From Black Box to Explainable Portfolio Optimization: Tracing Allocations to Views and Constraints

J. Landais*, R. Perchet†, F. Soupé‡ and Raul Leote de Carvalho§

* Quantitative analyst in Multi-Asset team of the Quant Research Group at BNP Paribas Asset Management
† Head of Multi-Asset team of the Quant Research Group at BNP Paribas Asset Management
‡ Co-head of the Quant Research Group at BNP Paribas Asset Management
§ Deputy head of the Quant Research Group at BNP Paribas Asset Management

February 13th 2026

Keywords: Robust optimization; Portfolio construction; Mean-variance optimization; Multi-asset; Asset Allocation

JEL Classification: G11, C61

† Author for correspondence.

From Black Box to Explainable Portfolio Optimization: Tracing Allocations to Views and Constraints

J. Landais[*], R. Perchet[†], F. Soupé[‡] and Raul Leote de Carvalho[§]

* Quantitative analyst in Multi-Asset team of the Quant Research Group at BNP Paribas Asset Management
† Head of Multi-Asset team of the Quant Research Group at BNP Paribas Asset Management
‡ Co-head of the Quant Research Group at BNP Paribas Asset Management
§ Deputy head of the Quant Research Group at BNP Paribas Asset Management

February 13th 2026

Abstract: This paper introduces a comprehensive framework to make portfolio optimization fully transparent and explainable in the context of tactical asset allocation. The method breaks down the optimized multi-asset portfolio into a sum of intuitive sub-portfolios, each capturing a distinct driver of the final allocation. Specifically, the decomposition separates the contributions from: (i) a sub-portfolio that replicates the strategic allocation using the available funds, accounting for mismatches between the investable universe and the indices used for the strategic asset allocation; (ii) sub-portfolios with the tilts expressing the tactical allocation views; (iii) a sub-portfolio with tilts arising from expected fund alphas, net of fees; (iv) a set of sub-portfolios quantifying the impact of binding constraints; and (v) a funding sub-portfolio with the adjustment to ensure full investment. We introduce additional practical refinements to enhance the interpretability, robustness, and practical usability of the decomposition. These include a more intuitive normalization of the sub-portfolios and mechanisms for redistributing extreme constraint effects.

Keywords: Robust optimization; Portfolio construction; Mean-variance optimization; Multi-asset; Asset Allocation

JEL Classification: G11, C61

I. Introduction

The construction of tactical asset allocation (TAA) portfolios in institutional investment practice presents a host of practical challenges that are often under-appreciated in academic literature. While strategic asset allocation (SAA) and TAA views are typically formulated in terms of traditional core indices representing broad asset classes, the actual implementation of TAA must be carried out using investment funds, often using a mix of active and passive funds. Not only may these funds have benchmarks that differ from the indices used for SAA and TAA views, but active funds also introduce additional sources of alpha and idiosyncratic risk. This mismatch between the benchmarks used to express investment views and evaluate performance, and the actual investable universe, complicates the translation of allocation decisions into portfolios that are not only implementable and robust, but also transparent and faithful to the original investment views.

Further complicating matters, TAA portfolios must adhere to a range of constraints, such as restrictions on leverage and short positions, as well as the requirement to remain fully invested. Additional complexity arises because investment views are usually not internally consistent with the correlation structures assumed in risk models, rendering traditional mean-variance optimization (MVO) approaches ineffective or even unusable for TAA construction. Moreover, tactical views are frequently difficult to express as precise numerical expected returns and may not be available for all assets in the investable universe. These realities underscore the need for portfolio construction frameworks that are robust to estimation errors and portfolio constraints, and that can accommodate incomplete, qualitative, or inconsistent investment views in a transparent and explainable manner.

In response to these challenges, this paper introduces a framework that brings full transparency to the portfolio optimization process by enabling an explicit attribution of each fund's allocation in the TAA portfolio to its underlying determinants: the SAA used for performance evaluation, the tactical investment views, the expected fund alphas adjusted for ongoing costs, and the binding constraints. We propose a rigorous analytical framework and provide linear decompositions that allow practitioners to systematically trace and interpret the origin of each weight in the final constrained TAA portfolio. Having detailed step-by-step explanations of how the optimizer operates and, specifically, how it constructs each portfolio weight from these building blocks, is crucial for establishing trust in the optimizer, especially among stakeholders with lower appetite for quantitative methods.

Beyond transparency, the proposed framework provides clear and tangible evidence of the notion of value for money. By decomposing each allocation into economically interpretable components and linking them to realized constraints and costs, the framework shows how client value is created (or eroded) at the margin, thereby offering verifiable proof of value for money in terms both of design and of explanation. This strengthens accountability and enables portfolio managers to clearly articulate and rigorously justify portfolio outcomes to stakeholders, regulators, and clients.

We begin by recalling the derivation of the basic framework in the context of MVO. We then extend the application of the framework to the case of robust portfolio optimization (RPO). Building on the foundational work of Lobo et al. (1998) and Ben-Tal and Nemirovski (1998), RPO addresses the issue of estimation error in portfolio inputs by seeking allocations that remain effective under worst-case deviations. Subsequent developments, including the two-step max–min formulation of Tütüncü and König (2004), the quadratic error structure of Ceria and Stubbs (2006), and the proportional error-covariance approach of Scherer (2006), have further enhanced the tractability and interpretability of robust solutions. More recently, Heckel et al. (2016) demonstrated across a variety of (RPO) formulations that, as uncertainty in expected returns increases, robust portfolios interpolate between the mean-variance optimal allocation and various risk-based allocations, with the specific risk-based approach depending on the RPO formulation used.

Recent contributions by Issaoui et al. (2021) have adapted RPO frameworks to the realities of SAA and TAA, providing practical guidelines for uncertainty calibration and the integration of qualitative investment views. These frameworks have been further extended by Somefun et al. (2022) to accommodate core-satellite portfolio approaches and thematic investments, and by Mallouli et al. (2025) to the construction of TAA portfolios using active and passive funds while accounting for ongoing fund charges and tracking error constraints. However, none of these works provides the analytical framework capable of decomposing the optimized portfolio into its underlying drivers, allowing each portfolio weight to be explained by contributions from the impact of the original tactical investment decisions, the expected fund alphas, the mismatch between core indices and funds used for implementation, and the myriad of portfolio constraints.

The main contribution of this paper is to extend the linear decomposition of the constrained mean–variance solution, where the optimized portfolio is expressed as the sum of the unconstrained solution and constraint-specific corrections, to the RPO setting proposed by Issaoui et al. (2021), and further adapted by Mallouli et al. (2025) for allocations across active and passive funds. We show that the decomposition remains valid when the covariance matrix is replaced by its robustified counterpart and, crucially, when optimization is anchored on implied returns derived from an unconstrained tactical portfolio, thereby accommodating qualitative TAA views. This extension renders state-of-the-art robust TAA optimization fully transparent and explainable.

