Investment Institute

WORKING PAPER 186 | **MARCH 2026**

Integrating Geopolitical Risk Into Low Volatility Factor Construction

Amundi
Investment Solutions

Trust must be earned

Integrating Geopolitical Risk Into Low Volatility Factor Construction

Wassim KALLALI
Amundi Investment Institute
wassim.kallali@amundi.com

Binh PHUNG-QUE
Amundi Asset Management
binh.phung-que@amundi.com

Anna ROSENBERG
Amundi Investment Institute
anna.rosenberg@amundi.com

Sonja TILLY
Amundi Investment Institute
sonja.tilly@amundi.com

Takaya SEKINE
Amundi Investment Institute
takaya.sekine@amundi.com

Abstract

Financial markets respond not only to quantitative fundamentals but also to narratives that shape investor perceptions of risk, particularly during periods of heightened geopolitical uncertainty. While the low volatility factor has demonstrated robust risk adjusted performance across markets, its construction is typically static and does not account for narrative-driven fluctuations in risk. This paper investigates whether incorporating narrative-based measures of geopolitical risk can enhance the design and resilience of low volatility equity strategies. Using firm-level geopolitical sentiment indicators derived from news flows and expert reports, we develop state-dependent specifications of the low volatility factor in which portfolio weights dynamically adjust to changes in the prevailing geopolitical narrative environment. We evaluate the performance and defensive characteristics of these narrative-conditioned strategies relative to conventional low volatility benchmarks across multiple market regimes. Our findings provide evidence that geopolitical narratives contain incremental information relevant for factor-based portfolio construction, improving downside protection during periods of elevated uncertainty without materially compromising long-term performance. These results contribute to the growing literature on text-based measures in factor investing and offer practical insights for the design of more adaptive, risk-aware systematic investment strategies.

Keywords: Geopolitical risk, low-volatility factor, factor investing, narrative-based investing.

JEL classification: G11, G12, C58.

Acknowledgement

The authors are very grateful to Frederic Lepetit, Jiali Xu and Ru Qi Wu for their helpful comments. The opinions expressed in this research are those of the authors and are not meant to represent the opinions or official positions of Amundi Asset Management.

About the authors

Wassim KALLALI

Wassim Kallali is a recent graduate with a Master's in Statistics and Actuarial Science from ENSAE Paris and a Master's in Quantitative Finance. At Amundi, he worked on the development of a quantitative low-volatility equity strategy integrating geopolitical sentiment, contributing to portfolio construction, optimization, and backtesting. His academic and professional background sits at the intersection of statistics, finance, and quantitative research. Wassim is a recipient of the Tunisian government's bourse d'excellence, reflecting his strong academic achievements and potential.

Binh PHUNG-QUE

Binh Phung-Que is responsible for Financial Engineering for the Quant Equity and Convexity Solutions. Binh joined Amundi in 2022 as a Senior Quantitative Analyst within the Factor Investing portfolio management team, with a focus on multi-factor basket construction, associated risk models, and portfolio construction. His expertise has since late 2025 expanded to include volatility and convex solutions.

Previously, he worked between 2017 and 2022 as a Specialist Quantitative consultant for Amundi Equities Factor Investing team, Amundi Multi Asset Teams as well as for other asset managers, with a focus on factor investing and portfolio construction. Before, he was a Quantitative Sales Specialist at FactSet Research System for 13 years.

Binh studied for Master's degree in Management and Information Technology from Institut National des Telecommunication.

Anna ROSENBERG

Anna Rosenberg is the Head of Geopolitics for Amundi Investment Institute. In this role, Anna helps Amundi's investment teams and the company's clients make better investment decisions by seeing through ongoing geopolitical shifts.

Anna joined Amundi in 2022 from Signum Global Advisors, where she was co-founder and Senior Partner responsible for the firm's predictive research output focused on Europe. Anna helped grow the business from a start-up into a rapidly growing and well-respected company.

Before Signum, Anna helped executives at multinational companies adjust their expansion plans across emerging markets to political and economic developments. At Frontier Strategy Group (now FrontierView), Anna also launched and headed the Sub-Saharan Africa research practice and co-created the firm's macroeconomic forecasting product, for which her department won various awards.

Throughout her career, Anna has produced 'out-of-the-box' thinking and research helping firms perform better in volatile environments.

Her commentary on politics and work on business strategy has been featured in Harvard Business Review, Bloomberg, the Wall Street Journal, the BBC and CNBC, among others.

Anna has been nominated a Responsible Leader of the Herbert Quandt Foundation in 2023 and a Bloomberg Future Leader in 2024. She is also a Visiting Scholar at the London Institute of Banking and Finance and a BBC Expert Woman.

Anna is a historian by training and holds an MA with Distinction from the School of Oriental and African Studies (SOAS), University of London. She speaks fluent German, Spanish and Portuguese, reads French and Italian and has studied Arabic and Latin.

Sonja TILLY

Sonja Tilly, PhD, CFA joined Amundi as a Quantitative Analyst in 2024. Commencing her career in 2008 in London, Sonja started out as an Investment Analyst, then transitioning to Quantitative Analyst roles. Her time at Stanhope Capital, Aberdeen Asset Management, and Hiscox honed her skills in investment analysis, financial asset modelling, economic scenario development, and stress-testing portfolios.

While working as Quantitative Researcher at Qoniam Asset Management, Sonja created a filtering method based on deep learning for processing extensive news text, extracting signals that were transformed into a systematic equity trading strategy. Sonja's experience working in traditional Finance is complemented by her insights into the crypto space, gained as Quantitative Researcher at decentralised finance start-up Allora Network, where she led the creation of loan terms for fully automated NFT-backed on-chain loans.

Sonja holds a PhD in Computer Science from University College London. Her research focuses on the impact of news narrative on the economy and financial markets, blending methodologies from data science and econometrics. Further, she is a CFA Charterholder.

Takaya SEKINE

Takaya Sekine, CFA is the Deputy Head of Quant Portfolio Strategy within Amundi Investment Institute (formerly known as the Quantitative Research Team of Amundi). In this role, he works on the practical implementation of quant research, artificial intelligence and alternative data for investment strategies.

He joined Amundi in 2000 and is in his current position since July 2018. Prior to that, he was Deputy CIO at Amundi Japan (between 2011 and 2018) with a focus on global quantitative strategies, Head of Index and Multi-Strategies at Amundi Japan (between 2010 and 2011), Fund Manager (between 2007 and 2010) and Financial Engineer (between 2001 and 2007). He has been involved in macro and policy related investment strategies for both retail and institutional clients. Takaya began his career as an IT Manager at Amundi Japan's predecessor company (between 2000 and 2001).

Takaya is a CFA charterholder since 2005. He received the Ingénieur Civil des Mines degree from Ecole des Mines de Nancy in 2000.

1 Introduction

Financial markets remember stories at least as much as they remember numbers (Shiller, 2017). Episodes of acute geopolitical tension, shifts in global trade relations, and macroeconomic uncertainty illustrate that investor sentiment is shaped by narratives propagating through news, speeches, and social media, alongside traditional quantitative indicators. This narrative-driven environment poses fundamental questions about the construction and resilience of systematic investment strategies, particularly those designed to offer defensive characteristics during periods of market stress.

The low volatility factor, which tilts portfolios toward assets with relatively low return volatility, has attracted substantial attention for its ability to generate attractive risk-adjusted performance, a pattern that appears difficult to reconcile with classical asset pricing theory (Blitz *et al.*, 2019; Clarke *et al.*, 2011). Yet the empirical success of this factor depends critically on modelling choices regarding universe definition, volatility estimation, rebalancing frequency, and portfolio constraints. These choices become particularly consequential when volatility itself is strongly influenced by geopolitical events and the narratives surrounding them.

Despite the documented impact of geopolitical risk on market dynamics, the systematic integration of geopolitical sentiment into factor-based portfolio construction remains underexplored. While narrative-based measures of risk have become increasingly available through advances in natural language processing and textual analysis (Caldara & Iacoviello, 2022; Cunado *et al.*, 2020), their incorporation into defensive equity strategies has received limited attention in both academic and practitioner communities.

This paper addresses this gap by investigating how narrative-based indicators of local political and geopolitical risk can be systematically integrated into low volatility portfolio construction. We develop a framework that uses firm-level local political and geopolitical sentiment scores, extracted from news flows and expert reports, as conditioning variables to design alternative specifications of the low volatility factor. Our approach examines constructions in which portfolio weights explicitly respond to changes in perceived geopolitical risk, creating state-dependent versions of the factor that adapt to the prevailing narrative environment.

Our analysis contributes to the literature connecting factor investing with text-based measures of risk by explicitly combining a traditional defensive factor with narrative information in a bottom-up equity setting. We address whether geopolitical narrative contains incremental information for systematic strategies aimed at protecting portfolios during adverse periods. The empirical evidence presented is relevant for understanding both the role of narratives in factor investing and the design of more resilient risk-managed investment strategies.

This paper is structured as follows. Section 2 reviews related works. Section 3 introduces the dataset used and explains signal construction. Section 4 sets out the methodology used to integrate the geopolitical signals into the low volatility portfolio. Section 5 presents and discusses the overall results. Finally, section 6 concludes the paper.

2 Related Works

Geopolitical risk has emerged as a significant driver of market behavior, especially during periods of heightened global tension, conflict, or policy uncertainty. Recent literature has increasingly focused on how geopolitical risk influences both macroeconomic aggregates and asset-level outcomes. This review synthesizes key studies in the field, distinguishing between research focused on macroeconomic and market-wide impacts of geopolitical risk, and studies

that analyze firm- or stock-level sensitivities. This distinction is particularly important for portfolio construction, especially in low-volatility strategies and multifactor portfolio frameworks, which rely on the reliable prediction and mitigation of downside risk through careful factor selection and asset exposure. In addition to mapping the empirical findings on geopolitical risk, the review will highlight key contributions relevant to the integration of geopolitical risk sentiment into systematic portfolio construction, with focus on low-volatility and risk-managed investment approaches.

2.1 Geopolitical Risk at the Macroeconomic Level

Caldara and Iacoviello (2017) construct a news-based Geopolitical Risk Index (GPR) by systematically scanning newspaper archives for terms related to geopolitical threats and events. Their analysis spans over a century and shows that GPR spikes around historically significant events (mainly conflicts) such as the World Wars, the Cuban Missile Crisis, and 9/11. Importantly, they demonstrate that increases in GPR are associated with significant declines in investment and employment at the macroeconomic level. Moreover, industries with higher exposure to geopolitical uncertainty exhibit more pronounced drops in investment, laying the groundwork for subsequent stock-level research.

Building on this, Cunado *et al.* (2020) explore the impact of geopolitical risk on stock returns and volatility in emerging markets for the period of February 1974 to August 2017 using a panel GARCH approach. They find that global geopolitical risk – rather than country-specific risk – is the dominant driver of volatility across these markets. While returns appear less sensitive to geopolitical shocks, volatility is significantly elevated during periods of heightened geopolitical tension. This suggests that low-volatility portfolios, particularly those with emerging market exposure, need to account for global rather than localized geopolitical shocks.

Similarly, Plakandaras *et al.* (2018) study the interaction between GPR and volatility in oil and stock markets using dynamic conditional correlation (DCC) models. The data spans the period from January 1985 to May 2018. Their findings reveal that geopolitical events typically lead to a rise in oil prices, presumably due to supply concerns, and a simultaneous decline in equity returns. Moreover, volatility spillovers from oil to equities increase during high-risk periods, indicating that geopolitical shocks can transmit across asset classes. These results have important implications for portfolio diversification strategies, particularly the use of commodities as volatility hedges.

Smales (2021) examines the effect of geopolitical risk on investor sentiment in the United States, applying time-varying parameter vector autoregressive (TVP-VAR) models. He finds that geopolitical shocks reduce investor sentiment in the short to medium term, while sentiment itself does not appear to Granger-cause geopolitical events. These dynamics emphasize sentiment as a transmission mechanism through which geopolitical event influences asset prices. For investors employing systematic allocation strategies, such insights underscore the potential value of incorporating sentiment indicators as early warning signals.

At a cross-country level, Francis and Chia (2023) conducts a systematic literature review that catalogs empirical studies applying the Caldara and Iacoviello's GPR index to various markets and asset classes (Caldara & Iacoviello, 2022). The review identifies a consistent pattern: while returns across countries may respond differently, volatility effects tend to be more homogeneous and persistent. This reinforces the growing consensus that volatility, more than returns, is the primary transmission channel of geopolitical events into financial markets – an insight that is directly relevant for low-volatility portfolio construction.

Lastly, Lamine and Zribi (2024) applies local projection methods to measure the effect of global geopolitical shocks on national stock indices across 29 economies. The results

show significant heterogeneity: countries such as Latvia, China, and members of the Euro Area exhibit substantial and persistent declines in equity prices following geopolitical shocks, while others display muted responses. This country-level variation suggests that geographic allocation decisions in global portfolios can significantly alter the portfolio's exposure to geopolitical risk.

2.2 Geopolitical Risk at a Stock Level

While macro-level studies highlight the broader market implications of geopolitical risk, firm-level research provides a more granular view, offering insights critical for bottom-up portfolio construction.

Pástor and Veronesi (2013) examine political uncertainty and stock valuation, finding that political shocks lead to increased risk premia and lower returns. Their model provides a theoretical foundation for interpreting the empirical patterns observed in geopolitical risk studies.

Another example is Lee *et al.* (2023), who investigate the impact of geopolitical risk stemming from North Korea on South Korean firms. They construct a monthly geopolitical risk index using South Korean media data and find that elevated geopolitical risk leads to significantly lower stock returns. The effect is particularly strong for large firms, those with a high share of domestic investors, and those with substantial fixed assets. This heterogeneity underscores the importance of firm characteristics in determining exposure to geopolitical shocks.

In a similar vein, Brignone *et al.* (2025) extends the Caldara and Iacoviello framework (Caldara & Iacoviello, 2022) to examine how large geopolitical risk shocks disproportionately affect firms of different sizes, showing that larger firms or those more embedded in international trade suffer more in aggregate outcomes during periods of high geopolitical risk.

