while with the second one it is possible to observe how the metrics varied in the past quarters.

For both pieces of analysis, we used daily market prices of portfolio constituents for the past 151 trading days (covering mid-August 2025 to mid-March 2026). All the analysis has been conducted with Python.


Head of Quantitative Research

Ianina Tarca: ianina.tarca@studbocconi.it

Quantitative Analysts

Lorenzo Giambra: lorenzo.giambra@studbocconi.it

Emilio De Col: emilio.decol@master.unibocconi.it



Parametric Approach

We propose to analyse VaR and ES separately for each asset included in the portfolio and then, to estimate the VaR and ES for the whole fund by taking into account the correlation between portfolio constituents.

Parametric approach is based on the assumption that returns of a financial security follow some theoretical distribution. Thus, VaR and ES can be expressed as an α -percentile of the distribution. The crucial step to accurately estimate VaR and ES is to select the appropriate distribution of returns and estimate its parameters.

It is possible to state that stock returns do not follow Gaussian distribution due to the presence of "fat tails": unexpected events might have a huge impact on the stock prices, so it is possible to observe extreme values more frequently than a Normal distribution would predict. For this reason, we assume that stock returns follow a Student-t distribution, thus, the parameters to be estimated are the mean μ , volatility σ and number of degrees of freedom ν .

To obtain more valid and robust results, we proceed with two alternative parameter estimation approaches – (a) simple approach, and (b) time-series modelling approach. For all parts of analysis, we use the last 151 return observations, which correspond to 7-month window.

Parametric Approach

Time-series modelling approach

Under the simple approach, we estimate the above-mentioned parameters in the following way:

1. We assume that the mean historical daily return of each security is a good estimate for the expected future return. Thus, μ is estimated as a simple average of daily returns.
2. Volatility of returns σ is calculated as a simple standard deviation of returns.
3. Number of degrees of freedom ν is selected in a way that it best approximates the empirical distribution of returns. In order to do that, we used the Kolmogorov-Smirnov statistic that, for a given empirical cumulative distribution function F and a proposal F_n , is:

$$D_n = \sup_x |(F_n - F)|$$

Ideally it should be equal to 0 for a perfect fit, so our goal is to minimize it by proposing different ν for Student-t distribution.

Time-series modelling approach

Because the volatility of returns is not constant over time, it is often modelled by conditional heteroscedasticity processes. The most common way to model volatility is through a Generalized Autoregressive Conditional Heteroscedasticity model GARCH(p,q), where the forecast of the next-period volatility depends on the previous p shocks to stock returns (derived from some mean model) and previous q forecasts of volatility:

$$\sigma_{t+1|t}^2 = \omega + \sum_{i=1}^p \alpha_i \epsilon_{t-i}^2 + \sum_{j=1}^q \beta_j \sigma_{t-j+1|t-j}^2$$

The advantage of GARCH model is that it allows to better estimate the current forecast of return volatility by putting more weight on more recent information. Thus, in the periods of market turbulence GARCH model will produce higher volatility forecasts than the simple average of squared deviations from the mean (see the graph at the bottom).

Because the portfolio is composed exclusively of equity instruments traded on liquid markets, we can assume that prices are efficient, and thus returns can be described by a constant mean model for GARCH(p,q) process, which implies that current mean estimates do not depend on previous returns or shocks. GARCH(p,q) then is estimated by Maximum Likelihood (MLE), which optimizes the distribution parameters. We subsequently use MLE estimates of distribution to derive VaR and ES.

Value-at-Risk

Once the parameters of stock returns are known, it is possible to calculate VaR. We estimate the VaR for 95% and 99% confidence level by applying the following formula:

$$VaR_\alpha = \sigma * T_\nu^{-1}(\alpha) + \mu$$

where σ is the estimated volatility of a security, $T_\nu^{-1}(\alpha)$ is the α -percentile of a Student-t distribution with ν degrees of freedom, and μ is the expected return of a stock

Expected Shortfall

Expected shortfall is defined as a conditional expectation of loss, given that the loss occurred. If we introduce the assumption of a continuous distribution of returns of a security, then parametric expected shortfall is simply defined as a tail conditional expectation, and thus can in general be defined by the following formula for any security X :

$$ES_\alpha(X) = -\frac{1}{\alpha} \int_0^\alpha VaR_\gamma(X) d\gamma$$

Under the assumption of Student-t distribution with ν degrees of freedom it can be proven that the expected shortfall would be given as:

$$ES_\alpha(X) = \sigma * \frac{\nu + (T_\nu^{-1}(\alpha))^2}{\nu - 1} \frac{\tau_\nu(T_\nu^{-1}(\alpha))}{\alpha} + \mu$$

where σ is the estimated volatility of a security, $T_\nu^{-1}(\alpha)$ is the α -percentile of a Student-t distribution with ν degrees of freedom, $\tau_\nu(\cdot)$ is the probability density function of Student-t distribution with ν degrees of freedom and μ is the expected return of a stock.