We develop the decomposition in several stages from theory to practical use cases: first, we recapitulate the linear decomposition for MVO with constraints; next, we generalize this result to RPO, introducing the robustified covariance matrix; then, we extend the framework to the practical case where the optimization is based on implied returns derived from an unconstrained tactical portfolio built from a selection of qualitative tactical views on individual asset classes; we then address the case where the asset universe is split between investable and non-investable assets, reflecting the realities of implementation using funds which do not necessarily match the core indices used for SAA and TAA investment views; and finally, we propose the decomposition of the weights of the TAA portfolio into contributions from an SAA replication, the investments views, and the different portfolio constraints including the funding constraint, constraints on the maximum and minimum weights of funds in the final portfolio, and other typical linear constraints, e.g., on environmental, social and governance (ESG) scores.

We further refine the framework to enhance its practical utility, showing how to include a more useful normalization of contributions and to explicitly separate the contributions from investment views into contributions from tactical views and from expected fund alphas. Finally, we propose an enhancement of the approach showing how to dispatch the sub-portfolios from maximum and minimum weight constraints back to their underlying sources, which is useful when these constraints merely offset large unconstrained tilts. This adjustment ensures that the decomposition remains interpretable even in the presence of strong offsetting effects, preserving transparency in cases where binding constraints would otherwise obscure the true drivers of the optimized allocation.

In the Results section, we demonstrate the framework using a TAA case study benchmarked to an index-based SAA. Using real-world funds, we show how the optimized portfolio can be broken down into an exact sum of intuitive sub-portfolios, each reflecting a distinct investment rationale as described above. We also examine the impact of imposing a minimum allocation to Sustainable Investments, highlighting how the decomposition makes the constraint's effect explicit while maintaining the transparency of the overall optimization process.

By providing a transparent, robust, and analytically tractable decomposition of TAA portfolios, our framework bridges the gap between quantitative optimization and the practical demands of professional asset managers and investors. By applying it to the most advanced approaches to TAA portfolio optimization, it enables practitioners and investors to construct portfolios that are not only robust to estimation errors and portfolio constraints, but also fully explainable in terms of the underlying investment views and the constraints that drive their construction.

II. Materials and methods

Consider an investment universe *A* with n_A financial assets divided into investable assets *I*, active and passive funds in our case, and non-investable assets *N*, i.e., core indices.

In the remainder of this document, *A* denotes vectors and matrices spanning the full universe of financial assets, investable and non-investable. All other vectors and matrices are restricted to investable assets only, i.e., the active and passive funds.

Given the optimal portfolio allocation to active and passive funds, $\boldsymbol{w}_{\text{TAA}}$, the objective of this paper is to come up with a linear decomposition:

$$\boldsymbol{w}_{\text{TAA}} = \sum_{i=1}^{n_{criteria}} \boldsymbol{w}_{\text{TAA},i} \quad (1)$$

where each sub-portfolio $\boldsymbol{w}_{\mathrm{TAA},i}$ captures the impact of each criterion i used in the portfolio construction problem. The sum of all sub-portfolios equals exactly the optimized portfolio $\boldsymbol{w}_{\mathrm{TAA}}$. Hence, this decomposition should enable a transparent attribution of the final portfolio weights to the underlying criteria used.

In our TAA problem, we start from a SAA portfolio defined on non-investable core indices, and we express tactical views on a subset of these indices, reflecting the convictions from an investment committee. However, the final TAA portfolio must be implemented using only funds from a pre-selected list of passive and active funds, where the active funds are chosen specifically for their expected positive alpha. Our objective here is to construct a TAA portfolio that faithfully integrates the SAA, the tactical views on core indices, the expected alphas of the selected funds, and the impact of all portfolio constraints, including weight bounds, linear constraints such as ESG score requirements, and the full-investment constraint.

In Table 1, we summarize the type of output we seek. This table consolidates the results illustrated later in Table 7. For each fund, the final TAA portfolio weight is exactly decomposed into the sum of its weights across several sub-portfolios: (i) replicating the SAA using the available funds (SAA Min TE); (ii) reflecting the tactical allocation views (Views); (iii) tilting towards funds with higher expected alpha after fees (Alpha); (iv) capturing the impact of constraints such as sustainability requirements (Constraints); and (v) ensuring that the final portfolio remains fully invested (Funding). In the sections that follow, we build this decomposition step by step, ultimately relying on an RPO framework.

Table 1: Target output decomposition of a given TAA portfolio

	Portfolio Weights					
	TAA	SAA Min TE	Views	Alpha	Constraints	Funding
					30% Minimum Allocation to Sustainable Investments	
Equity Europe Mid-large Active Fundamental	2.1%	2.7%	1.0%	1.5%	-3.2%	0.1%
Equity Europe Mid-large Passive Index	25.9%	17.5%	6.2%	-1.6%	3.8%	0.0%
Equity USA Growth Active Fundamental	6.4%	4.7%	0.0%	0.6%	0.7%	0.3%
Equity USA Mid-large Passive Index	4.8%	8.0%	0.0%	-0.6%	-2.3%	-0.2%
Equity Japan Mid-large Active Fundamental	5.5%	4.7%	0.6%	0.6%	-0.2%	-0.2%
Equity Japan Mid-large Passive Index	6.1%	5.4%	0.7%	-0.5%	0.7%	-0.2%
Equity Emerging Mid-large Active Fundamental	4.4%	2.9%	0.9%	0.7%	-0.1%	0.0%
Equity Emerging Mid-large Passive Index	5.4%	4.5%	1.6%	-0.6%	-0.2%	0.1%
Bonds Global Aggregate Active Fundamental	0.0%	18.4%	-12.1%	0.6%	-3.8%	-3.1%
Bonds EUR IG Active Fundamental	0.0%	4.6%	0.0%	3.3%	-6.8%	-1.2%
Bonds EUR IG Passive Index	28.4%	11.3%	11.9%	-3.6%	10.3%	-1.5%
Bonds USD IG Passive Index	11.0%	15.4%	-5.2%	-0.4%	1.1%	0.2%
Portfolio Weight Sum	100.0%	100.0%	5.7%	0.0%	0.0%	-5.7%

Notes: IG: Investment Grade. Sustainable Investment allocation is calculated from each fund's minimum exposure in Table A2. The results are based on the example in Table 8.

II.A. Linear decomposition of mean-variance optimization solution

The MVO tactical active portfolio $\boldsymbol{a} = \boldsymbol{w}_{\text{TAA}} - \boldsymbol{w}_{\text{SAA}}$ can be found by solving:

$$\max_{\boldsymbol{a}} \left(\boldsymbol{\mu}^\top \boldsymbol{a} - \lambda(\boldsymbol{a}^\top \boldsymbol{\Sigma} \boldsymbol{a})\right) \tag{2}$$

with $\boldsymbol{\mu}^\top \boldsymbol{a}$ the expected active return, $\sqrt{\boldsymbol{a}^\top \boldsymbol{\Sigma} \boldsymbol{a}}$ the tracking error and given $\lambda > 0$ the risk aversion. The unconstrained solution to (2) is well known and given by:

$$\boldsymbol{a}^{(0)} = \frac{1}{2\lambda} \boldsymbol{\Sigma}^{-1} \boldsymbol{\mu} \tag{3}$$

Now let us see the impact of adding multiple linear constraints:

$$\boldsymbol{B}^\top \boldsymbol{a} = b \tag{4}$$

where $\boldsymbol{B}$ is an $n_A \times m$ matrix with the constraint vectors in columns and b is an m vector with the required matching values.