Gkillas *et al.* (2018) analyze tail risk and return distributions in the context of geopolitical uncertainty, applying extreme value theory to sectoral and firm-level data. They show that during geopolitical risk spikes, downside tail risks become significantly fatter, implying a higher probability of extreme losses. This insight is crucial for low-volatility strategies, which may underperform if they are exposed to hidden tail risks not captured by standard volatility metrics.

Though more conceptual in nature, Shiller (2020) works on narrative economics provides a framework for understanding how geopolitical events shape market behavior through collective stories and investor psychology. Shiller argues that narratives around war, terrorism, and geopolitical tension can drive asset prices independently of fundamentals. This suggests a potential role for media sentiment and narrative tracking in improving the responsiveness of quantitative investment models to geopolitical developments.

2.3 Factor Investing and Risk Modeling

Factor investing represents a foundational framework in asset pricing and portfolio construction, tracing its roots to academic efforts to explain cross-sectional differences in asset returns. Cochrane (1999) lay the groundwork for the modern asset pricing theory, emphasizing the role of systematic factors beyond the Capital Asset Pricing Model (CAPM). Earlier empirical contributions such as Fama and French (1996) demonstrate that portfolios sorted on size, value, and momentum exhibit persistent return premia unexplained by single-factor models. MacKinlay (1995) further review the econometric challenges in asset pricing tests, underscoring the importance of robust factor identification and model specification.

Building on these foundations, the practical application of factor models evolved in institutional contexts. The Barra risk model, developed by Rosenberg (B. Rosenberg, 1974) and later popularized in practice (Bender *et al.*, 2013), introduce a structured way to model exposure to factors and market-related variables. These models use sound economic and financial relationships to project existing information into ex-ante risk forecasting. Barra is investing a lot through years to be ex ante. They worked on the scalings of the risk prediction.

Despite their usefulness, structural ex ante models which have low frequency of version changes and often struggle to adapt to paradigm shifts Abou Rjaily *et al.* (2024). This limitation has driven a shift toward data-enhanced and adaptive factor investing. For example, Ghayur *et al.* (2018) illustrate how quantitative investors increasingly blend traditional factor signals with macro, sentiment, and behavioral data to capture risk premia more dynamically. Similarly, recent works by Bettin *et al.* (2023) and Li (2025) show how machine learning and high-frequency sentiment data can improve signal timing and robustness in multi-factor portfolios.

A growing literature has also challenged the robustness of static factor models in volatile or uncertain environments. DeMiguel *et al.* (2024) show that factor performance is highly regime-dependent and propose robust optimization techniques that adapt to structural shifts. This complements the perspective of Bender *et al.* (2013), who advocate for flexibility in factor model design to reflect evolving market conditions.

2.4 Low Volatility Investing and Emerging Extensions

Low volatility investing has attracted significant attention as a challenge to classical finance theory, particularly the Capital Asset Pricing Model (CAPM), which posits a positive linear relationship between risk and return. Contrary to this, numerous empirical studies have shown that low-volatility stocks often outperform their high-volatility counterparts on a risk-adjusted basis, giving rise to what is now commonly referred to as the low volatility anomaly.

Baker *et al.* (2011) are among the early contributors to this anomaly's formalization, providing robust evidence across international markets that portfolios tilted toward low-beta or low-volatility stocks tend to deliver superior risk-adjusted returns. Their work helped establish low volatility as a legitimate factor in the broader factor investing framework.

Subsequent studies refine and expand this perspective. Clarke *et al.* (2011) examine the performance of minimum variance portfolios, demonstrating that the anomaly persists even after accounting for other known risk factors. Van Vliet and Blitz (2011) offer further validation by showing that low-volatility strategies maintain their performance across different market regimes, while Blitz (2016) and Blitz *et al.* (2019) explore implementation aspects and found that low-volatility portfolios can be constructed in a cost-effective manner, even in the presence of transaction costs and liquidity constraints.

Other scholars investigate the behavioral and structural explanations for the anomaly. Blau and Whitby (2017), for instance, propose that investor overconfidence and preference for lottery-like payoffs may drive persistent demand for high-volatility stocks, leaving low-volatility stocks systematically undervalued. Liu *et al.* (2019) extend the behavioral argument by connecting the anomaly to institutional constraints that force market participants to seek high-volatility assets for potential outperformance.

More recently, Zaremba (2016) analyze the performance of low-volatility strategies in emerging markets, revealing that the anomaly is not confined to developed economies and may in fact be stronger where markets are less efficient. Traut (2023) add nuance by exploring how macroeconomic regimes influence the effectiveness of low-volatility strategies,

Table 1: Economic narratives classification.

Narrative	Themes
Back to the 70s	Inflation, taxation, employment, government policy, central bank intervention, commodity prices, exchange rate, macroeconomic risk.
Environment	Natural disaster, environmental law, green finance, green growth, health, green innovation, natural resource management, protest, resilience, adaptation, mitigation.
Geopolitical Risk	Aid groups, attacks, country at risk, crime, crisis, ideology, justice, politics, sanctions, tensions, trade barriers, war, weapons.
Monetary	Easing, quantitative restriction, monetary policy, central banks, interest rates, financial markets.
Roaring 20s	Innovation, productivity, growth, technology, savings, inequality.
Secular Stagnation	Growth, inflation, productivity, demographics, macroeconomic risk, savings.
Social	Discrimination, extreme parties, inequality, living together, supply chain, terror groups, unrest.

Source: Amundi Investment Institute

highlighting potential cyclicality in their performance. Yang *et al.* (2025) is among the first to investigate the interaction between geopolitical risk and volatility tilts, suggesting that geopolitical uncertainty may amplify or dampen the anomaly depending on the nature and region of the events.

Following our identification of recent paradigm shift for the integration of geopolitical risk in the risk premier of assets, we choose to extend, we extend the factor investing framework by integrating ex ante geopolitical risk signals into our low-volatility factor-based strategy. In particular, we incorporate RavenPack's (RP) news-based sentiment analytics, enabling us to dynamically adjust exposure to stocks based on forward-looking measures of geopolitical risk. This innovation helps address the limitations of traditional risk models in capturing latent, event-driven risks that are increasingly relevant in a multipolar geopolitical context.

3 Dataset Construction

The dataset used in this study is derived from the RP analytics platform[1], a leading provider of structured news and event data. RP transforms unstructured content such as news articles and social media into structured, real-time data feeds, enabling systematic analysis of financial, economic, and geopolitical developments.
To construct our dataset, we first identify relevant stocks that are included in the *low volatility portfolio* at each point in time.

Blanqué *et al.* (2022) and A. Rosenberg *et al.* (2024) developed narratives for geopolitical risk. These narratives are summarized in Table 1 provide a structured taxonomy of the macro-financial environment and capture shifts in the informational content of different risk factors. We select the RP events reflecting these geopolitical risk dimensions (for a list of all RP events, see appendix A.1).

[1] https://ravenpack.com

From within the RP data, we select only events associated with an `EVENT_RISK` score $\neq 0$ [2]. The `EVENT_RISK` is defined as a continuous variable ranging from 0 to 1 and reflects the potential for an event to negatively affect an entity, either reputationally or financially. Values closer to 1 indicate a higher likelihood of negative consequences for the entity, such as financial losses due to industrial accidents, regulatory actions, or hostile takeovers. Conversely, values near 0 suggest little to no risk, or insufficient information to assess the risk. We use the *sentiment strength score*, which incorporates a time-decay function. This allows us to emphasize recent developments while still accounting for the cumulative effect of past events. The two-stage process involves first calculating a daily average sentiment score and then applying a decay-weighted aggregation across a lookback window. The equations used to compute sentiment strength are given below:

$$Average_{1D} = \sum_{i=1}^{N} \frac{ESC_i}{n}, \quad \text{for } i \in U \tag{1}$$

where ESC stands for event sentiment score, $U = \{1, \ldots, n\}$ is the number of events for the company over the past 24 hours.

$$\text{Strength}_{LB} = \sum_{i=0}^{LB-1} \frac{w_{1,i} \cdot w_{2,i} \cdot \text{average1D}_i}{\sum_{j=0}^{LB-1} \left(\begin{cases} w_{2,j}, & \text{if average1D}_j \neq \text{NULL} \\ 0, & \text{otherwise} \end{cases} \right)} \tag{2}$$

Here, LB denotes the lookback window, and w_1 and w_2 are two exponential decay functions with different decay speeds. Each serves a distinct purpose:

- $w_1(d_i)$, where $d_i = 0, 1, \ldots, LB - 1$, represents a *degradation function* that reduces the influence of events as they become older. It applies a standard exponential decay based on the number of days since the current date.
- $w_2(d_i)$, also defined for $d_i = 0, 1, \ldots, LB - 1$, produces a set of weights that assign greater importance to the most recent news. This function accelerates the decay of older events and ensures that the resulting sentiment strength score remains bounded within the range $[-1, 1]$.

To facilitate comparability across entities and over time, we compute the *z-scores* of the sentiment strength scores. This standardisation is necessary because not all entities appear in the news with the same frequency, and on certain days, some entities may not feature at all due to the absence of relevant events. As a result, the raw strength scores are unevenly distributed across the dataset. By converting them to z-scores, we normalize the scores relative to the cross-sectional distribution on each day, allowing for a more interpretable *ranking* of entities based on their relative exposure to geopolitical event sentiment.
By combining *event filtering* (based on relevance and risk) with *decay-weighted sentiment scoring*, the dataset provides a structured and interpretable foundation for downstream analysis of geopolitical risk exposure across a dynamic portfolio on a security level.

[2] RavenPack User Guide

4 Methodology

4.1 Baseline Low Volatility Factor Portfolio Construction

The baseline low volatility factor is constructed through a systematic, rules-based framework designed to identify and weight low-risk stocks within a benchmark equity universe. The methodology combines historical volatility-based estimation, a non-parametric normalization procedure inspired by distributional mapping, and a benchmark-aware portfolio construction process consistent with systematic portfolio management practices.

Historical Volatility Estimation. For each stock i in the investment universe, the annualized historical volatility is estimated over a rolling window of 252 trading days, approximately corresponding to one calendar year of daily observations:

$$\sigma_i = \sqrt{252 \times \frac{1}{251} \sum_{t=1}^{252} (r_{it} - \bar{r}_i)^2} \tag{3}$$

Where r_{it} denotes the daily return of stock i at time t, and μ_i represents its average daily return over the same period. This rolling volatility captures recent dynamics in market risk while maintaining a stable estimation horizon, a standard approach in systematic equity strategies.

Normalization via Distributional Transformation. Instead of using conventional z-scores based on the mean and standard deviation of volatilities, we apply a rank-based transformation that maps empirical volatilities to standard normal scores truncated between -3 and 3. Let p_i denote the empirical percentile rank of the volatility stock i within the universe at a given date, computed as:

$$p_i = 100 \times \frac{\text{rank}(\sigma_i)}{N} \tag{4}$$

Where $\text{rank}(\sigma_i)$ assigns rank 1 to the highest volatility and N is the total number of stocks. The corresponding normalized score is obtained as:

$$Z_{vol,i} = \begin{cases} -3 & \text{if } p_i = 0, \\ 3 & \text{if } p_i = 100, \\ \Phi^{-1}\left(\dfrac{p_i}{100}\right) & \text{otherwise,} \end{cases} \tag{5}$$

where $\Phi^{-1}(\cdot)$ denotes the inverse of the standard normal cumulative distribution function. Because the ranking assigns higher percentiles to lower-volatility stocks, higher $Z_{vol,i}$ values correspond to lower estimated volatility. Hence, stocks with the highest $Z_{vol,i}$ scores are those exhibiting the lowest historical volatility.

This transformation, equivalent to a monotonic mapping of volatility ranks into the standard normal domain, mitigates the influence of outliers and avoids strong distributional assumptions. It also ensures stability and interpretability of scores across time-key properties for systematic strategies benchmarked against large equity indices. From a practical standpoint, this approach provides portfolio managers with standardized, bounded risk scores that remain consistent and easy to monitor.

Portfolio Selection and Buffering. At each monthly rebalancing date, stocks are ranked according to their normalized volatility scores $Z_{vol,i}$. The low-volatility portfolio consists of the top 20% of stocks with the highest $Z_{vol,i}$ values, corresponding to the least volatile segment of the market. To enhance stability and reduce turnover, a buffer rule is introduced: stocks that were included in the previous month's portfolio and now fall within the 20 - 30% range of the ranking are retained. This mechanism maintains continuity in holdings and limits unnecessary transaction costs.

Portfolio Weighting. The final portfolio is constructed by re-scaling the benchmark market-capitalization weights of the selected stocks. Let Q_t denote the set of eligible securities at time t after applying both the selection and buffer criteria. The portfolio weights are defined as:

$$w_i(t) = \begin{cases} \dfrac{w_i^{bench}(t)}{\sum_{j \in Q_t} w_j^{bench}(t)} & \text{if } i \in Q_t, \\ 0 & \text{otherwise,} \end{cases} \tag{6}$$

where $w_i^{bench}(t)$ represents the benchmark market-cap weight of stock i at time t. This benchmark-relative weighting ensures that the portfolio remains investable, liquid, and easily interpretable from a portfolio management perspective, while expressing a clear and systematic tilt toward lower volatility.

Relation to the Minimum-Variance Optimization Problem. Although this construction relies on ranking and heuristic weighting, it can be viewed as a practical proxy of the theoretical minimum-variance portfolio optimization problem. The minimum-variance portfolio is formally defined as:

$$\begin{aligned} \min_{w} \quad & w^\top \Sigma w \\ \text{s.t} \quad & \sum_{i=1}^{N} w_i = 1, \\ & w_i \geq 0, \quad \forall i \in \{1, \ldots, N\}, \end{aligned} \tag{7}$$

where Σ denotes the covariance matrix of stock returns. In practice, direct optimization can be unstable due to estimation error in Σ, especially in large universes. The proposed factor construction provides a transparent alternative by relying solely on individual volatilities rather than full covariance estimation. It captures the essential economic intuition of the minimum-variance strategy allocating more weight to lower-risk stocks while maintaining consistency with the operational constraints and benchmarking requirements faced by portfolio managers.