We estimate the ES for 95% and 99% confidence level.

Portfolio VaR and ES

Considering the correlation between the stocks, we estimate the VaR and ES of the whole portfolio for 95% and 99% confidence level by applying the following formulas:

$$VaR_{\alpha,ptf} \approx \sqrt{VaR_\alpha * \rho * VaR_\alpha'}$$

$$ES_{\alpha,ptf} \approx \sqrt{ES_\alpha * \rho * ES_\alpha'}$$

where VaR_α and ES_α are column vectors of individual stock VaR and ES, respectively and ρ is the correlation matrix between securities

The approximation arises because of the assumption of Student-t distribution of returns – the formulas above become an equality the closer the distribution of returns is to the Gaussian.

	Simple approach	GARCH
VaR_{95%}	-1.59%	-1.66%
VaR_{99%}	-2.61%	-2.98%
ES_{95%}	-2.25%	-2.55%
ES_{99%}	-3.45%	-4.34%

Long-Short Equity Fund – Risk Metrics Analysis

TOP & BOTTOM 5 stocks (simple approach)

	VaR 95	VaR 99	ES 95	ES 99		VaR 95	VaR 99	ES 95	ES 99
Italgas SpA	-2.09%	-3.03%	-2.67%	-3.50%	Norwegian Cruise Line Holdings Ltd	-4.64%	-9.10%	-7.74%	-14.12%
Elia Group SA	-2.15%	-3.13%	-2.75%	-3.62%	Paramount Skydance Corp	-4.97%	-9.38%	-8.03%	-14.34%
Symrise AG	-2.62%	-3.73%	-3.30%	-4.29%	Ciena Corp	-5.58%	-10.44%	-8.80%	-14.87%
Warner Bros Discovery Inc	-2.03%	-3.38%	-2.89%	-4.38%	Lumentum Holdings Inc	-8.73%	-13.84%	-11.92%	-16.99%
Kenvue Inc	-2.28%	-3.76%	-3.23%	-4.86%	Super Micro Computer Inc	-6.78%	-13.25%	-11.28%	-20.53%

Bootstrapping

When estimating a certain metric, one of the main problems in Statistics is the lack of the whole population data and the consequent use of only a sample. In our case the population data is the complete historical price data of the securities that are part of our portfolio, in which we only have the data of recent years.

Bootstrapping is a statistical technique that by having only a sample of the population data, provides estimates of statistical metrics that are closer to the ones obtained from the population data.

Given a sample of size n , implementing bootstrap is very simple:

- Sample with replacement n times from the original sample (note that one observation could be selected more than once);
- Compute the metric of interest (in our case the VaR or ES) on this newly created sample and save it;
- Repeat the previous steps M times with $M \rightarrow +\infty$ (we have selected $M=300$ for instance);
- Average and compute the standard error of the metrics estimated in each step.

With this method, by estimating the expected shortfall and the standard errors, we can retrieve a more insightful view of our portfolio, but in this case, we are losing the risk contribution of each stock that we had in the previous case.

	Estimate	Standard error
VaR_{95%}	-1.48%	0.33%
VaR_{99%}	-2.73%	0.43%
ES_{95%}	-2.24%	0.36%
ES_{99%}	-3.05%	0.41%

TOP & BOTTOM 5 stocks (GARCH)

	VaR 95 (GARCH)	VaR 99 (GARCH)	ES 95 (GARCH)	ES 99 (GARCH)
Warner Bros Discovery Inc	-0.57%	-1.19%	-1.02%	-2.04%
Elia Group SA	-1.93%	-3.09%	-2.67%	-3.89%
Italgas SpA	-1.67%	-2.93%	-2.49%	-3.98%
Kenvue Inc	-1.77%	-3.28%	-2.97%	-4.80%
P/F Bakkafrost	-2.73%	-4.14%	-3.61%	-5.00%

	VaR 95 (GARCH)	VaR 99 (GARCH)	ES 95 (GARCH)	ES 99 (GARCH)
Gartner Inc	-3.95%	-9.19%	-8.10%	-18.10%
Super Micro Computer Inc	-4.82%	-11.35%	-9.95%	-22.22%
Expedia Group Inc	-5.76%	-12.73%	-11.05%	-23.30%
Intel Corp	-8.65%	-17.42%	-14.88%	-28.09%
Norwegian Cruise Line Holdings Ltd	-7.11%	-16.04%	-13.87%	-29.43%



DISCLAIMER - This is an academic paper related to an academic project that does not represent any investment recommendation nor offer any solicitation to buy or sell securities or to adopt an investment strategy. The opinions expressed are subject to change. References to specific securities, asset classes and financial markets are for illustrative purposes only and are not to be intended to and should not be interpreted as recommendations. Reliance upon information in this material is at the sole risk and discretion of the reader.