The zero-sum constraint arising from the fact that both $\boldsymbol{w}_{\text{TAA}}$ and $\boldsymbol{w}_{\text{SAA}}$ are fully invested portfolios can be written as one of such constraints by taking one column of $\boldsymbol{B}$ to be $\mathbf{1}$ (the n_A-sized vector with all coefficients set to 1 and the corresponding $b = 0$. Examples of constraints that fit this form include not only this zero-sum constraint, $\mathbf{1}^\top \boldsymbol{a} = 0$, but other typical constraints such as imposing country $\boldsymbol{c}^\top \boldsymbol{a} = 0$ or sector, $\boldsymbol{s}^\top \boldsymbol{a} = 0$, neutrality, or targeting a given active exposure on some characteristic $\boldsymbol{q}$ of the underlying assets $\boldsymbol{q}^\top \boldsymbol{a} = q_0$.

With the vector of Lagrange multipliers $\boldsymbol{\delta} = [\delta_1, \dots, \delta_m]^\top$, the Lagrangian for all these constraints:

$$\mathcal{L}(\boldsymbol{a}, \boldsymbol{\delta}) = \boldsymbol{\mu}^\top \boldsymbol{a} - \lambda \boldsymbol{a}^\top \boldsymbol{\Sigma} \boldsymbol{a} - \boldsymbol{\delta}^\top (\boldsymbol{B}^\top \boldsymbol{a} - \boldsymbol{b}) \tag{5}$$

Applying first order conditions in $\boldsymbol{a}$ results in:

$$\boldsymbol{\mu} - 2\lambda \boldsymbol{\Sigma} \boldsymbol{a} - \boldsymbol{B} \boldsymbol{\delta} = 0 \tag{6}$$

which when solved for $\boldsymbol{a}$ produces a linear decomposition of the portfolio into the sum of the unconstrained portfolio and terms with the Lagrange multipliers (Grinold and Khan (2000), Boyd and Vandenberghe (2004), Meucci (2005)):

$$\begin{aligned} \boldsymbol{a} &= \frac{1}{2\lambda} \boldsymbol{\Sigma}^{-1} (\boldsymbol{\mu} - \boldsymbol{B} \boldsymbol{\delta}) \\ &= \boldsymbol{a}^{(0)} - \frac{1}{2\lambda} \boldsymbol{\Sigma}^{-1} \boldsymbol{B} \boldsymbol{\delta} \end{aligned} \tag{7}$$

enforcing the constraints $\boldsymbol{B}^\top \boldsymbol{a} = \boldsymbol{b}$:

$$\boldsymbol{B}^\top \boldsymbol{a}^{(0)} - \frac{1}{2\lambda} \boldsymbol{B}^\top \boldsymbol{\Sigma}^{-1} \boldsymbol{B} \boldsymbol{\delta} = \boldsymbol{b} \tag{8}$$

and then solving for $\boldsymbol{\delta}$ leads to:

$$\boldsymbol{\delta} = 2\lambda\, (\boldsymbol{B}^\top \boldsymbol{\Sigma}^{-1} \boldsymbol{B})^{-1} (\boldsymbol{B}^\top \boldsymbol{a}^{(0)} - \boldsymbol{b}) \tag{9}$$

which when plugged back into (7) results in:

$$\boldsymbol{a}^\star = \boldsymbol{a}^{(0)} - \boldsymbol{\Sigma}^{-1} \boldsymbol{B}\, (\boldsymbol{B}^\top \boldsymbol{\Sigma}^{-1} \boldsymbol{B})^{-1} \left(\boldsymbol{B}^\top \boldsymbol{a}^{(0)} - \boldsymbol{b}\right) \tag{10}$$

If we write the constraint matrix in column format as a collection of vectors $\boldsymbol{B} = [\boldsymbol{B}_1, \dots, \boldsymbol{B}_m]$ then $\boldsymbol{B}^\top \boldsymbol{a} = \boldsymbol{b}$ is equivalent to $\boldsymbol{B}_k^\top \boldsymbol{a} = \boldsymbol{b}_k$ and applying the Karush-Kuhn-Tucker (KKT) conditions allows us to re-write (7) as:

$$\boldsymbol{a}^\star = \boldsymbol{a}^{(0)} - \frac{1}{2\lambda} \boldsymbol{\Sigma}^{-1} \sum_{k=1}^{m} \boldsymbol{B}_k \delta_k \tag{11}$$

which is an explicit sum of the unconstrained portfolio and one correction term per constraint:

$$\boldsymbol{a}^\star = \boldsymbol{a}^{(0)} - \sum_{k=1}^{m} \frac{1}{2\lambda} \boldsymbol{\Sigma}^{-1} \boldsymbol{B}_k \delta_k \tag{12}$$

with the multipliers δ_k jointly determined by the linear equation:

$$(\boldsymbol{B}^\top \boldsymbol{\Sigma}^{-1} \boldsymbol{B})\boldsymbol{\delta} = 2\lambda(\boldsymbol{B}^\top \boldsymbol{a}^{(0)} - \boldsymbol{b}) \tag{13}$$

In the Appendix C we extend the framework to include turnover penalties.

II.B. Linear decomposition using robust portfolio optimization solution

Here we repeat the exercise for the more general case of RPO formulated as (e.g. Ceria and Stubbs (2006)):

$$\max_{\boldsymbol{a}} \left(\boldsymbol{\mu}^\top \boldsymbol{a} \;-\; \lambda\, \boldsymbol{a}^\top \boldsymbol{\Sigma} \boldsymbol{a} \;-\; \kappa \sqrt{\boldsymbol{a}^\top \boldsymbol{\Omega} \boldsymbol{a}}\right) \tag{14}$$

subject to constraints $\boldsymbol{B}^\top \boldsymbol{a} = \boldsymbol{b}$ as before and where $\boldsymbol{\Omega}$ is the uncertainty matrix of returns and κ is the aversion to uncertainty in returns.

Common choices for the uncertainty matrix include specifications proportional to the identity matrix, to a diagonal matrix built from estimated asset variances, or to the full variance–covariance matrix $\boldsymbol{\Sigma}$. Heckel et al. (2016) analyzed the properties of the optimal solution to equation (14) under both low- and high-uncertainty regimes for these different forms of $\boldsymbol{\Omega}$. More recently, Yin et al. (2020) and Mallouli et al. (2025) provided theoretical background and empirical justification for why an uncertainty matrix proportional to the diagonal matrix of estimated asset variances is often the most appropriate choice for standard portfolio optimization problems.

The presence of the square-root term makes the problem non-quadratic, so it is no longer possible to proceed as above. However, we can recover a similar algebraic structure by introducing one scalar auxiliary variable to make the problem quadratic again, with a modified covariance, conditional on that scalar (Boyd and Vandenberghe (2004), and Ben-Tal, El Ghaoui and Nemirovski (2009)).