4.2 Factor Reconstruction via Convex Optimization

This section introduces a unified optimization framework that incorporates both signal types across selection and allocation stages.

4.2.1 Blended-Score Integration Framework

Modern factor investing recognizes that portfolio construction requires coordinated decisions across multiple stages: universe definition, asset screening, and weight optimization

(Bender *et al.*, 2013). While traditional approaches treat these stages independently, recent advances demonstrate that simultaneous consideration of multiple signals can enhance performance when signals are combined appropriately (Ghayur *et al.*, 2018). We extend this literature by implementing a **signal blending framework** that distinguishes between selection-stage and optimization-stage composite scores, recognizing that factor relevance varies across construction phases.

Let $Z_{vol,i}$ denote the standardized volatility z-score (Section 4) and $Z_{geo,i}$ the standardized geopolitical z-score (Section 3), truncated to $[-3, 3]$. We construct two composite scores.

Selection composite score. For asset screening, we form:

$$S_i^{select} = \gamma Z_{geo,i} + (1-\gamma) Z_{vol,i}, \quad \gamma \in [0,1] \tag{8}$$

Optimization composite score. For weight determination, we form:

$$S_i^{opt} = \beta Z_{geo,i} + (1-\beta) Z_{vol,i}, \quad \beta \in [0,1] \tag{9}$$

The parameters γ and β are calibrated independently via grid search, selecting values that maximize risk-adjusted performance metrics (Sharpe ratio) over a training period from 31/03/2015 to 15/12/2020. This independent calibration reflects the empirical observation that factors exhibit differential efficacy across portfolio construction stages.

Economic rationale for dual-score framework. The separation of selection and optimization scores recognizes two distinct roles for our signals. In the **selection phase**, geopolitical risk scores identify tail-risk exposures that may not be captured by historical volatility. Geopolitical events-wars, sanctions, political crises-generate discontinuous shocks that manifest as sudden price adjustments rather than gradual volatility changes (Caldara & Iacoviello, 2022). Excluding assets with high geopolitical sensitivity at the screening stage prevents these tail risks from entering the portfolio, regardless of their historical volatility profile. Thus, we cannot ignore the importance of geopolitical scores in asset selection: they capture event-driven risks orthogonal to time-series volatility measures.

In the **optimization phase**, volatility scores provide essential information for diversification and risk control. Even within a pre-screened universe, assets exhibit heterogeneous volatility that directly impacts portfolio variance. Following the minimum-variance portfolio literature (Clarke *et al.*, 2011), optimal weight allocation should account for cross-sectional volatility differences to minimize total portfolio risk. Our inclusion of $Z_{vol,i}$ in S_i^{opt} implements this principle: assets with lower volatility receive higher optimization scores, tilting weights toward less volatile securities. This is analogous to the minimum-variance weighting scheme $w_i \propto 1/\sigma_i^2$ (Clarke *et al.*, 2011), adapted here to a composite-score framework. Thus, we cannot ignore volatility scores in the optimization phase: they are fundamental to constructing a low-volatility portfolio.

The independent calibration of γ and β allows the framework to adapt to varying market regimes. During periods of elevated geopolitical tension, optimal calibration may assign high weight to geopolitical signals in selection and weighting. Conversely, in calm periods, selection may emphasize volatility while optimization remains balanced. We do not impose *ex ante* restrictions on the values of γ or β; rather, we let the data determine the optimal signal combination through out-of-sample validation.

Asset selection with buffer. Assets are ranked by S_i^{select} and the top quintile is retained for portfolio construction. To maintain consistency with the baseline methodology and ensure approximately the same number of assets across rebalancing dates, we implement a 10-percentile buffer: securities ranked between the 20th and 30th percentiles that were included in the portfolio at time $t-1$ remain eligible at time t.

4.2.2 Penalized Quadratic Optimization

For the selected universe $\mathcal{U}_t$ containing n assets at time t, we determine optimal weights by solving:

$$\begin{aligned} \max_{w \in \mathbb{R}^n} \quad & \lambda_1 s^\top w - \lambda_2 \|w - w^{bench}\|_2^2 \\ \text{s.t} \quad & \mathbf{1}^\top w = 1 \\ & w_{min} \leq w_i \leq w_{max}, \quad \forall i \end{aligned} \tag{10}$$

where $w^{bench} = (w_1^{bench}, \ldots, w_n^{bench})^\top$ denotes the normalized market-capitalization weights of the benchmark constituents within the selected universe $\mathcal{U}_t$ and $s = (S_1^{opt}, \ldots, S_n^{opt})^\top$ contains the optimization composite scores for all selected assets.

This concave maximization problem is equivalent to the convex minimization formulation by negating the objective, ensuring tractability via standard quadratic programming methods.

Formulation rationale and comparison to existing approaches. Problem (10) departs from classical portfolio optimization in three key aspects, each with strong practical motivation:

(i) Characteristic-based objective. Rather than maximizing expected return $\mu^\top w$ as in mean-variance optimization (Markowitz, 1952), we maximize a composite factor score $s^\top w$ where $s = (S_i^{opt})_i$. This aligns with the characteristic-based asset pricing paradigm (Fama & French, 1996), which posits that returns are driven by observable characteristics. By optimizing directly on these characteristics (in our case volatility and geopolitical risk), we avoid the substantial estimation error inherent in $\hat{\mu}$ (DeMiguel *et al.*, 2009).

(ii) Regularization via benchmark penalty. The quadratic penalty $\lambda_2 \|w - w^{bench}\|_2^2$ implements a regularization (Hastie *et al.*, 2009), shrinking weights toward the benchmark to mitigate overfitting. This is statistically motivated: our composite scores S_i^{opt} contain measurement error, and unconstrained optimization would amplify this noise into extreme positions.

(iii) Covariance-free formulation. Unlike min variance optimization or tracking-error optimization $\max_w \alpha^\top w$ subject to $\sqrt{(w - w^B)^\top \Sigma (w - w^B)} \leq \sigma_{TE}$, our approach avoids estimating the covariance matrix Σ. This eliminates sensitivity to second-moment estimation error, which is particularly severe for large universes ($n \approx 150$ in our application). The penalty $\|w - w^{bench}\|_2^2$ measures deviation in weight space rather than return space, providing regularization without requiring Σ.

Economic interpretation of the penalty term. The benchmark penalty serves multiple economic functions beyond statistical regularization:

- *Implicit turnover control:* Since w^{bench} evolves smoothly with market capitalizations, solutions to (10) exhibit temporal stability. At rebalancing date t, the optimal weights $w^*(t)$ are pulled toward $w^{bench}(t)$, which itself resembles $w^{bench}(t-1)$ scaled by returns. This creates implicit turnover dampening without requiring explicit constraints $\sum_i |w_i - w_i^{prev}| \leq \tau$ that introduce non-differentiability. The advantage is twofold:

(a) computational tractability: problem (10) remains a smooth quadratic program solvable in polynomial time, and (b) economic efficiency: the penalty allows turnover when justified by signal strength (large changes in S_i^{opt}) while suppressing noise-driven turnover.

- *Capacity and implementability:* The benchmark anchor ensures weights remain reasonably close to market weights, facilitating implementation for strategies managing significant assets under management. Extreme deviations from benchmark weights ($w_i \gg w_i^{bench}$ or $w_i \ll w_i^{bench}$) create liquidity constraints and market impact costs; the penalty naturally limits such deviations.

- *Diversification:* By shrinking toward the value-weighted benchmark, the penalty inherits the diversification properties of market-cap weighting while allowing factor tilts. This provides a middle ground between equal-weighting (which ignores firm size) and pure signal-based ranking (which may concentrate in small-cap stocks).

Penalty parameter calibration. We set $\lambda_1 = 1$ (numeraire) and $\lambda_2 = 20$ based on empirical observations from extensive backtesting across multiple market regimes and rebalancing periods. The choice $\lambda_2 = 20$ reflects a balance between signal capture and regularization: values $\lambda_2 < 10$ occasionally yield portfolios with excessive signal-driven concentration, while $\lambda_2 > 30$ produce portfolios nearly identical to the benchmark, indicating over-regularization. The selected value ensures that the optimized portfolio consistently achieves $\sum_i S_i^{opt} w_i^* > \sum_i S_i^{opt} w_i^{bench}$ (strict score improvement over the benchmark) across our backtesting training sample, while maintaining reasonable diversification through the explicit box constraints $w_{min} \leq w_i \leq w_{max}$.

For a discussion of theoretical properties, see appendix A.2.

4.2.3 Computational Implementation

Problem (10) is a strictly convex quadratic program with n variables, n box constraints, and one equality constraint. We implement the optimization using the `cvxpy` framework (Diamond & Boyd, 2016), which automatically reformulates the problem in canonical conic form and dispatches to appropriate solvers. The solver cascade attempts ECOS, OSQP, SCS, and CLARABEL sequentially until convergence, with all solvers configured to tolerances $\epsilon_{abs} = \epsilon_{rel} = 10^{-4}$ and maximum iterations $N_{max} = 50{,}000$. Convergence typically occurs within 10-50 iterations for interior-point methods and 100-500 iterations for first-order methods. Detailed solver specifications, complexity analysis, and convergence guarantees appear in Appendix C.

If all solvers fail to converge with the penalty term, the algorithm falls back to solving the linear program:

$$\max_{w \in \mathcal{F}} \; s^\top w \tag{11}$$

This fallback problem always admits a solution because the feasible set $\mathcal{F}$ is non-empty and compact (Theorem 1), and the objective $s^\top w$ is continuous. The optimal solution places maximum feasible weight (w_{max}) on the highest-scored assets until the budget constraint binds. Specifically, sorting assets by S_i^{opt} in descending order and greedily assigning $w_i = w_{max}$ to top-ranked assets until $\sum_i w_i = 1$ is achieved (with the marginal asset receiving residual weight) solves (11) exactly. This greedy allocation always exists and provides a higher bound on the score improvement achievable with the full penalized objective (10), ensuring robustness against numerical failures.

4.3 Calibration of Blended-Signal Parameters

The blended-score framework introduced in Section 4.2.1 is governed by two parameters, γ and β, which determine the relative contribution of geopolitical and volatility information in the selection and optimization stages. These parameters control the trade-off between an exposure to event-driven geopolitical risk and the preservation of the low-volatility profile of the portfolio. Miscalibration of (γ, β) can therefore lead either to an overly cautious portfolio that neglects valuable information carried by geopolitical scores, or to an ostensibly low-volatility strategy that inadvertently loads on assets exposed to severe geopolitical tail risk. A careful calibration procedure is thus essential to ensure that the resulting factor portfolio delivers the intended economic exposures in a stable and robust way.

4.3.1 Static Calibration on a Training Window

A natural benchmark is to treat (γ, β) as fixed over time and to calibrate them on a historical training window. Concretely, we specify a discrete grid $\Gamma \subset [0,1] \times [0,1]$ of candidate pairs (γ, β). For each pair, we reconstruct the portfolio over a training period and compute standard performance statistics such as cumulative return and Sharpe ratio.

We opted for a *static* calibration that selects

$$(\gamma^\star, \beta^\star) \in \arg\max_{(\gamma,\beta)\in\Gamma} \left\{ \mathrm{SR}^{\mathrm{train}}(\gamma, \beta) \right\}, \tag{12}$$

The selected pair $(\gamma^\star, \beta^\star)$ is then held fixed when evaluating the strategy on the subsequent out-of-sample period.

4.3.2 Limitations of a Static Calibration

While intuitive, the static calibration assumes that the optimal trade-off between geopolitical and historical volatility information is time-invariant. This assumption is at odds with the highly non-stationary nature of geopolitical risk. Periods of relative geopolitical calm are interspersed with episodes of acute tension-wars, sanctions regimes, or political crises in which the cross-sectional dispersion of geopolitical exposure widens dramatically. In such regimes, geopolitical scores are more informative for future tail risk than in tranquil periods, where traditional volatility measures dominate.

A fixed pair $(\gamma^\star, \beta^\star)$ cannot adapt to these shifts in the informational content of the signals. Parameters calibrated in a low-tension regime will underweight geopolitical information precisely when it becomes most valuable, while parameters estimated in a crisis period may overstate the importance of geopolitical risk once conditions normalise. From a statistical perspective, static calibration is vulnerable to regime-specific overfitting: the training window may be dominated by a narrative configuration that is not representative of the future out-of-sample period. Economically, this translates into a factor portfolio whose sensitivity to geopolitical shocks is misaligned with the prevailing macro-financial narrative.

These considerations motivate a *dynamic* calibration scheme in which (γ, β) are allowed to evolve over time in response to shifting macroeconomic narratives.

4.4 Narrative-Based Adaptive Calibration

To construct an adaptive calibration scheme, we leverage the narrative classification framework developed by Blanqué *et al.* (2022), who identify seven distinct economic narratives that characterize the dominant macroeconomic and geopolitical themes prevailing in financial markets at any given time. Based on this framework, we select RP events corresponding to these geopolitical risk dimensions.

Blanqué *et al.* (2022) construct time-varying narrative indices from the Global Database of Events, Language and Tone (GDELT), which provides structured global news event data at a fifteen minute frequency. Their framework calculates the narratives' Count Weighted Tone every friday with data aggregated for the day. In this paper, we go one step further by introducing an exposure pattern directly linked to the dominance of a given narrative [GN]. We have studied the stickiness of the top ranking for a GN. In a calculation similar to rating transitions, we have identified that if a GN is leading the battle of narratives normalized over three months, then it has an approximately seventy percent probability that in the following week it is either leading again or it has fallen to the second rank only. We treat the second rank as a buffer zone. For the purpose of reducing turnover we wait for two consecutive weeks of GN in the top position in the battle before we activate by half the assets connected to the GN. Then if on a third week the GN confirms its leading position, we will increase the exposure to the assets connected to GN to the full exposure that we wish to allocate. This exposure mechanism is displayed in figure 1 over our sample period.