For that, we can use the identity which is valid for any $q \geq 0$:

$$\sqrt{q} = \min_{t>0} \left(\frac{q}{2t} + \frac{t}{2}\right) \tag{15}$$

applying it to the uncertainty term:

$$-\kappa\sqrt{\boldsymbol{a}^\top \boldsymbol{\Omega} \boldsymbol{a}} = \max_{t>0} \left[-\frac{\kappa}{2}\left(\frac{\boldsymbol{a}^\top \boldsymbol{\Omega} \boldsymbol{a}}{t} + t\right)\right] \tag{16}$$

Therefore, the RPO problem is equivalent to the joint maximization:

$$\max_{a,\, t>0} \left(\boldsymbol{\mu}^\top \boldsymbol{a} - \lambda\, \boldsymbol{a}^\top \boldsymbol{\Sigma} \boldsymbol{a} - \frac{\kappa}{2t}\, \boldsymbol{a}^\top \boldsymbol{\Omega} \boldsymbol{a} - \frac{\kappa}{2}\, t\right) \tag{17}$$

subject to constraints $\boldsymbol{B}^\top \boldsymbol{a} = \boldsymbol{b}$. At the optimum, t will satisfy the condition $t^\star = \sqrt{(\boldsymbol{a}^\star)^\top \boldsymbol{\Omega} \boldsymbol{a}^\star}$.

Because the gradient of $-\lambda \boldsymbol{a}^\top \Sigma \boldsymbol{a}$ is $-2\gamma \boldsymbol{\Sigma} \boldsymbol{a}$, and the gradient of $-\frac{\kappa}{2t} \boldsymbol{a}^\top \boldsymbol{\Omega} \boldsymbol{a}$ is $-\frac{\kappa}{t} \boldsymbol{\Omega} \boldsymbol{a}$, we can now introduce the modified variance covariance matrix defined for $t > 0$:

$$\boldsymbol{Q}(t) = 2\lambda\, \boldsymbol{\Sigma} + \frac{\kappa}{t}\, \boldsymbol{\Omega} \tag{18}$$

Then, dropping the constant $-\frac{\kappa}{2} t$ with respect to $\boldsymbol{a}$ in (17), the conditional problem becomes:

$$\max_{\boldsymbol{a}} \left(\boldsymbol{\mu}^\top \boldsymbol{a} - \frac{1}{2} \boldsymbol{a}^\top \boldsymbol{Q}(t)\, \boldsymbol{a}\right) \tag{19}$$

subject to the constraints $\boldsymbol{B}^\top \boldsymbol{a} = \boldsymbol{b}$. With this, we recover the same structure as for the original mean–variance constrained problem, but now with $\boldsymbol{\Sigma}$ replaced by $\boldsymbol{Q}(t)$, up to a scaling.

For a given fixed $t > 0$, and applying first order constraints, the unconstrained robust portfolio satisfies:

$$\boldsymbol{a}^{(0)}(t) = \boldsymbol{Q}(t)^{-1} \boldsymbol{\mu} \tag{20}$$

and the solution that satisfies the constraints $\boldsymbol{B}^\top \boldsymbol{a} = \boldsymbol{b}$ can be written in a form equivalent to (10):

$$\boldsymbol{a}^\star(t) = \boldsymbol{a}^{(0)}(t) - \boldsymbol{Q}(t)^{-1} \boldsymbol{B}\, (\boldsymbol{B}^\top \boldsymbol{Q}(t)^{-1} \boldsymbol{B})^{-1} \left(\boldsymbol{B}^\top \boldsymbol{a}^{(0)}(t) - \boldsymbol{b}\right) \tag{21}$$

As before, using the constraint matrix in column format $\boldsymbol{B} = [\boldsymbol{B}_1, \dots, \boldsymbol{B}_m]$ and the Lagrange multipliers $\boldsymbol{\delta} = [\delta_1, \dots, \delta_m]^\top$, we recover an equation equivalent to (12):

$$\boldsymbol{a}^\star(t) = \boldsymbol{a}^{(0)}(t) - \sum_{k=1}^{m} (\boldsymbol{Q}(t)^{-1} \boldsymbol{B}_k \delta_k(t)) \tag{22}$$

with a similar linear decomposition as before but with $\boldsymbol{Q}(t)$ replacing $\boldsymbol{\Sigma}$, up to a scaling, and with the multipliers δ_k now jointly determined by the linear equation:

$$(\boldsymbol{B}^\top \boldsymbol{Q}(t)^{-1}\boldsymbol{B})\boldsymbol{\delta} = (\boldsymbol{B}^\top \boldsymbol{a}^{(0)} - \boldsymbol{b}) \tag{23}$$

In the Appendix C we show how to add turnover penalties to this framework.

II.C. Linear decomposition when optimizing from tactical allocation implied returns

Practitioners often anchor their process on a pre-defined TAA portfolio because even the unconstrained MVO solution based on expected returns is often highly sensitive to correlation estimates. Small changes in pair-wise correlations of asset returns can lead to extremely large changes in the overall allocations, and producing accurate enough expected returns is simply not possible. While RPO can reduce correlation sensitivity, it does not solve the deeper challenge of expressing views in the form of a numerical return forecast, particularly for assets outside the practitioner's scope of analysis. Consequently, starting from a pre-selected unconstrained TAA portfolio often remains the preferred approach, even under RPO, as it provides a practical and robust foundation that reflects strategic convictions without requiring explicit return estimates for all assets. In such a case, optimization is still useful and required to change the portfolio so that it meets all required constraints.

Issaoui et al. (2021) introduced a methodology for constructing unconstrained TAA portfolios based on the observation that investment committees (IC) across the industry typically express tactical views in terms of the anticipated direction and strength of bets for the assets under consideration:

$$\boldsymbol{S}_{\text{directional}} = (S^1_{\text{directional}}, \dots, S^{n_A}_{\text{directional}})^\top \tag{24}$$

A straightforward approach to constructing an unconstrained active TAA portfolio is to base allocations directly on the directional scores that reflect the views of the investment committee. Issaoui et al. (2021) proposes that such an unconstrained TAA portfolio assigns active weights according to a risk budgeting methodology, ensuring that each view is represented proportionally to its conviction and associated risk:

$$\boldsymbol{a}^{(IC)} = \boldsymbol{S}_{\text{directional}} \times (RB\,\boldsymbol{\sigma}^{-1}) \tag{25}$$

where $\boldsymbol{\sigma} = (\sigma_1, \dots, \sigma_{n_A})^\top$ is the vector of asset volatilities and RB denotes the total risk budget. This formulation ensures that each directional view allocates a portion of the portfolio's tracking error proportional to its conviction and direction, according to:

$$RB \cdot S_i = a_i^{IC} \cdot \sigma_i \tag{26}$$

If, in Section II.A, we replace $\boldsymbol{\mu}$ with the implied returns $\bar{\boldsymbol{\mu}} = 2\lambda\boldsymbol{\Sigma}\boldsymbol{a}^{(IC)}$ obtained from the unconstrained TAA portfolio in (25), then the first component, $\boldsymbol{a}^{(0)}$, in equations (10) and (11) is essentially a scaled representation of the original TAA active portfolio $\boldsymbol{a}^{(IC)}$. The second term introduces the deviations required to enforce the specified constraints.

Conversely, in Section II.B, replacing $\boldsymbol{\mu}$ with the pseudo-implied returns $\bar{\boldsymbol{\mu}} = \boldsymbol{Q}\boldsymbol{a}^{(IC)}$, as suggested by Issaoui et al. (2021) and with $\boldsymbol{Q}$ defined in (18), leads to a similar interpretation in equations (21) and (22). The initial term again reflects a proportional version of the starting TAA portfolio, while the subsequent correction term incorporates the necessary tilts to satisfy the imposed restrictions.

II.D. Linear decomposition for practical application with tactical allocation to funds

In practical applications, the TAA portfolio must be fully invested in a selection of passive and/or active funds, whereas both the SAA portfolio and the tactical investment views are typically formulated using a universe of standard core asset class indices, many of which may not be directly investable or may not even correspond exactly to the benchmarks of the selected funds. Moreover, the selection of active funds is typically carried out by a team that is independent from the investment committee. These funds are added to the investable universe based on their expected alpha, regardless of the tactical views expressed by the committee. As a result, the inclusion of active funds is driven by their potential to generate excess returns, rather than alignment with specific tactical positions.

This context creates several challenges for portfolio construction. The first is managing the mismatch between the investable funds and the core indices used for SAA and investment views. In practice, the goal is to build a portfolio invested exclusively in funds that most faithfully represent the intended tactical views, while simultaneously satisfying all required constraints and controlling tracking error risk relative to the SAA portfolio.

A second challenge is achieving the right balance between passive and active funds. This involves accounting for the expected alpha generated by active funds, net of ongoing costs, and weighing this against the expected excess returns implied by tactical investment views. Addressing this trade-off is essential for constructing a portfolio that considers both sources of potential outperformance. This issue was recently explored by Mallouli et al. (2025) in the context of RPO. Here, we adopt their formulation to show how the linear decomposition framework can be adapted to enhance transparency in a real-world fund allocation application.

In particular, we emphasize the attribution of deviations in the final constrained TAA portfolio to four main sources: (i) the mismatch between the investable fund universe and the asset class indices used for both the SAA and to express tactical views; (ii) the actual tactical investment views, with a clear separation between contributions from tactical views and from active fund expected alpha; (iii) linear constraints, such as those imposing maximum and minimum allocations to each fund; and (iv) the funding constraint, which ensures that the TAA portfolio remains fully invested.

Consider an investment universe of n_{tot} assets. Let N be the subset of the non-investable assets, i.e., the asset class indices used for SAA and for expressing tactical investment views. Let I be the subset of investable assets, which is the selection of investable funds, including active and passive funds. Let the A exponent in vectors and matrices below refer to the full investment universe with all the I investable assets and all the N non-investable assets. Matrices and vectors below without this exponent span only the selected investable funds in I.

With these definitions, let $\boldsymbol{w}_{\text{SAA}}^{A} = \begin{bmatrix} \boldsymbol{w}_{\text{SAA}} \\ \mathbf{0} \end{bmatrix}$ be a fully invested SAA portfolio that allocates only to non-investable assets, i.e., with zero allocation to the investable funds. The construction of a tactical portfolio $\boldsymbol{w}_{\text{TAA}}^{A} = \begin{bmatrix} \mathbf{0} \\ \boldsymbol{w}_{\text{TAA}} \end{bmatrix}$, composed only of investable funds and with zero allocation to all non-investable assets, and with $\boldsymbol{a}^{\boldsymbol{A}} = \boldsymbol{w}_{\text{TAA}}^{A} - \boldsymbol{w}_{\text{SAA}}^{A}$, can be formulated as the following RPO problem where:

$$\max_{\boldsymbol{w}_{\text{TAA}}^{A}} (\bar{\boldsymbol{\mu}}^{A})^{\top}\boldsymbol{a}^{A} - \lambda(\boldsymbol{a}^{A})^{\top}\boldsymbol{\Sigma}^{A}\boldsymbol{a}^{A} - \kappa\sqrt{(\boldsymbol{a}^{A})^{\top}\boldsymbol{\Omega}^{A}\boldsymbol{a}^{A}} \tag{27}$$

subject to the linear constraints:

$$\mathbf{1}^{\top}\boldsymbol{w}_{\text{TAA}}^{A} = 1 \tag{28}$$

$$w_{\text{TAA},i}^{A} \geq w_{min,i}^{A}, i = 1, \dots, n_A \tag{29}$$

$$w_{\text{TAA},i}^{A} \leq w_{max,i}^{A}, i = 1, \dots, n_A \tag{30}$$

$$w_{\text{TAA},i}^{A} = 0 \text{ if } i \in N \tag{31}$$

$$\left(\boldsymbol{w}_{\text{TAA}}^{A}\right)^{\top}\boldsymbol{O}^{A,j} \geq O_{min,j}^{A}, j = 1, \dots, n_C \tag{32}$$

$$\left(\boldsymbol{w}_{\text{TAA}}^{A}\right)^{\top}\boldsymbol{O}^{A,j} \leq O_{max,j}^{A}, j = 1, \dots, n_C \tag{33}$$

where (28) is the funding constraint, (29) and (30) set the minimum $w_{min,i}^{A}$ and maximum and $w_{max,i}^{A}$ weights allowed for the asset i, and (31) are the constraints that restrict the allocation of the portfolio to selected funds only. $\boldsymbol{O}^{A,j}$ is the vector with the asset values of a given characteristic j used to create linear constraints and $\left(\boldsymbol{w}_{\text{TAA}}^{A}\right)^{\top}\boldsymbol{O}^{A,j}$ represents the value of the portfolio for that characteristic while $O_{min,j}^{A}$ and $O_{max,j}^{A}$ in equations (32) and (33) are the maximum and minimum values, respectively, that bound the portfolio value.

As proposed by Mallouli et al. (2025), the implied returns spanning all assets, $\bar{\boldsymbol{\mu}}^{A}$, are calculated so as to render the tactical active portfolio selected by the investment committee $\boldsymbol{a}^{(IC)}$ efficient under (27) plus the independent expected fund alphas $\boldsymbol{\alpha}_{fb}$ net of ongoing costs:

$$\bar{\boldsymbol{\mu}}^{A} = 2\lambda\boldsymbol{\Sigma}^{A}\boldsymbol{a}^{A(IC)} + \kappa\frac{\boldsymbol{\Omega}^{A}\boldsymbol{a}^{A(IC)}}{\sqrt{\left(\boldsymbol{a}^{A(IC)}\right)^{\top}\boldsymbol{\Omega}\boldsymbol{a}^{A(IC)}}} + \frac{1}{\gamma}\boldsymbol{\alpha}_{fb}^{A} \tag{34}$$

where $\boldsymbol{a}^{A(IC)} = \begin{bmatrix}\boldsymbol{a}^{(IC)} \\ \mathbf{0}\end{bmatrix}$ is a vector spanning all assets with the tactical portfolio from the investment committee, $\boldsymbol{\alpha}_{fb}^{A} = \begin{bmatrix}\mathbf{0} \\ \boldsymbol{\alpha}_{fb}\end{bmatrix}$ is a vector spanning all assets with the expected net alphas for each fund relative to their respective benchmark, and $1/\gamma$ can be interpreted as an overall confidence in the expected fund alphas.