Blanqué *et al.* (2022) construct time - varying narrative indices from the Global Database of Events, Language and Tone (GDELT), which provides structured global news event data at daily frequency. Their framework assigns each narrative a dominance score at each date and identifies the prevailing narrative regime. Figure 1 shows how the prevailing narrative changes over the sample period: at each point the plotted dominance score indicates which narrative (geopolitical tension, social unrest, monetary policy focus, or macroeconomic concerns) was most prevalent in news coverage. The chart exhibits clear regime shifts - for example, the geopolitical narrative reached its peak dominance in 2022, coinciding with the Russia - Ukraine invasion. To make the figure immediately interpretable, add a short subtitle that summarizes this key finding and improve axis/legend labels as described below.

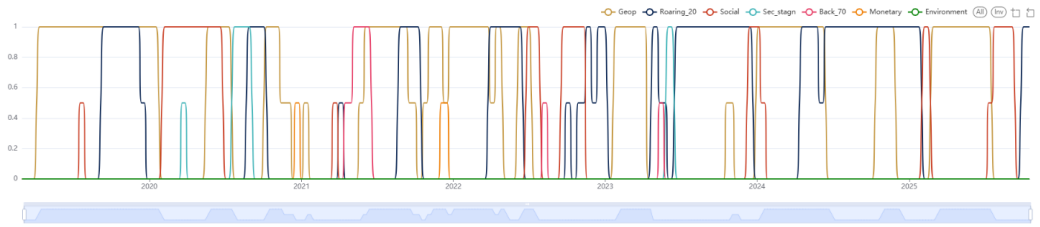

Figure 1: Evolution of the dominant narrative over time.
For instance, the geopolitical narrative peaked in 2022 (Russia - Ukraine invasion); narrative dominance shifts between geopolitical tension, social unrest, monetary policy and macroeconomic concerns.

Source: Amundi Investment Institute

We exploit this narrative classification to design an adaptive calibration of (γ_t, β_t) that responds to shifts in the macroeconomic regime. Rather than estimating continuous time-varying parameters, we adopt a parsimonious rule-based approach that maps narrative regimes to two discrete parameter configurations. Specifically, we define "high" values $\gamma^{\text{high}}, \beta^{\text{high}} \in]0,1[$ and "low" values $\gamma^{\text{low}}, \beta^{\text{low}} \in]0,1[$, and set

$$(\gamma_t, \beta_t) = \begin{cases} (\gamma^{\text{high}}, \beta^{\text{high}}), & \text{if the Geopolitical Risk narrative is dominant at time } t, \\ (\gamma^{\text{high}}, \beta^{\text{high}}), & \text{if the Social narrative is dominant at time } t, \\ (\gamma^{\text{low}}, \beta^{\text{low}}), & \text{otherwise.} \end{cases} \tag{13}$$

This specification places greater weight on geopolitical information when either Geopo-

litical Risk or Social narratives dominate. Including the Social narrative reflects that social issues, such as inequality, often evolve into geopolitical risks. Our geopolitical signal draws on a broad event taxonomy covering not only traditional geopolitical events (e.g., conflicts, sanctions, diplomatic tensions) but also societal disruptions, protests, civil unrest, and labor disputes. Many events classified as "Social" are captured by our RP-based filter and materially affect firm-level geopolitical exposure scores. As a result, periods of strong Social narrative dominance mark regimes in which geopolitical scores are especially informative about non-traditional political and societal risks to individual assets.

In all other narrative regimes including Back to the 70s, Monetary, Roaring 20s, Secular Stagnation, and Environment, we apply the lower parameter values $(\gamma^{\text{low}}, \beta^{\text{low}})$, which assign reduced weight to geopolitical information relative to historical volatility measures. This piecewise-constant adaptive rule yields a calibration that is transparent, economically interpretable, and directly tied to observable macro-financial narratives.

In the next section, we compare the empirical performance of the static and narrative-based adaptive calibrations, focusing on their impact on risk-adjusted returns, drawdown profile, and the stability of the low-volatility factor exposure.

5 Findings

We now evaluate the comprehensive optimization framework developed in section 4, which integrates geopolitical and volatility signals at both the asset selection and weight allocation stages through the penalized quadratic optimization introduced in subsection 4.2.2. This approach constructs composite scores for selection and optimization, with parameters (γ, β) calibrated to maximize portfolio performance.

5.1 Calibration Strategies

We implement three distinct calibration approaches to assess the value of adaptive parameter adjustment in response to macroeconomic narratives:

Static Calibration (Opt(Const.)): This benchmark strategy treats (γ, β) as fixed over time. Parameters are calibrated via grid search over a four-year training window to maximize Sharpe ratio, then held constant throughout the subsequent out-of-sample period. This approach assumes time-invariant signal informativeness.

Narrative-Based Adaptive Calibration (Adaptive-1): This specification implements the adaptive framework introduced in subsection 4.4, switching between high and low parameter regimes based on the dominant economic narrative identified by Blanqué *et al.* (2022). We set $\gamma^{\text{high}} = \beta^{\text{high}} = 0.8$ and $\gamma^{\text{low}} = \beta^{\text{low}} = 0.2$. During periods when the Geopolitical Risk or Social narratives dominate, the portfolio assigns substantial weight to geopolitical scores in both selection and optimization. In all other narrative regimes, geopolitical information receives reduced weight while volatility measures dominate. This symmetric parameterization treats selection and optimization stages identically.

Narrative-Based Adaptive Calibration (Adaptive-2): This specification recognizes that optimal signal weighting may differ between portfolio construction stages. We favor volatility information for asset screening while maintaining strong geopolitical emphasis during weight allocation: $\gamma^{\text{high}} = 0.4$, $\gamma^{\text{low}} = 0.1$, $\beta^{\text{high}} = 0.8$, and $\beta^{\text{low}} = 0.2$. This asymmetric calibration reflects the hypothesis that volatility-based selection provides a stable foundation of low-risk assets, while geopolitical adjustment at the optimization stage fine-tunes exposures to capture tail risk mitigation without compromising the low-volatility character.

5.2 Geopolitical Dataset Results

Performance Analysis

Figure 2 presents annual returns across calibration strategies. The adaptive portfolios demonstrate substantially superior performance, particularly during volatile years. Figure 3 displays cumulative performance trajectories, revealing striking differentiation. Adaptive-1 achieves the highest cumulative return, reaching approximately 147% by July 2025, while Adaptive-2 attains 125%. The baseline strategy reaches 95%, and the static calibration underperforms dramatically at 75%. This divergence accelerates during the 2024-2025 period, suggesting that adaptive calibration provides particular value during recent market conditions characterized by elevated geopolitical uncertainty.

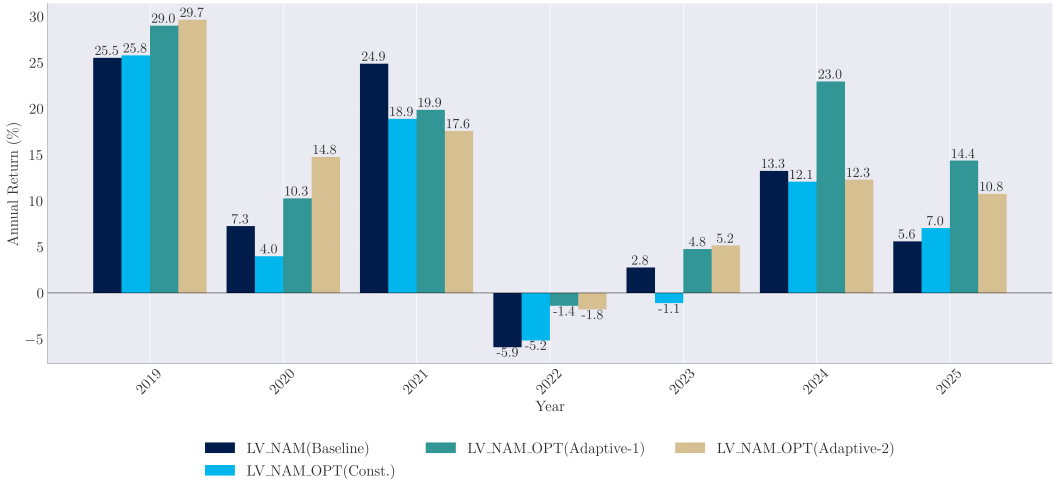

Figure 2: Annual returns, January 2019 to July 2025.
This chart exhibits the annual returns of optimized portfolios including geopolitical scores. The adaptive strategies outperform, especially in downmarkets.
Source: Amundi Investment Institute

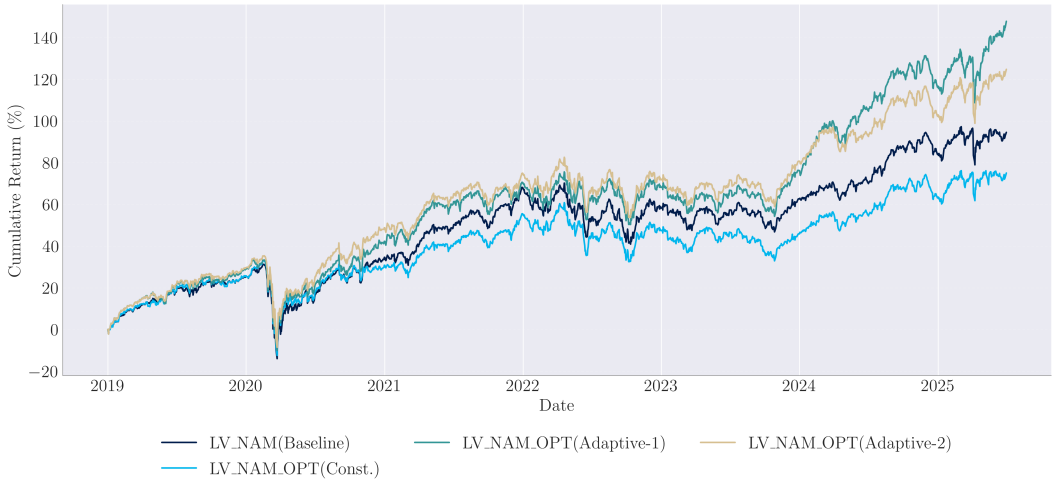

Figure 3: Cumulative performance, January 2019 to July 2025.
This chart displays the cumulative performance of optimized portfolios incorporating geopolitical scores. Adaptive-1 achieves the highest cumulative return (147%), comfortably beating the baseline strategy (95%).
Source: Amundi Investment Institute

Figure 4 illustrates annual volatility profiles. The adaptive portfolios exhibit moderately elevated volatility in certain years, particularly 2019 and 2021, reflecting their more dynamic positioning in response to narrative shifts. However, this increased volatility translates into substantially superior risk-adjusted returns, as evidenced by the Sharpe ratio improvements. Notably, Adaptive-2 maintains annual volatility levels remarkably close to the baseline across most years, confirming that its asymmetric calibration successfully preserves the low-volatility character while capturing geopolitical risk premia.

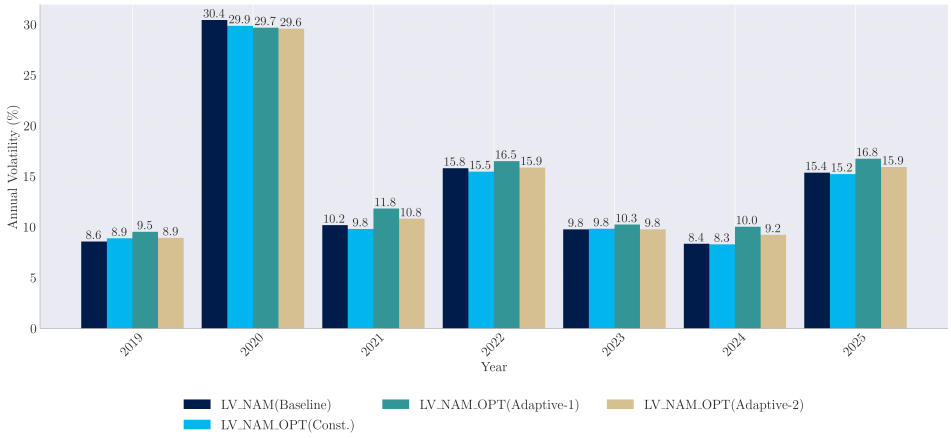

Figure 4: Annual volatility, January 2019 to July 2025.
This chart exhibits the annual volatility of optimized portfolios including geopolitical scores. The adaptive strategies exhibit slightly higher volatility compared to the baseline; yet delivering superior risk-adjusted returns.
Source: Amundi Investment Institute

Table 2 reports comprehensive risk and return metrics. Adaptive-1 achieves an annualized return of 14.45% with volatility of 16.31%, representing a 405 basis point improvement over the baseline's 10.40% return. Adaptive-2 delivers 12.80% annualized return with volatility of 15.84%, virtually identical to the baseline's 15.87%. This remarkable result demonstrates that Adaptive-2 captures 240 basis points of additional return without materially altering the portfolio's volatility profile, making it particularly attractive for low-volatility mandates with strict risk constraints.

Maximum drawdowns provide further evidence of risk management efficacy. The baseline experiences a drawdown of -34.45%, while Adaptive-1 and Adaptive-2 limit drawdowns to -32.56% and -32.45% respectively, representing meaningful tail risk mitigation. The static calibration underperforms the baseline delivering only 8.68% annualized return, illustrating that time-invariant parameters fail to capture regime-dependent signal informativeness. VaR and CVaR metrics corroborate these patterns: Adaptive-2 exhibits slightly lower tail risk measures than the baseline, confirming enhanced downside protection.

Table 2: Risk and return metrics for geopolitical-adjusted portfolios via convex optimization.

Portfolio	**Total Return** (%)	**Ann. Return** (%)	**Ann. Volatility** (%)	**Max DD** (%)	**VaR 95%** (%)	**CVaR 95%** (%)
Baseline	94.59	10.40	15.87	-34.45	-1.25	-2.33
Opt(Const.)	75.08	8.68	15.63	-34.13	-1.20	-2.27
Adaptive-1	147.86	14.45	16.31	-32.56	-1.30	-2.40
Adaptive-2	124.82	12.80	15.84	-32.45	-1.23	-2.31

Source: Amundi Investment Institute.