In the appendix we provide more information about the choice of covariance matrix $\boldsymbol{\Sigma}^{A}$ and the uncertainty matrix $\boldsymbol{\Omega}^{A}$ spanning all non-investable assets in N and all funds in I, as well as the choices of risk aversion λ, aversion to uncertainty κ and confidence in fund alphas γ.

Since the weights of the non-investable assets in $\boldsymbol{w}_{\text{TAA}}^{A}$ are set to zero, we can re-write (27) as:

$$\max_{\boldsymbol{w}_{\text{TAA}}} (\bar{\boldsymbol{\mu}} + 2\lambda\boldsymbol{\Sigma}_{I}^{A}\boldsymbol{w}_{\text{SAA}}^{A})^{\top}\boldsymbol{w}_{\text{TAA}} - \lambda\boldsymbol{w}_{\text{TAA}}^{\top}\boldsymbol{\Sigma}\boldsymbol{w}_{\text{TAA}} - \kappa\sqrt{(\boldsymbol{a}^{A})^{\top}\boldsymbol{\Omega}^{A}\boldsymbol{a}^{A}} \tag{35}$$

$$\boldsymbol{w}_{\text{TAA}}{}^{\top}\mathbf{1} = 1 \tag{36}$$

$$w_{\text{TAA},i} \geq w_{min,i} \text{ , } i = 1, \dots, n_I \tag{37}$$

$$w_{\text{TAA},i} \leq w_{max,i} \text{ , } i = 1, \dots, n_I \tag{38}$$

$$\boldsymbol{w}_{\text{TAA}}{}^{\top}\boldsymbol{O}^j \geq O_{min,j} \text{ , } j = 1, \dots, n_C \tag{39}$$

$$\boldsymbol{w}_{\text{TAA}}{}^{\top}\boldsymbol{O}^{\boldsymbol{j}} \leq O_{max,j} \text{ , } j = 1, \dots, n_C \tag{40}$$

where n_I is the number of funds in I. Note that the returns vector is $\overline{\boldsymbol{\mu}} + 2\lambda\boldsymbol{\Sigma}_I^A\boldsymbol{w}_{\text{SAA}}^A$ where $\boldsymbol{\Sigma}_I^A$ is the submatrix of $\boldsymbol{\Sigma}^A$ obtained by removing the rows corresponding to non-investable assets. By refining the problem, we have added the term $2\lambda\boldsymbol{\Sigma}_I^A\boldsymbol{w}_{\text{SAA}}^A$ to the returns vector, which accounts for how well the investable assets track the SAA allocation.

With $\boldsymbol{\Omega}_I^A$ the submatrix of $\boldsymbol{\Omega}^A$ obtained by removing the rows corresponding to non-investable assets, and noting that $\boldsymbol{\Sigma}$ is a positive definite symmetric matrix and $\boldsymbol{\Omega}$ is at least positive semidefinite and symmetric, the KKT theorem ensures the existence of a unique solution and provides the following conditions:

- Stationarity:

$$2\lambda\boldsymbol{\Sigma}\boldsymbol{w}_{\text{TAA}} + \kappa\frac{\boldsymbol{\Omega}_I^A\boldsymbol{a}^A}{\sqrt{(\boldsymbol{a}^A)^{\top}\boldsymbol{\Omega}^A\boldsymbol{a}^A}} - \overline{\boldsymbol{\mu}} - 2\lambda\boldsymbol{\Sigma}_I^A\boldsymbol{w}_{\text{SAA}}^A - \sum_{i=1}^{n_I}\left(\delta_{min,i} - \delta_{max,i}\right)\boldsymbol{e}_i - \sum_{j=1}^{n_c}(\varepsilon_{min,j} - \varepsilon_{max,j})\boldsymbol{O}^j + \delta_{funding}\boldsymbol{1} = 0 \tag{42}$$

  where $\delta_{funding}$ is the Lagrangian multiplier associated with the funding constraint (35), $\delta_{min,i}$ and $\delta_{max,i}$ are the Lagrangian multipliers associated with the minimum and maximum constraints applied to fund i, and $\boldsymbol{e}_i$ and is the vector with zero everywhere except for the fund i which equals 1. $\varepsilon_{min,j}$ and $\varepsilon_{max,j}$ are the Lagrangian multipliers associated with the linear constraint on $\boldsymbol{O}^{\boldsymbol{j}}$.

- Primal feasibility:
  $\boldsymbol{1}^{\top}\boldsymbol{w}_{\text{TAA}} = 1$,
  $\boldsymbol{w}_{min} \leq \boldsymbol{w}_{\text{TAA}} \leq \boldsymbol{w}_{max}$,
  $O_{min,j} \leq \boldsymbol{w}_{\text{TAA}}{}^{\top}\boldsymbol{O}^j \leq O_{max,j}$ for all $j = 1, \dots, n_C$.

- Dual feasibility:
  $\delta_{min,i} \geq 0, i = 1, \dots, n_I$,
  $\delta_{max,i} \geq 0, i = 1, \dots, n_I$,
  $\varepsilon_{min,j} \geq 0, j = 1, \dots, n_C$,
  $\varepsilon_{max,j} \geq 0, j = 1, \dots, n_C$.

- Complementary slackness:
  $\delta_{max,i}(w_{\text{TAA},i} - w_{max,i}) = 0, i = 1, \dots, n_I$,
  $\delta_{min,i}(w_{\text{TAA},i} - w_{min,i}) = 0, i = 1, \dots, n_I$,
  $\varepsilon_{min,j}\left(\boldsymbol{w}_{\text{TAA}}{}^{\top}\boldsymbol{O}^j - O_{min,j}\right) = 0, j = 1, \dots, n_C$,
  $\varepsilon_{max,j}\left(\boldsymbol{w}_{\text{TAA}}{}^{\top}\boldsymbol{O}^j - O_{max,j}\right) = 0, j = 1, \dots, n_C$.