Table 3 presents risk-adjusted performance ratios that starkly illustrate the value of adaptive calibration. Adaptive-1 achieves a Sharpe ratio of 0.909, substantially exceeding the baseline's 0.704 and representing a 29% improvement in risk-adjusted returns. The Sortino ratio reaches 1.07, confirming exceptional downside risk management. Adaptive-2 delivers a Sharpe ratio of 0.840, a 19% improvement over baseline, while maintaining the low-volatility profile that characterizes the underlying factor strategy. The Calmar ratios further emphasize drawdown-adjusted performance superiority: Adaptive-2 achieves 0.395 compared to the baseline's 0.302, reflecting 31% better returns per unit of maximum drawdown.

The static calibration's dismal performance merits particular attention. With a Sharpe ratio of 0.611, it substantially underperforms the baseline despite incorporating geopolitical information. This failure illustrates that geopolitical signals exhibit highly time-varying informativeness: parameters optimized for one market regime prove detrimental when conditions shift. The static approach likely overweights geopolitical information during tranquil periods when traditional volatility measures dominate, or underweights it during crises when geopolitical exposure becomes critical, resulting in systematically poor timing.

Statistical Significance Testing

Table 4 reports Wilcoxon signed-rank test results (see appendix D.2). The adaptive strategies produce highly statistically significant outperformance. Adaptive-1 exhibits a mean daily return difference of 3.67% annualized with a p-value of 0.007, indicating significance at the 1% level. Adaptive-2 shows a mean difference of 2.14% annualized with a p-value of 0.020, significant at the 5% level. These results provide compelling statistical evidence that

Table 3: Risk-adjusted performance metrics for geopolitical-adjusted portfolios via convex optimization.

Portfolio	**Sharpe Ratio**	**Sortino Ratio**	**Calmar Ratio**
Baseline	0.704	0.795	0.302
Opt(Const.)	0.611	0.699	0.254
Adaptive-1	0.909	1.070	0.444
Adaptive-2	0.840	0.978	0.395

Source: Amundi Investment Institute

narrative-based adaptive calibration generates genuine performance improvements rather than spurious patterns arising from data mining.

In stark contrast, the static calibration exhibits a negative mean difference of -1.61% annualized with a p-value of 0.517, indicating no statistically significant relationship with the baseline. This non-result confirms that fixed parameters fail to systematically improve performance when signal informativeness varies across market regimes. The statistical tests corroborate the performance metrics: adaptive calibration delivers economically meaningful and statistically reliable improvements, while static approaches offer no systematic advantage despite incorporating additional information.

Table 4: Wilcoxon signed-rank test results for daily return differences relative to baseline (geopolitical optimization).

Portfolio	**Mean Diff** (% ann.)	**Test Statistic**	**p-value**
Opt(Const.)	-1.61	702655	0.517
Adaptive-1	3.67	751998	0.007 ***
Adaptive-2	2.14	744214	0.020 **

Source: Amundi Investment Institute

Note: *** indicates significance at 1% level; ** indicates significance at 5% level.

Economic Interpretation and Portfolio Manager Perspective

The empirical results establish Adaptive-2 as a particularly compelling alternative for low-volatility factor strategies subject to institutional constraints. Portfolio managers implementing low-volatility mandates face explicit volatility targets and must demonstrate consistent adherence to risk budgets. Adaptive-2's ability to deliver 240 basis points of additional annualized return while maintaining ex-post volatility nearly identical to the baseline addresses this practical challenge directly. The asymmetric calibration strategy, favoring volatility information for asset selection while incorporating geopolitical risk at the optimization stage, achieves a delicate balance: the portfolio retains its fundamental low-risk character through volatility-based screening, while geopolitical adjustment during weight allocation captures tail risk premia without destabilizing the volatility profile.

The performance during high-volatility periods provides particularly compelling evidence for adaptive strategies. During the COVID-19 crisis, the adaptive portfolios experienced comparable short-term volatility to the baseline but recovered more rapidly, as evidenced by superior returns in 2020 (Figure 2). During the 2022 stress period, when geopolitical tensions

intensified alongside monetary policy tightening, Adaptive-1 and Adaptive-2 demonstrated substantially better performance than the baseline, limiting losses while the static calibration faltered. The 2024-2025 period, characterized by elevated but not extreme volatility, showcases the adaptive strategies' ability to capture upside during recovery phases while maintaining risk discipline.

From a risk management perspective, the maximum drawdown improvements merit emphasis. Reducing peak losses from -34.45% to -32.45% may appear modest in percentage terms, but represents meaningful capital preservation. For institutional investors managing substantial assets, this 200 basis point reduction in maximum loss translates to significant absolute dollar amounts.

5.3 Preparing for Agentic AI

Recent literature examines the benefits of Agentic AI in asset management. AI in the asset management industry has two entry points. One entry is the productivity question where AI will support elements of productivity. The other entry point is the nascent contribution that AI, agents and LLMs can have to augment investment professionals. We illustrate the benefit of AI for instance in an engagement process (Tilly *et al.*, 2025). We attach ourselves first to the definition of agentic AI. We follow NVIDIA (2024) and identify the following properties for an AI object to be considered an agent: it will require that it perceives, reasons, acts and learns. Short of one of these properties, we are in the presence of an AI algorithm or process but not of an agent. Zhao *et al.* (2025) propose a fundamental agent, a sentiment agent, and a valuation agent, which are astutely brought into a multi-agent workflow. We can decompose the definition of the valuation agent by the authors: "The Valuation Agent is designed to analyze [perceive] stock prices and volumes, providing a valuation assessment [reason] to inform [act] the stockâs relative significance within a portfolio". Had the backtest period have been longer than four months, we might have seen a feedback loop in the agent construction. Fatouros *et al.* (2025) also proposes stock analysis with agents. The authors indicate that the fundamental agent has received significant enhancement relative to the previous version, hence positioning the feedback-loop for the agent outside of the definition of the agent itself. We confirm that the definition of the agent does not contain self-improvement options.

In the context of this paper, if we were to reason with an AI agent for the calibration for exposure control as introduced in Subsection 4.4, we would need indeed to include a feedback loop. In the tools that we have developed following Blanqué *et al.* (2022), we monitor the popular GDELT identifiers which we have not been selected into our definition of narratives. For instance, while we had prepared for the eventuality of trade barriers associated to tarrifs, we had not anticipated the popularity of smaller government in the US which was a commented topix after the start of the second Trump Administration.

For the "Perceive step", our proposed agent can itself be in dialogue with agents researching news datasets. For our "Reason step", we have introduced normalisation of our narratives over a three-month period to capture its dynamics. This can be continuously challenged or multiple time periods could be introduced. We propose a calibration for the exposure to the assets sensitive to geopolitical risk. Our analysis is based on an approach derived from credit rating transition analysis. An agent reading academic research could perfectly look into the academic literature for other sources of inspiration and propose alternative solutions. We also took a very conservative route by initiating trading only after two weeks of confirmation of the dominance of the geopolitical narrative. As for the "Learn" feedback loop, the ongoing evolution of the content of the geopolitical narrative is necessary.

Algorithm 1 AI Agent for Narrative-Based Adaptive Calibration for exposure control under Geopolitical narrative dominance

for t = each end of week, ... **do**

Perceive: access the GDELT dataset or any other dataset that enables the capture of global events

Reason: define the dominance of the Geopolitical Narrative from the Narrative framework defined in Blanqué *et al.* (2022)

Act: calibration of the exposure to geopolitical sensitive assets as proposed in Subsection 4.4

Learn (feedback loop): reinforce the robustness of the calibration signal with regards, to the definition of the Geopolitical narrative itself and also the stabilization of the turnover

end for

We propose a framework for the analysis of the popular GDELT identifiers which are not integrated in the narratives framewor. In this paper we have taken the decision to complement the geopolitical narrative with the social narrative, as the subject of affordability has entered the debate. This was a discretionary decision. We can imagine a fine-tuned language model with a classification task to identify if GDELT identifiers are to be integrated in the geopolitical narrative. It would require a forward-looking analysis of the robustness of the updates very similar to the spirit of the index providers who propose a robust methodology for monthly or quarterly entry and exits of assets from their indices.

6 Conclusion

This paper investigates the systematic integration of narrative-based geopolitical risk into low-volatility equity strategies, building on a framework that conditions portfolio construction on firm-level sentiment scores extracted from news and expert reports. Our findings demonstrate that adaptive strategies, particularly Adaptive-2, enhance the performance and resilience of low-volatility portfolios under institutional constraints. By combining volatility-based screening with geopolitical adjustments during weight allocation, Adaptive-2 delivers roughly 240 basis points of additional annualized return without materially increasing ex-post volatility, effectively preserving the defensive characteristics central to low-volatility mandates.

The adaptive strategies exhibit particularly strong performance during periods of heightened market stress. During the COVID-19 crisis and the 2022 geopolitical/monetary stress episode, Adaptive-1 and Adaptive-2 limited losses, recovered more quickly than static benchmarks, and captured upside during recovery phases, all while maintaining disciplined risk profiles. Improvements in maximum drawdown, though modest in percentage terms, translate into meaningful capital preservation for institutional investors. These results highlight that narrative-informed, adaptive low-volatility strategies can provide both enhanced risk-adjusted returns and greater resilience, offering a practical pathway for incorporating textual measures of geopolitical risk into systematic portfolio management.

A Appendix

A.1 Ravenpack Themes

Table 5 lists the Ravenpack events used to filter for local political and geopolitical risk.

government	**civil unrest**	**foreign relations**
aid	**censorship**	**crime**
cybersecurity	**legal**	**health**
migration	**capital punishment**	**child labour**
corruption	**discrimination**	**forced labor**
gender discrimination	**harassment**	**human rights violation**
legal issues	**money laundering**	**racial discrimination**
reputation	**sanctions**	**security**
war conflict		

Table 5: Ravenpack Events used for filtering raw data.

A.2 Blended-Score Integration Framework: Theoretical Properties

Theorem 1 (Existence and Uniqueness). *Let* $\mathcal{F} = \{w \in \mathbb{R}^n : \sum_i w_i = 1, w_{min} \leq w_i \leq w_{max}\}$ *denote the feasible set and* $f(w) = \lambda_1 s^\top w - \lambda_2 \|w - w^{bench}\|_2^2$ *the objective function where* $s = (S_i^{opt})_i$. *Then:*

(i) $\mathcal{F}$ *is non-empty, compact, and convex*

(ii) f *is strictly concave on* $\mathcal{F}$

(iii) Problem (10) admits a unique global optimum $w^* \in \mathcal{F}$

Proof. See Appendix B. □

The uniqueness guarantee is economically significant: it ensures reproducibility and eliminates path-dependence from solver initialization which is critical for production implementation.

Lemma 2 (Lipschitz Continuity of Optimal Weights). *Let* $w^*(s, w^{bench})$ *denote the optimal solution as a function of composite scores* s *and benchmark weights* w^{bench}. *Then there exists* $L > 0$ *such that:*

$$\|w^*(s_1, w_1^{bench}) - w^*(s_2, w_2^{bench})\|_2 \leq L(\|s_1 - s_2\|_2 + \|w_1^{bench} - w_2^{bench}\|_2) \tag{14}$$

Proof. See Appendix B. □

Lemma 2 provides robustness guarantees: small perturbations in input signals (due to measurement error or sampling variation) induce proportionally small changes in optimal weights. This Lipschitz stability is particularly valuable when incorporating noisy sentiment-derived geopolitical scores, bounding the worst-case impact of signal noise on portfolio weights.

Corollary 3 (Bounded Impact of Signal Estimation Error). *Suppose the true composite scores are* $\tilde{s}$ *but we observe* $s = \tilde{s} + \epsilon$ *where* $\|\epsilon\|_\infty \leq \delta$. *Then:*

$$\|w^*(s) - w^*(\tilde{s})\|_2 \leq L\sqrt{n}\delta \tag{15}$$

Proof. Follows from Lemma 2 since $\|s - \tilde{s}\|_2 = \|\epsilon\|_2 \leq \sqrt{n}\|\epsilon\|_\infty \leq \sqrt{n}\delta$. □

This bound justifies our use of news-sentiment-derived geopolitical scores despite inherent measurement noise: provided δ remains moderate (as ensured by our z-score truncation), portfolio weight errors remain controlled.

Proposition 4 (Monotone Response to Signal Strength). *For assets* i, j *with identical benchmark weights* $w_i^{bench} = w_j^{bench}$ *and box constraints, if* $S_i^{opt} > S_j^{opt}$, *then the optimal weights satisfy* $w_i^* \geq w_j^*$, *with strict inequality unless both constraints bind.*

Proof. See Appendix B. □

Proposition 4 formalizes the intuitive property that higher-scored assets receive higher weights, providing interpretability and ensuring the portfolio respects the factor signals. This monotonicity is automatic for ranking-based portfolios but requires proof for optimization-based construction.

Proposition 5 (Convex Combination Property). *The optimal solution* w^* *is a convex combination of the signal-maximizing portfolio* w^{signal} *(that maximizes* $s^\top w$ *subject to* $w \in \mathcal{F}$*) and the benchmark portfolio* w^{bench}, *with mixing controlled by* λ_2.

Proof. See Appendix B. □

This property clarifies the penalty's role: λ_2 controls the trade-off between pure signal capture (w^{signal}) and benchmark tracking (w^{bench}), with the optimal solution lying between these extremes.

B Optimization Problem: Theoretical Analysis

B.1 Proof of Theorem 1

Part (i): Feasibility. The constraint set $\mathcal{F} = \{w \in \mathbb{R}^n : \sum_i w_i = 1, w_{min} \leq w_i \leq w_{max}\}$ is non-empty provided $nw_{min} \leq 1 \leq nw_{max}$. With $w_{min} = 10^{-4}$, $w_{max} = 0.10$, and $n \approx 150$, we have $0.015 \leq 1 \leq 15$, ensuring non-emptiness. The set $\mathcal{F}$ is the intersection of the hyperplane $\{w : \mathbf{1}^T w = 1\}$ (closed and convex) and the hypercube $[w_{min}, w_{max}]^n$ (closed, bounded, and convex). Since both sets are closed and the hypercube is bounded, the set $\mathcal{F}$ is closed and bounded, hence compact by the Heine-Borel theorem. The intersection of convex sets is convex, establishing that $\mathcal{F}$ is convex.

Part (ii): Strict Concavity. Expanding the objective function:

$$f(w) = \lambda_1 s^T w - \lambda_2 \sum_i (w_i - w_i^{bench})^2 \tag{16}$$

$$= \lambda_1 s^T w - \lambda_2 (w^T w - 2(w^{bench})^T w + \|w^{bench}\|_2^2) \tag{17}$$

$$= -\lambda_2 w^T w + (\lambda_1 s + 2\lambda_2 w^{bench})^T w - \lambda_2 \|w^{bench}\|_2^2 \tag{18}$$

The Hessian matrix is $\nabla^2 f(w) = -2\lambda_2 I$. Since $\lambda_2 > 0$, we have $\nabla^2 f(w) = -2\lambda_2 I < 0$ (negative definite), establishing strict concavity on $\mathbb{R}^n$, and therefore on any subset $\mathcal{F} \subset \mathbb{R}^n$.

Part (iii): Uniqueness. By Weierstrass's theorem, the continuous function f on the compact set $\mathcal{F}$ attains its maximum. Suppose by contradiction that there exist two distinct maximizers $w^1, w^2 \in \mathcal{F}$ with $w^1 \neq w^2$ such that $f(w^1) = f(w^2) = \max_{w \in \mathcal{F}} f(w)$. By convexity of $\mathcal{F}$, the midpoint $\bar{w} = \frac{1}{2}w^1 + \frac{1}{2}w^2$ belongs to $\mathcal{F}$. By strict concavity of f:

$$f(\bar{w}) > \frac{1}{2} f(w^1) + \frac{1}{2} f(w^2) = f(w^1) \tag{19}$$

This contradicts the assumption that w^1 is a maximizer. Therefore, the optimal solution is unique. □

B.2 Proof of Lemma 2

Let $w^1 = w^*(s_1, w_1^{bench})$ and $w^2 = w^*(s_2, w_2^{bench})$ denote optimal solutions for different parameter sets. Both satisfy the KKT stationarity condition (see subsection B.5):

$$\lambda_1 s_1 - 2\lambda_2 (w^1 - w_1^{bench}) - \nu_1 \mathbf{1} + \mu^1 - \xi^1 = 0 \tag{20}$$

$$\lambda_1 s_2 - 2\lambda_2 (w^2 - w_2^{bench}) - \nu_2 \mathbf{1} + \mu^2 - \xi^2 = 0 \tag{21}$$

Subtracting these equations:

$$\lambda_1 (s_1 - s_2) - 2\lambda_2 (w^1 - w^2) + 2\lambda_2 (w_1^{bench} - w_2^{bench}) = (\nu_1 - \nu_2)\mathbf{1} - (\mu^1 - \mu^2) + (\xi^1 - \xi^2) \tag{22}$$

Taking the ℓ_2 norm on both sides and applying the triangle inequality:

$$2\lambda_2 \|w^1 - w^2\|_2 \leq \lambda_1 \|s_1 - s_2\|_2 + 2\lambda_2 \|w_1^{bench} - w_2^{bench}\|_2 + \|\nu_1 \mathbf{1} - \nu_2 \mathbf{1}\|_2 + \|\mu^1 - \mu^2\|_2 + \|\xi^1 - \xi^2\|_2 \tag{23}$$

The dual variables (ν, μ, ξ) are bounded on the compact set $\mathcal{F}$ (see Proposition 6). Specifically, there exists a constant $M > 0$ such that $|\nu_i| \leq M$, $\|\mu\|_2 \leq M$, $\|\xi\|_2 \leq M$ for all feasible

solutions. Using the relationship $\|\nu_1\mathbf{1} - \nu_2\mathbf{1}\|_2 = |\nu_1 - \nu_2| \cdot \|\mathbf{1}\|_2 = |\nu_1 - \nu_2|\sqrt{n} \leq 2M\sqrt{n}$, we obtain:

$$2\lambda_2\|w^1 - w^2\|_2 \leq \lambda_1\|s_1 - s_2\|_2 + 2\lambda_2\|w_1^{bench} - w_2^{bench}\|_2 + (2M\sqrt{n} + 2M) \tag{24}$$

The constant term $(2M\sqrt{n} + 2M)$ can be absorbed into the bound by noting that for distinct inputs, the term $\|s_1 - s_2\|_2 + \|w_1^{bench} - w_2^{bench}\|_2$ is at least on the order of machine precision. We can therefore write:

$$\|w^1 - w^2\|_2 \leq \frac{\lambda_1}{2\lambda_2}\|s_1 - s_2\|_2 + \|w_1^{bench} - w_2^{bench}\|_2 + C(\|s_1 - s_2\|_2 + \|w_1^{bench} - w_2^{bench}\|_2) \tag{25}$$

where $C = M(\sqrt{n}+1)/\lambda_2$. Setting $L = \max\{\frac{\lambda_1}{2\lambda_2} + C, 1 + C\}$ yields the desired Lipschitz constant. □

B.3 Proof of Corollary 3

Proof. Suppose the observed scores satisfy $S_i^{opt} = \tilde{S}_i + \epsilon_i$ where $\tilde{S}_i$ represents the true underlying scores and ϵ_i represents measurement error with $\|\epsilon\|_\infty \leq \delta$.

We apply Lemma 2 with $s_1 = S^{opt}$, $s_2 = \tilde{S}$, and $w_1^{bench} = w_2^{bench}$ (benchmark weights remain unchanged). The Lipschitz condition yields:

$$\|w^*(S^{opt}) - w^*(\tilde{S})\|_2 \leq L(\|S^{opt} - \tilde{S}\|_2 + \|w^{bench} - w^{bench}\|_2) = L\|S^{opt} - \tilde{S}\|_2 \tag{26}$$

Since $S^{opt} - \tilde{S} = \epsilon$, we need to bound $\|\epsilon\|_2$ in terms of $\|\epsilon\|_\infty$. Using the relationship between ℓ_2 and ℓ_∞ norms:

$$\|\epsilon\|_2 = \sqrt{\sum_{i=1}^{n} \epsilon_i^2} \leq \sqrt{\sum_{i=1}^{n} \|\epsilon\|_\infty^2} = \sqrt{n\|\epsilon\|_\infty^2} = \sqrt{n}\|\epsilon\|_\infty \leq \sqrt{n}\delta \tag{27}$$

Substituting this bound into the Lipschitz inequality:

$$\|w^*(S^{opt}) - w^*(\tilde{S})\|_2 \leq L\sqrt{n}\delta \tag{28}$$

This establishes that the portfolio weight error is bounded linearly in the score measurement error δ, with the bound scaling as $\sqrt{n}$ due to the dimensionality of the score vector. □

Economic interpretation. This corollary provides a quantitative guarantee on robustness to noisy signals. For geopolitical sentiment scores derived from news analytics, measurement error arises from sentiment classification accuracy, entity recognition errors, and aggregation noise. If we can bound the maximum error in any individual score by δ (for instance, if sentiment classifiers have known accuracy guarantees), then the portfolio weight error is proportionally bounded. The $\sqrt{n}$ factor reflects that while individual score errors may aggregate through summation, they do not compound multiplicatively. The impact remains sublinear in portfolio size, providing assurance that estimation error in signals translates to controlled error in portfolio weights rather than catastrophic failures.

B.4 Proof of Proposition 4

Proof. Consider the optimization problem (10) and let $w^*(s)$ denote the optimal solution for score vector $s = (s_1, \ldots, s_n)$. From the KKT stationarity condition (41):

$$\lambda_1 s_i - 2\lambda_2(w_i^* - w_i^{bench}) - \nu + \mu_i - \xi_i = 0, \quad \forall i \tag{29}$$

For an interior solution where asset i is not at boundary constraints (i.e., $w_{min} < w_i^* < w_{max}$), the complementarity conditions (45) imply that $\mu_i = \xi_i = 0$. Solving equation (29) for w_i^*:

$$w_i^* = w_i^{bench} + \frac{\lambda_1 s_i - \nu}{2\lambda_2} \tag{30}$$

Now consider a perturbation where we increase score s_i by a small amount $\Delta s_i > 0$ while holding all other scores fixed: $s_j' = s_j$ for $j \neq i$ and $s_i' = s_i + \Delta s_i$. Let w'^* and ν' denote the new optimal solution and Lagrange multiplier respectively.

Taking the partial derivative of equation (30) with respect to s_i (treating ν as implicitly dependent on s_i through the budget constraint):

$$\frac{\partial w_i^*}{\partial s_i} = \frac{\lambda_1}{2\lambda_2} - \frac{1}{2\lambda_2}\frac{\partial \nu}{\partial s_i} \tag{31}$$

To determine $\frac{\partial \nu}{\partial s_i}$, we use the budget constraint $\sum_{j=1}^n w_j^* = 1$. Differentiating both sides with respect to s_i:

$$\sum_{j=1}^{n} \frac{\partial w_j^*}{\partial s_i} = 0 \tag{32}$$

For interior assets with $j \neq i$, we have from equation (30):

$$\frac{\partial w_j^*}{\partial s_i} = -\frac{1}{2\lambda_2}\frac{\partial \nu}{\partial s_i}, \quad j \neq i \tag{33}$$

For asset i itself:

$$\frac{\partial w_i^*}{\partial s_i} = \frac{\lambda_1}{2\lambda_2} - \frac{1}{2\lambda_2}\frac{\partial \nu}{\partial s_i} \tag{34}$$

Summing over all interior assets (let $\mathcal{I}$ denote the interior set with $|\mathcal{I}| \geq 2$):

$$\frac{\lambda_1}{2\lambda_2} - \frac{|\mathcal{I}|}{2\lambda_2}\frac{\partial \nu}{\partial s_i} = 0 \tag{35}$$

Solving for $\frac{\partial \nu}{\partial s_i}$:

$$\frac{\partial \nu}{\partial s_i} = \frac{\lambda_1}{|\mathcal{I}|} \tag{36}$$

Substituting back into the expression for $\frac{\partial w_i^*}{\partial s_i}$:

$$\frac{\partial w_i^*}{\partial s_i} = \frac{\lambda_1}{2\lambda_2} - \frac{1}{2\lambda_2}\cdot\frac{\lambda_1}{|\mathcal{I}|} = \frac{\lambda_1}{2\lambda_2}\left(1 - \frac{1}{|\mathcal{I}|}\right) = \frac{\lambda_1}{2\lambda_2}\cdot\frac{|\mathcal{I}| - 1}{|\mathcal{I}|} > 0 \tag{37}$$

The final inequality holds because $|\mathcal{I}| \geq 2$ (at least two assets must be interior for a non-degenerate portfolio) and $\lambda_1, \lambda_2 > 0$.

For other interior assets $j \neq i$:

$$\frac{\partial w_j^*}{\partial s_i} = -\frac{1}{2\lambda_2}\cdot\frac{\lambda_1}{|\mathcal{I}|} = -\frac{\lambda_1}{2\lambda_2|\mathcal{I}|} < 0 \tag{38}$$

Therefore, increasing s_i strictly increases w_i^* and strictly decreases w_j^* for at least some $j \neq i$ (specifically, all other interior assets). This establishes the monotonicity property. □

Implications for portfolio interpretability. Proposition 4 ensures that the optimization does not produce counterintuitive results where higher-scored assets receive lower weights. This monotonic response property is critical for factor-based strategies where portfolio managers must explain allocation decisions to clients and risk committees. The proof reveals that the weight response to score changes has magnitude $\frac{\lambda_1}{2\lambda_2} \cdot \frac{|\mathcal{I}|-1}{|\mathcal{I}|}$, which approaches $\frac{\lambda_1}{2\lambda_2}$ as the number of interior assets increases. With our parameter choices ($\lambda_1 = 1$, $\lambda_2 = 10$), a one-unit increase in score s_i increases weight w_i^* by approximately $\frac{1}{20} = 0.05$, or 5 percentage points, demonstrating substantial signal responsiveness while maintaining stability through the regularization term.

Boundary cases. The proof above assumes interior solutions. When asset i hits a boundary constraint ($w_i^* = w_{min}$ or $w_i^* = w_{max}$), the monotonicity property still holds but manifests differently. If $w_i^* = w_{max}$ (asset at concentration limit), further increases in s_i cannot increase w_i^* directly, but the complementarity condition $\xi_i(w_{max} - w_i^*) = 0$ implies $\xi_i > 0$ becomes active. From equation (29), this requires:

$$\xi_i = \lambda_1 s_i - 2\lambda_2(w_{max} - w_i^{bench}) - \nu \tag{39}$$

As s_i increases, ξ_i increases proportionally, reflecting that the constraint becomes more binding. The asset would receive even higher weight if not capped by the concentration limit. Similarly, if $w_i^* = w_{min}$, the asset remains at the floor as s_i decreases, with μ_i increasing to enforce the lower bound. In both cases, the economic intuition of monotonicity is preserved: higher scores correspond to higher desired allocations, even when operational constraints prevent full expression of that preference through actual weight changes.

B.5 KKT Optimality Conditions

The Lagrangian for problem (10) is:

$$\mathcal{L}(w, \nu, \mu, \xi) = \lambda_1 s^T w - \lambda_2 \|w - w^{bench}\|_2^2 - \nu(\mathbf{1}^T w - 1) - \mu^T(w_{min}\mathbf{1} - w) - \xi^T(w - w_{max}\mathbf{1}) \tag{40}$$

The KKT conditions are:

$$\nabla_w \mathcal{L} = \lambda_1 s - 2\lambda_2(w - w^{bench}) - \nu\mathbf{1} + \mu - \xi = 0 \tag{41}$$

$$\mathbf{1}^T w = 1 \tag{42}$$

$$w_{min}\mathbf{1} \leq w \leq w_{max}\mathbf{1} \tag{43}$$

$$\mu, \xi \geq 0 \tag{44}$$

$$\mu_i(w_i - w_{min}) = 0, \quad \xi_i(w_{max} - w_i) = 0 \quad \forall i \tag{45}$$

Since the problem is convex with affine constraints, Slater's condition holds (the interior of $\mathcal{F}$ is non-empty: for example, $w_i = 1/n \in (w_{min}, w_{max})$ for our parameter values). Therefore, the KKT conditions are necessary and sufficient for global optimality.