By factoring the stationarity condition, we get:

$$\left[2\lambda\boldsymbol{\Sigma} + \kappa\frac{\boldsymbol{\Omega}}{\sqrt{(\boldsymbol{a}^A)^\top\boldsymbol{\Omega}^A\boldsymbol{a}^A}}\right]\boldsymbol{w}_{\text{TAA}} - \kappa\frac{\boldsymbol{\Omega}_I^A\boldsymbol{w}_{SAA}^A}{\sqrt{(\boldsymbol{a}^A)^\top\boldsymbol{\Omega}^A\boldsymbol{a}^A}} - \overline{\boldsymbol{\mu}} - 2\lambda\boldsymbol{\Sigma}_I^A\boldsymbol{w}_{\text{SAA}}^A - \sum_{i=1}^{n_I}\left(\delta_{min,i} - \delta_{max,i}\right)\boldsymbol{e}_i - \sum_{j=1}^{n_c}(\varepsilon_{min,j} - \varepsilon_{max,j})\boldsymbol{O}^j + \delta_{funding}\boldsymbol{1} = 0 \quad (43)$$

If we reintroduce the matrix $\boldsymbol{Q}^A$ defined in (18):

$$\boldsymbol{Q}^A = 2\lambda\boldsymbol{\Sigma}^A + \kappa\frac{\boldsymbol{\Omega}^A}{\sqrt{(\boldsymbol{a}^A)^\top\boldsymbol{\Omega}^A\boldsymbol{a}^A}} \quad (44)$$

Note that this matrix depends on the optimized weights of $\boldsymbol{w}_{\text{SAA}}^A$, hence the stationarity condition cannot give a closed-form solution. Still, and with $\boldsymbol{Q}_I^A$ the submatrix of $\boldsymbol{Q}^A$ obtained by removing the rows corresponding to non-investable assets, and with $\boldsymbol{Q}$ submatrix of $\boldsymbol{Q}^A$ with only columns and rows for the investable assets, then, by replacing (44) in (43), the optimal portfolio is:

$$\boldsymbol{w}_{\text{TAA}} = \boldsymbol{Q}^{-1}\left[\overline{\boldsymbol{\mu}} + \boldsymbol{Q}_I^A\boldsymbol{w}_{\text{SAA}}^A + \sum_{i=1}^{n_I}\left(\delta_{min,i} - \delta_{max,i}\right)\boldsymbol{e}_i + \sum_{j=1}^{n_c}(\varepsilon_{min,j} - \varepsilon_{max,j})\boldsymbol{O}^j - \delta_{funding}\boldsymbol{1}\right] \quad (45)$$

We thus recover the linear decomposition given in (22) for the example under consideration.

II.E. Interpretation of the linear decomposition of a tactical allocation to funds

Equation (45) allows us to decompose the allocation to funds in the $\boldsymbol{w}_{\text{TAA}}$ into a sum of portfolios:

$$\boldsymbol{w}_{\text{TAA}} = \boldsymbol{w}_{\text{SAA}\,replication} + \boldsymbol{w}_{returns} + \boldsymbol{w}_{constraints} + \boldsymbol{w}_{funding} \quad (46)$$

By introducing the following definitions:

$$\boldsymbol{w}_{\text{SAA}\,replication} = \boldsymbol{Q}^{-1}\boldsymbol{Q}_I^A\boldsymbol{w}_{\text{SAA}}^A \quad (47)$$

$$\boldsymbol{w}_{returns} = \boldsymbol{Q}^{-1}\overline{\boldsymbol{\mu}} \quad (48)$$

$$\boldsymbol{w}_{constraints} = \boldsymbol{Q}^{-1}\left[\sum_{i=1}^{n_I}\left(\delta_{min,i} - \delta_{max,i}\right)\boldsymbol{e}_i + \sum_{j=1}^{n_c}(\varepsilon_{min,j} - \varepsilon_{max,j})\boldsymbol{O}^j\right] \quad (49)$$

$$\boldsymbol{w}_{funding} = -\boldsymbol{Q}^{-1}\delta_{funding}\boldsymbol{1} \quad (50)$$

$\boldsymbol{w}_{\text{SAA}\,replication}$ can be interpreted as the minimum tracking error portfolio that replicates $\boldsymbol{w}_{\text{SAA}}^A$ constrained to invest only in funds and constructed under uncertainty, i.e., using this new matrix $\boldsymbol{Q}$ instead of the original full $\boldsymbol{\Sigma}$.

$\boldsymbol{w}_{returns}$ is the fully unconstrained portfolio that maximizes the Sharpe ratio based on the implied returns $\overline{\boldsymbol{\mu}}$, with $\boldsymbol{Q}$ replacing the original full $\boldsymbol{\Sigma}$. This portfolio tilts away from the allocation $\boldsymbol{w}_{\text{SAA}\,replication}$ to take advantage of the tactical investment views and alphas of the funds in $\overline{\boldsymbol{\mu}}$.

$\boldsymbol{Q}^{-1}\boldsymbol{e}_i$ is the minimum-variance portfolio that holds asset i with weight 1, and all other weights are chosen to minimize total variance (i.e., to hedge the risk of asset i as much as possible), while $\boldsymbol{Q}^{-1}\boldsymbol{O}^j$ is the minimum-variance portfolio that achieves a unit exposure to the characteristic $\boldsymbol{O}^j$. In this way, $\boldsymbol{w}_{constraints}$ is the sum of all such portfolios scaled to fulfil all max and min weight constraints and all other linear constraints on $\boldsymbol{O}^j$.

$\boldsymbol{w}_{funding}$ is the fully unconstrained minimum variance portfolio scaled by $\delta_{funding}$, added or removed so that the final portfolio will be forced to have weights adding to 1.

II.F. Towards a more practical linear decomposition

Equations (45) and (46) follow naturally from the framework in Sections II.A and II.B. However, it is possible to rearrange or regroup the terms in the decomposition, provided they still sum to the final allocation $\boldsymbol{w}_{\text{TAA}}^{A}$, if an alternative breakdown offers greater clarity or is better suited for interpreting the resulting portfolio. Here we propose some changes that, in our view, increase transparency.

II.F.1 Minimum tracking error portfolio for SAA

The first concerns the term $\boldsymbol{w}_{\text{SAA}\,replication}$ defined in equation (47). This portfolio is invested in the available funds and seeks to replicate the SAA allocation to non-investable core indices but still optimized using the matrix $\boldsymbol{Q}$ rather than the original full covariance matrix $\boldsymbol{\Sigma}$. However, as it is, this portfolio is not necessarily fully invested. Adding a constraint to ensure it becomes fully invested enhances the transparency of the decomposition by transforming it into a more standard minimum tracking error portfolio. To achieve this, we introduce the following new definition:

$$\begin{aligned}\boldsymbol{w}_{\text{SAA}\,minTE} &= \boldsymbol{Q}^{-1}[\boldsymbol{Q}_I^A\boldsymbol{w}_{\text{SAA}}^A + \delta_{\min TE}\boldsymbol{1}] \\ &= \boldsymbol{w}_{\text{SAA}\,replication} + \boldsymbol{Q}^{-1}\delta_{\min TE}\boldsymbol{1} \qquad (51)\end{aligned}$$

with $\boldsymbol{w}_{\text{SAA}\,minTE}$ spanning only investable assets and with the Lagrangian:

$$\delta_{minTE} = \frac{1-(\boldsymbol{1}^\top\boldsymbol{Q}^{-1}\boldsymbol{Q}_I^A\boldsymbol{w}_{\text{SAA}}^A)}{\boldsymbol{1}^\top\boldsymbol{Q}^{-1}\boldsymbol{1}} \qquad (52)$$

If we replace $\boldsymbol{w}_{\text{SAA}\,replication}$ with $\boldsymbol{w}_{\text{SAA}\,minTE}$ in decomposition (46), we must adjust the funding term in (50) by subtracting $\boldsymbol{Q}^{-1}\delta_{\min TE}\boldsymbol{1}$. This correction ensures that the overall sum remains consistent. The fully invested minimum tracking error portfolio serves as the new baseline for constructing the fund allocation in $\boldsymbol{w}_{\text{TAA}}$, while the remaining components in (46) introduce tilts that reflect investment views and constraints.