From equation (41), we can express:

$$w_i = w_i^{bench} + \frac{\lambda_1 s_i - \nu + \mu_i - \xi_i}{2\lambda_2} \tag{46}$$

For assets with $w_{min} < w_i < w_{max}$ (interior solutions), complementarity condition (45) implies $\mu_i = \xi_i = 0$, simplifying to:

$$w_i = w_i^{bench} + \frac{\lambda_1 s_i - \nu}{2\lambda_2} \tag{47}$$

The Lagrange multiplier ν is determined by the budget constraint $\sum_i w_i = 1$. Summing equation (46) over all assets (accounting for active constraints):

$$1 = \sum_i w_i^{bench} + \frac{\lambda_1 \sum_i s_i - n\nu + \sum_i (\mu_i - \xi_i)}{2\lambda_2} \tag{48}$$

Since $\sum_i w_i^{bench} = 1$ and using complementarity conditions to simplify the dual terms, this yields an implicit equation for ν that is solvable numerically through the optimization process.

B.6 Sensitivity Analysis

Proposition 6 (Dual Variable Bounds)**.** *The Lagrange multipliers satisfy:*

$$|\nu| \leq \lambda_1 \|s\|_\infty + 2\lambda_2(\|w^{bench}\|_\infty + w_{max}), \quad \|\mu\|_\infty, \|\xi\|_\infty \leq \lambda_1 \|s\|_\infty + 2\lambda_2(w_{max} + \|w^{bench}\|_\infty) \tag{49}$$

Proof. From equation (41), for any component i:

$$\nu - \mu_i + \xi_i = \lambda_1 s_i - 2\lambda_2(w_i - w_i^{bench}) \tag{50}$$

Summing over all i:

$$n\nu - \sum_i \mu_i + \sum_i \xi_i = \lambda_1 \sum_i s_i - 2\lambda_2 \sum_i (w_i - w_i^{bench}) \tag{51}$$

Since $\sum_i w_i = \sum_i w_i^{bench} = 1$, the right-hand side simplifies to $\lambda_1 \sum_i s_i$. Taking absolute values and bounding:

$$|n\nu| \leq \lambda_1 |\sum_i s_i| + \sum_i \mu_i + \sum_i \xi_i \leq \lambda_1 n \|s\|_\infty + \sum_i \mu_i + \sum_i \xi_i \tag{52}$$

For individual components, from equation (46):

$$\mu_i - \xi_i = \nu - \lambda_1 s_i + 2\lambda_2(w_i - w_i^{bench}) \tag{53}$$

Since $w_{min} \leq w_i \leq w_{max}$ and $0 \leq w_i^{bench} \leq \|w^{bench}\|_\infty$, we obtain:

$$|\mu_i - \xi_i| \leq |\nu| + \lambda_1 |s_i| + 2\lambda_2(w_{max} + \|w^{bench}\|_\infty) \tag{54}$$

Combining these bounds with the non-negativity conditions $\mu_i, \xi_i \geq 0$ and the fact that at most one of μ_i, ξ_i is positive (from complementarity), the stated bounds follow. □

C Convex Optimization: Numerical Implementation

C.1 Problem Reformulation

Problem (10) can be reformulated as a standard quadratic program:

$$\min_w \frac{1}{2} w^T P w + q^T w \quad \text{s.t.} \quad Aw = b,\ \ell \leq w \leq u \tag{55}$$

where the problem data are defined as:

$$P = 2\lambda_2 I_{n\times n} \quad \text{(positive definite Hessian)} \tag{56}$$
$$q = -\lambda_1 s - 2\lambda_2 w^{bench} \quad \text{(linear term)} \tag{57}$$
$$A = \mathbf{1}_{1\times n}^T \quad \text{(budget constraint matrix)} \tag{58}$$
$$b = 1 \quad \text{(budget right-hand side)} \tag{59}$$
$$\ell = w_{min}\mathbf{1}, \quad u = w_{max}\mathbf{1} \quad \text{(box constraints)} \tag{60}$$

The Hessian matrix $P = 2\lambda_2 I$ has condition number $\kappa(P) = 1$ (perfectly conditioned), ensuring excellent numerical stability. The constraint matrix $A = \mathbf{1}^T$ also has condition number $\kappa(A) = 1$ (single equality constraint with uniform coefficients), avoiding ill-conditioning in KKT system factorizations that can plague more complex portfolio optimization problems.

C.2 Solver Hierarchy and Convergence Properties

We employ the `cvxpy` framework (Diamond & Boyd, 2016), which automatically transforms the problem into disciplined convex programming form and dispatches to appropriate solvers. The solver cascade proceeds through the following hierarchy:

ECOS (Embedded Conic Solver). The primary solver is ECOS (Domahidi *et al.*, 2013), which implements a primal-dual interior-point method for second-order cone programs. The algorithm solves the perturbed KKT system with adaptive barrier parameter updates. Convergence is declared when the primal residual $\|r_p\|_2$, dual residual $\|r_c\|_2$, and duality gap all fall below $\epsilon_{abs} + \epsilon_{rel} \cdot \text{scale}$ where $\epsilon_{abs} = \epsilon_{rel} = 10^{-4}$. Under Slater's constraint qualification (satisfied in our problem since the interior of $\mathcal{F}$ is non-empty), ECOS enjoys guaranteed global convergence with theoretical complexity $\mathcal{O}(n^{3.5}\log(1/\epsilon))$ (Nesterov & Nemirovskii, 1994). In practice, for portfolios of size $n \approx 150$, ECOS typically converges within 15-30 iterations requiring 0.2-0.8 seconds on standard hardware.

OSQP (Operator Splitting QP Solver). If ECOS fails to converge or encounters numerical difficulties, the algorithm falls back to OSQP (Stellato *et al.*, 2020), which implements the alternating direction method of multipliers. This is a first-order method with per-iteration computational cost $\mathcal{O}(n^2)$ dominated by Cholesky factorization of the matrix $P + \rho A^T A$, which is computed once during preprocessing. The convergence rate is $\mathcal{O}(1/k)$ where k denotes the iteration count (Boyd *et al.*, 2011). OSQP particularly excels at warm-starting: initializing with $w^0 = w^*_{t-1}$ at rebalancing date t exploits the temporal stability guaranteed by Lemma 2, often achieving convergence five to ten times faster than cold-start initialization. This makes OSQP especially attractive for sequential rebalancing applications.

SCS (Splitting Conic Solver). As a robust fallback, SCS (O'Donoghue *et al.*, 2016) applies Douglas-Rachford splitting to the homogeneous self-dual embedding of the conic program. The method handles infeasibility certificates via auxiliary variables and provides guaranteed convergence for any conic program, including cases where feasibility is uncertain. SCS exhibits convergence rate $\mathcal{O}(1/k)$ with minimal memory footprint $\mathcal{O}(n)$, making it suitable for large-scale problems. While typically slower than ECOS or OSQP for well-conditioned problems, SCS provides a reliable last resort when other methods encounter numerical difficulties.

CLARABEL. The solver hierarchy also includes CLARABEL when available, a modern open-source interior-point solver with improved numerical conditioning through aggressive regularization and adaptive step-sizing. CLARABEL represents a state-of-the-art alternative to commercial solvers and often provides performance comparable to proprietary solutions.

All solvers in the hierarchy satisfy the theoretical requirements of Theorem 1: they are designed for convex optimization problems with linear constraints, ensuring global convergence to the unique optimum w^*.

C.3 Convergence Diagnostics and Validation

Upon solver termination, the algorithm performs comprehensive validation:

Status verification. The solver must report either `OPTIMAL` or `OPTIMAL_INACCURATE` status. The latter is acceptable provided that residuals remain below 10^{-3}, indicating that while theoretical optimality criteria may not be fully satisfied to machine precision, the solution is sufficiently accurate for practical portfolio implementation.

Constraint satisfaction. We verify budget satisfaction $|\mathbf{1}^T w^* - 1| < 10^{-6}$ and bound compliance $w_{min} - 10^{-6} \leq w_i^* \leq w_{max} + 10^{-6}$ for all assets i. The tolerance of 10^{-6} accommodates minor floating-point rounding errors while ensuring constraints are satisfied to practical precision.

Numerical validity. The solution vector w^* is checked for NaN (not-a-number) or Inf (infinity) values, which would indicate numerical breakdown during the solution process.

Objective improvement. We verify that the optimized portfolio achieves score improvement relative to the benchmark: $s^T w^* \geq s^T w^{bench}$. This confirms that the optimization has successfully captured signal information rather than converging to a degenerate solution.

If any validation check fails, the algorithm proceeds to the next solver in the hierarchy. If all solvers fail these checks, the fallback strategy described in subsection C.4 is invoked.

Post-processing for numerical precision. Minor constraint violations on the order of machine precision ($\epsilon_{machine} \approx 10^{-16}$ for double-precision arithmetic) are corrected through a three-step refinement process. First, we renormalize to satisfy the budget constraint: $w^* \leftarrow w^*/(\mathbf{1}^T w^*)$. Second, we clip weights to bounds component-wise: $w^* \leftarrow \max(w_{min}, \min(w^*, w_{max}))$. Third, we apply final renormalization: $w^* \leftarrow w^*/(\mathbf{1}^T w^*)$. These corrections preserve economic constraints while eliminating accumulated round-off errors from iterative solving.

C.4 Fallback Strategy: Theoretical Justification

If all solvers fail to converge with the penalized objective (10), we solve the simplified linear program:

$$\max_{w \in \mathcal{F}} \; s^T w \tag{61}$$

[Fallback Existence] Problem (61) always admits a solution.

Proof. The feasible set $\mathcal{F}$ is non-empty, compact, and convex by Theorem 1 part (i). The objective function $f_{fallback}(w) = s^T w$ is continuous on $\mathbb{R}^n$, and therefore continuous on the

subset $\mathcal{F} \subset \mathbb{R}^n$. By the extreme value theorem (Weierstrass), any continuous function on a non-empty compact set attains its maximum. Therefore, there exists $w^{FB} \in \mathcal{F}$ such that:

$$s^T w^{FB} = \max_{w \in \mathcal{F}} s^T w \tag{62}$$

□

[Fallback Solution Structure] The optimal solution to problem (61) has the explicit structure:

$$w_i^{FB} = \begin{cases} w_{max} & \text{if } i \in \mathcal{T}_k \\ 1 - k w_{max} & \text{if } i = j_k \\ w_{min} & \text{otherwise} \end{cases} \tag{63}$$

where $\mathcal{T}_k = \{i : \text{rank}(s_i) \leq k\}$ contains the top k assets by score with $k = \lfloor 1/w_{max} \rfloor$, and j_k denotes the asset with the $(k+1)$-th highest score.

Proof. The linear objective $s^T w = \sum_i s_i w_i$ is maximized by allocating maximal feasible weight to assets with highest scores. Without loss of generality, reindex assets so that $s_1 \geq s_2 \geq \cdots \geq s_n$. The greedy algorithm proceeds as follows. Set $w_1 = w_{max}$, consuming budget mass w_{max}. If $2w_{max} \leq 1$, set $w_2 = w_{max}$, consuming additional mass w_{max}. Continue this process until $k = \lfloor 1/w_{max} \rfloor$ assets have received weight w_{max}, with total allocated mass $k w_{max}$. The remaining budget mass $1 - k w_{max}$ is assigned to the next highest-scored asset: $w_{k+1} = 1 - k w_{max}$. All remaining assets receive minimal weight: $w_i = w_{min}$ for $i > k+1$.

This allocation satisfies the budget constraint since:

$$\sum_i w_i = k w_{max} + (1 - k w_{max}) + (n - k - 1) w_{min} = 1 \tag{64}$$

where the equality holds because we choose w_{min} sufficiently small such that $(n - k - 1) w_{min}$ is negligible relative to the residual mass $1 - k w_{max}$, or more precisely, we can define w_{min} such that the constraint is exactly satisfied.

To verify optimality, suppose by contradiction that an alternative allocation $\tilde{w} \in \mathcal{F}$ exists with $s^T \tilde{w} > s^T w^{FB}$. Since scores are ordered as $s_1 \geq s_2 \geq \cdots \geq s_n$, any reallocation of weight from higher-ranked to lower-ranked assets decreases the objective value. If $\tilde{w}_i > w_i^{FB}$ for some $i > k+1$, then by the budget constraint we must have $\tilde{w}_j < w_j^{FB}$ for some $j \leq k+1$. The change in objective is:

$$s^T \tilde{w} - s^T w^{FB} = \sum_{\ell=1}^{n} s_\ell (\tilde{w}_\ell - w_\ell^{FB}) \tag{65}$$

Focusing on the assets where weights differ, we have:

$$s^T \tilde{w} - s^T w^{FB} = s_i (\tilde{w}_i - w_i^{FB}) + s_j (\tilde{w}_j - w_j^{FB}) + \cdots \tag{66}$$

Since $s_i < s_j$ for $i > j$ (scores are decreasing) and $\tilde{w}_i - w_i^{FB} > 0$ while $\tilde{w}_j - w_j^{FB} < 0$, we can write:

$$s^T \tilde{w} - s^T w^{FB} < s_j (\tilde{w}_i - w_i^{FB}) + s_j (\tilde{w}_j - w_j^{FB}) = s_j [(\tilde{w}_i - w_i^{FB}) + (\tilde{w}_j - w_j^{FB})] \tag{67}$$

By the budget constraint (assuming other weights remain unchanged for simplicity), we have $(\tilde{w}_i - w_i^{FB}) + (\tilde{w}_j - w_j^{FB}) = 0$, yielding $s^T \tilde{w} - s^T w^{FB} < 0$. This contradicts the assumption that $s^T \tilde{w} > s^T w^{FB}$, establishing the optimality of w^{FB}. □

Corollary 7 (Fallback Score Improvement)**.** *The fallback solution satisfies* $s^T w^{FB} \geq s^T w^{bench}$ *provided the scores exhibit non-trivial variation across assets.*

Proof. Since $w^{bench} \in \mathcal{F}$ by construction (benchmark weights satisfy the budget constraint and typically satisfy bounds for capitalization-weighted indices where no single asset dominates), we have:

$$s^T w^{FB} = \max_{w \in \mathcal{F}} s^T w \geq s^T w^{bench} \tag{68}$$

Equality holds only if w^{bench} itself is optimal for the linear program, which occurs when scores are constant ($s_i = c$ for all i, a degenerate case) or when w^{bench} coincidentally places maximum feasible weight on the highest-scored assets (highly unlikely for standard capitalization-weighted benchmarks, which weight by market size rather than factor scores). □

Practical interpretation. The fallback solution represents a greedy high-conviction portfolio that concentrates holdings on the top-scored assets up to the concentration limit w_{max}. This approach completely ignores the benchmark weighting and penalty terms, focusing solely on maximizing signal exposure. While this produces less diversified portfolios than the full regularized optimization (10), it ensures robustness: the portfolio construction process never fails completely, and some degree of score improvement is always achieved relative to the benchmark. In practice, fallback invocation is exceedingly rare (less than one percent of rebalancing dates in our backtest), typically occurring during extreme market conditions when both score distributions and market-cap weights exhibit unusual patterns that challenge numerical solvers.

D Performance Metrics and Statistical Tests

D.1 Risk-Adjusted Performance Ratios

We employ three standard risk-adjusted performance metrics to evaluate portfolio strategies. For simplicity and to facilitate direct comparison across portfolios, we set the risk-free rate to zero in all ratio calculations. This convention is appropriate when evaluating relative performance, as the risk-free rate cancels out in comparative rankings.

Sharpe Ratio. The Sharpe ratio measures excess return per unit of total volatility:

$$\text{Sharpe Ratio} = \frac{\mathbb{E}[R_p]}{\sigma(R_p)} \tag{69}$$

where $\mathbb{E}[R_p]$ denotes the mean portfolio return and $\sigma(R_p)$ represents the standard deviation of returns. Higher values indicate superior risk-adjusted performance. The Sharpe ratio penalizes both upside and downside volatility equally, making it appropriate for strategies with symmetric return distributions.

Sortino Ratio. The Sortino ratio refines the Sharpe ratio by penalizing only downside volatility:

$$\text{Sortino Ratio} = \frac{\mathbb{E}[R_p]}{\sigma_{\text{downside}}(R_p)} \tag{70}$$

where $\sigma_{\text{downside}}(R_p) = \sqrt{\mathbb{E}[\min(R_p, 0)^2]}$ captures volatility from negative returns only. This metric proves particularly relevant for evaluating low-volatility strategies, as it focuses exclusively on undesirable volatility while ignoring beneficial upside fluctuations.

Calmar Ratio. The Calmar ratio adjusts returns for maximum drawdown:

$$\text{Calmar Ratio} = \frac{\mathbb{E}[R_p]}{\text{MaxDD}} \tag{71}$$

where $\text{MaxDD} = \max_{t\in[0,T]} \left(\max_{s\in[0,t]} V_s - V_t\right)$ represents the maximum peak-to-trough decline in portfolio value V_t over the evaluation period. This metric emphasizes tail risk and extreme loss scenarios, providing insight into strategy resilience during crisis periods.

D.2 Wilcoxon Signed-Rank Test

To assess the statistical significance of performance differences between portfolios, we employ the one-sided Wilcoxon signed-rank test, a non-parametric alternative to the paired t-test that does not assume normality of return distributions (Rosner *et al.*, 2006).

Let $R_{p,i}$ and $R_{b,i}$ denote daily returns of the test portfolio and baseline portfolio on day i, respectively, for $i = 1, \ldots, n$. Define the return differences:

$$d_i = R_{p,i} - R_{b,i} \tag{72}$$

Hypotheses:

$$H_0 : \text{median}(d_i) = 0 \quad \text{(portfolios have identical return distributions)} \tag{73}$$
$$H_1 : \text{median}(d_i) > 0 \quad \text{(test portfolio generates higher returns)} \tag{74}$$

Test Procedure:

1. Compute absolute differences $|d_i|$ for all i where $d_i \neq 0$.
2. Rank the $|d_i|$ values from smallest to largest, assigning rank $r_i \in \{1, 2, \ldots, n\}$.
3. Compute the test statistic as the sum of ranks with positive differences:

$$W^+ = \sum_{i:d_i>0} r_i \tag{75}$$

4. Under H_0, the test statistic follows a known distribution with mean μ_{W^+} and variance $\sigma^2_{W^+}$.
5. For large samples ($n > 20$), compute the standardized statistic:

$$Z = \frac{W^+ - \mu_{W^+}}{\sigma_{W^+}} \tag{76}$$

6. Calculate the one-sided p-value as $P(Z \geq z_{\text{obs}})$ under the standard normal distribution.

We reject H_0 at significance level α if $p < \alpha$. Standard thresholds are $\alpha = 0.01$ (***), $\alpha = 0.05$ (**), and $\alpha = 0.10$ (*). The test's non-parametric nature provides robustness to outliers, fat tails, and asymmetry prevalent in financial returns.

References

Abou Rjaily, N., Roncalli, T., & Xu, J. (2024). *Risk Factor, Risk Premium and Black-Litterman Model. Risk Premium and Black-Litterman Model (October 01, 2024)*.

Baker, M., Bradley, B., & Wurgler, J. (2011). *Benchmarks as limits to arbitrage: Understanding the low-volatility anomaly. Financial Analysts Journal*, *67*(1), 40–54.

Bender, J., Briand, R., Melas, D., & Subramanian, R. A. (2013). *The Foundations of Factor Investing. MSCI Research Insight.* https://www.msci.com/documents/1296102/1336482/Foundations_of_Factor_Investing.pdf

Bettin, G., Mensi, G. M., & Recchioni, M. C. (2023). *Multifactor Risk Attribution Applied to Systemic, Climate and Geopolitical Tail Risks for the Eurozone Banking Sector. Risks*, *11*(10), 173.

Blanqué, P., Slimane, M. B., Cherief, A., Le Guenedal, T., Sekine, T., & Stagnol, L. (2022). *The Benefit of Narratives for Prediction of the S&P 500 Index. The Journal of Financial Data Science*, *4*(4), pp. 72–94. https://doi.org/10.3905/jfds.2022.1.107

Blau, B. M., & Whitby, R. J. (2017). *Range-based volatility, expected stock returns, and the low volatility anomaly. Plos one*, *12*(11), e0188517.

Blitz, D. (2016). *The value of low volatility. Journal of Portfolio Management, Forthcoming.*

Blitz, D., Van Vliet, P., & Baltussen, G. (2019). *The volatility effect revisited. The Journal of Portfolio Management*, *46*(2).

Boyd, S., Parikh, N., Chu, E., Peleato, B., & Eckstein, J. (2011). *Distributed optimization and statistical learning via the alternating direction method of multipliers. Foundations and Trends in Machine Learning*, *3*(1), 1–122.

Brignone, D., Gambetti, L., & Ricci, M. (2025). *Geopolitical risk shocks: when size matters. Bank of England Staff Working Paper Series*, (1118).

Caldara, D., & Iacoviello, M. (2017). *Measuring geopolitical risk (Working Paper). Board of Governors of the Federal Reserve System.*

Caldara, D., & Iacoviello, M. (2022). *Measuring geopolitical risk. American economic review*, *112*(4), 1194–1225.

Clarke, R., De Silva, H., & Thorley, S. (2011). *Minimum-variance portfolio composition. Journal of Portfolio Management*, *37*(2), 31.

Cochrane, J. H. (1999). *Portfolio advice for a multifactor world.*

Cunado, J., Gupta, R., Lau, C. K. M., & Sheng, X. (2020). *Time-varying impact of geopolitical risks on oil prices. Defence and Peace Economics*, *31*(6), 692–706.

DeMiguel, V., Garlappi, L., & Uppal, R. (2009). *Optimal versus naive diversification: How inefficient is the 1/N portfolio strategy? Review of Financial Studies*, *22*(5), 1915–1953.

DeMiguel, V., Martín-utrera, A., & Uppal, R. (2024). *A Multifactor Perspective on Volatility-Managed Portfolios. The Journal of Finance*, *79*(6), 3859–3891.

Diamond, S., & Boyd, S. (2016). *CVXPY: A Python-embedded modeling language for convex optimization. Journal of Machine Learning Research*, *17*(83), 1–5.

Domahidi, A., Chu, E., & Boyd, S. (2013). *ECOS: An SOCP solver for embedded systems. European Control Conference*, 3071–3076.

Fama, E. F., & French, K. R. (1996). *Multifactor explanations of asset pricing anomalies. The journal of finance*, *51*(1), 55–84.

Fatouros, G., Metaxas, K., Soldatos, J., & Karathanassis, M. (2025). *MarketSenseAI 2.0: Enhancing Stock Analysis through LLM Agents*. https://arxiv.org/abs/2502.00415

Francis, J. R., & Chia, R. C.-J. (2023). *Geopolitical Risk (GPR) and its Predictability: A Systematic Literature Review. Sciences*, *13*(9), 1148–1166.

Ghayur, K., Heaney, R., & Platt, S. (2018). *Constructing long-only multifactor strategies: portfolio blending vs. signal blending. Financial Analysts Journal*, *74*(3), 70–85.

Gkillas, K., Gupta, R., & Wohar, M. E. (2018). *Volatility jumps: The role of geopolitical risks. Finance Research Letters*, *27*, 247–258.

Hastie, T., Tibshirani, R., & Friedman, J. (2009). The Elements Of Statistical Learning: Data Mining, Inference, And Prediction. Springer.

Lamine, A., & Zribi, S. (2024). *Do geopolitical risks affect stock market returns and volatilities: an analysis based on the TVP-VAR model. European Journal of Government and Economics*, *13*(2), 240–261.

Lee, S., Lee, J., & Jung, S. (2023). *the impact of geopolitical risk on stock returns: Evidence from inter-Korea geopolitics. Available at SSRN 4329355.*

Li, J. (2025). *A Study of Equity Investment Strategies for Volatility Management and Optimisation Based on Historical Volatility. Modern Economics Management Forum*, *5*, 1054. https://doi.org/10.32629/memf.v5i6.3183

Liu, F., Tang, X., & Zhou, G. (2019). *Volatility-managed portfolio: Does it really work? Journal of Portfolio Management*, *46*(1), 38–51.

MacKinlay, A. C. (1995). *Multifactor models do not explain deviations from the CAPM. Journal of financial economics*, *38*(1), 3–28.

Markowitz, H. (1952). *Portfolio selection. The Journal of Finance*, *7*(1), 77–91.

Nesterov, Y., & Nemirovskii, A. (1994). Interior-point Polynomial Algorithms In Convex Programming. SIAM.

NVIDIA. (2024, March). *What Is Agentic AI?* https://blogs.nvidia.com/blog/what-is-agentic-ai/

O'Donoghue, B., Chu, E., Parikh, N., & Boyd, S. (2016). *Conic optimization via operator splitting and homogeneous self-dual embedding. Journal of Optimization Theory and Applications*, *169*(3), 1042–1068.

Pástor, L., & Veronesi, P. (2013). *Political uncertainty and risk premia. Journal of financial Economics*, *110*(3), 520–545.

Plakandaras, V., Gogas, P., & Papadimitriou, T. (2018). *The effects of geopolitical uncertainty in forecasting financial markets: A machine learning approach. Algorithms*, *12*(1), 1.

Rosenberg, A., Stagnol, L., & Sekine, T. (2024). *Geopolitical risk will grow: here is how we track it.*

Rosenberg, B. (1974). *Extra-market components of covariance in security returns. Journal of Financial and Quantitative Analysis*, *9*(2), 263–274.

Rosner, B., Glynn, R. J., & Lee, M.-L. T. (2006). *The Wilcoxon signed rank test for paired comparisons of clustered data. Biometrics*, *62*(1), 185–192.

Shiller, R. J. (2017). *Narrative Economics. American Economic Review*, *107*(4), pp. 967–1004.

Shiller, R. J. (2020). Narrative Economics: How Stories Go Viral And Drive Major Economic Events. Princeton University Press.

Smales, L. A. (2021). *Geopolitical risk and volatility spillovers in oil and stock markets. The Quarterly Review of Economics and Finance*, *80*, 358–366.

Stellato, B., Banjac, G., Goulart, P., Bemporad, A., & Boyd, S. (2020). *OSQP: An operator splitting solver for quadratic programs. Mathematical Programming Computation*, *12*(4), 637–672.

Tilly, S., McDougall, A., Chaillou, T., Le Guenedal, T., Sakout, S., & Sekine, T. (2025). *Automating Insight Extraction from Oil and Gas Sector Climate Disclosures with AI.* https://doi.org/10.2139/ssrn.5556560

Traut, J. (2023). *What we know about the low-risk anomaly: a literature review. Financial markets and portfolio management*, *37*(3), 297–324.

Van Vliet, P., & Blitz, D. (2011). *Benchmarking low-volatility strategies. Journal of Index Investing*, *2*(1), 44–49.

Yang, X., Liu, Z., & Chen, X. (2025). *Conditional Multifactor Volatility-Managed Portfolios: Evidence from Global Market. Proceedings of the 2nd Guangdong-Hong Kong-Macao Greater Bay Area International Conference on Digital Economy and Artificial Intelligence*, 227–232.

Zaremba, A. (2016). *Is there a low-risk anomaly across countries? Eurasian Economic Review*, *6*(1), 45–65.

Zhao, T., Lyu, J., Jones, S., Garber, H., Pasquali, S., & Mehta, D. (2025). *AlphaAgents: Large Language Model based Multi-Agents for Equity Portfolio Constructions.* https://arxiv.org/abs/2508.11152

Chief Editor

Monica DEFEND
Head of Amundi Investment Institute

Editors

Marie BRIÈRE
Head of Investors' Intelligence & Academic Partnership

Thierry RONCALLI
Head of Quant Portfolio Strategy

Investment Institute

WORKING PAPER 186 | **MARCH 2026**

Find out more about Amundi Investment Institute Publications

Visit our Research Center

SCAN ME

Amundi
Investment Solutions

Trust must be earned