II.F.2 Contributions from investment views towards the allocation

Here we focus on how to change the term $\boldsymbol{w}_{returns}$ to render it more informative. This could be done in different ways depending on how investment views are formulated. In the example above where we use implied returns defined in (34), derived from a given unconstrained TAA active portfolio built from individual views on each asset class and corrected with the expected net alphas from funds, we can write those implied returns as:

$$\overline{\boldsymbol{\mu}}^A = \boldsymbol{Q}^A\boldsymbol{a}^{A(IC)} + \frac{1}{\gamma}\boldsymbol{\alpha}_{fb}^A \qquad (53)$$

We can decompose the first term into a sum of individual views by splitting the portfolio $\boldsymbol{a}^{A(IC)}$ into a sum of the individual asset weights allocated to each non-investable asset:

$$\boldsymbol{a}^{A(IC)} = \begin{bmatrix} a_1 \\ 0 \\ \vdots \\ 0 \\ 0 \\ 0 \\ \vdots \\ 0 \end{bmatrix} + \begin{bmatrix} 0 \\ a_2 \\ \vdots \\ 0 \\ 0 \\ 0 \\ \vdots \\ 0 \end{bmatrix} + \ldots + \begin{bmatrix} 0 \\ 0 \\ \vdots \\ a_N \\ 0 \\ 0 \\ \vdots \\ 0 \end{bmatrix} \tag{54}$$

which leads to:

$$\overline{\boldsymbol{\mu}}^A = \sum_{i=1}^{n_N} \overline{\boldsymbol{\mu}}^A_{views,i} + \frac{1}{\gamma}\boldsymbol{\alpha}^A_{fb} \tag{55}$$

where $\overline{\boldsymbol{\mu}}^A_{views,i}$ are the implied returns for each asset in N calculated from the matrix $\boldsymbol{Q}$ rather than the original full covariance matrix $\boldsymbol{\Sigma}$, and derived from each individual allocation to non-investable assets i in (54).

Because the final allocation is not allowed to be invested in the non-investable assets in N, we can drop the rows for such assets in (55) and use the bottom of the vectors for investable assets in I only in (48) for $\boldsymbol{w}_{returns}$ allowing us to decompose this portfolio into a sum of portfolios:

$$\begin{aligned} \boldsymbol{w}_{returns} &= \boldsymbol{Q}^{-1} \sum_{i=1}^{n_N} \overline{\boldsymbol{\mu}}_{views,i} + \frac{1}{\gamma}\boldsymbol{Q}^{-1}\boldsymbol{\alpha}_{fb} \\ &= \sum_{i=1}^{n_N} \boldsymbol{w}_{views,i} + \boldsymbol{w}_{\alpha} = \boldsymbol{w}_{views} + \boldsymbol{w}_{\alpha} \end{aligned} \tag{56}$$

The first term in (56) represents a sum of portfolios composed exclusively of investable assets. Each portfolio is constructed to replicate the target allocation to each non-investable asset in (55), derived from the investment views. This term gives insight into how each view on a non-investable asset can be replicated using investable assets, with the replication constructed by minimizing tracking error measured using the matrix $\boldsymbol{Q}$ instead of the original full covariance matrix $\boldsymbol{\Sigma}$.

The second term, $\boldsymbol{w}_{\alpha}$, corresponds to the allocation to investable assets that is optimal in the absence of any investment views. It is determined solely by the expected alphas of investable assets, adjusted for ongoing costs.

II.F.3 Normalization of contributions towards the allocation

In the problem at hand, $\boldsymbol{w}_{\text{TAA}}$ is also constrained to be fully invested, i.e. portfolio weights add to 100%. $\boldsymbol{w}_{\text{SAA}\,minTE}$ is also fully invested. In the absence of other constraints and with $\overline{\boldsymbol{\mu}}_{\boldsymbol{\alpha}} = \boldsymbol{0}$, the weights in $\boldsymbol{w}_{views}$ derived from the investment views have to be compensated by weights in $\boldsymbol{w}_{funding}$ so as to meet the funding constraints, i.e. that the final weights of $\boldsymbol{w}_{\text{TAA}}$ total 100%.

In this setting, $\boldsymbol{w}_{\text{TAA}}$ is constrained to be fully invested, meaning that the weights add to 100%. $\boldsymbol{w}_{\text{SAA}\,minTE}$ is also fully invested. When no additional constraints are imposed, and with $\boldsymbol{\alpha}_{fb} = 0$, the weights in $\boldsymbol{w}_{views}$, derived from the investment views, must be offset by weights in $\boldsymbol{w}_{funding}$ so that the final weights in $\boldsymbol{w}_{\mathrm{TAA}}$ sum to 100%.

To generalize to non-zero $\boldsymbol{\alpha}_{fb}$ and additional constraints, we introduce the definition of zero-sum portfolio $\widehat{\boldsymbol{w}}_j$ associated with portfolio $\boldsymbol{w}_j$:

$$\widehat{\boldsymbol{w}}_j = \boldsymbol{w}_j - \frac{\mathbf{1}^{\top}\boldsymbol{w}_j}{\mathbf{1}^{\top}\boldsymbol{w}_{funding}}\boldsymbol{w}_{funding} \tag{57}$$

where $\mathbf{1}^{\top}\boldsymbol{w}$ is the sum of weights of $\boldsymbol{w}$ and $\mathbf{1}^{\top}\boldsymbol{w}_{funding}$ the sum of weights of $\boldsymbol{w}_{funding}$. This definition will be useful below when formulating the final decomposition of $\boldsymbol{w}_{\mathrm{TAA}}$ in a form that makes each term in the decomposition more useful in enhancing the transparency of the final allocation obtained from the optimization problem.

II.F.4 Decomposition with normalized portfolios

With the definitions introduced above, it is useful to re-arrange the decomposition in (46) as:

$$\boldsymbol{w}_{\mathrm{TAA}} = \boldsymbol{w}_{\mathrm{SAA}\,minTE} + \boldsymbol{w}_{views} + \widehat{\boldsymbol{w}}_{\alpha} + \widehat{\boldsymbol{w}}_{constraints} + \widehat{\boldsymbol{w}}_{funding} \tag{58}$$

Here, $\boldsymbol{w}_{\mathrm{SAA}\,minTE}$ is fully invested, while $\widehat{\boldsymbol{w}}_{\alpha}$ and $\widehat{\boldsymbol{w}}_{constraints}$ are zero-sum portfolios that adjust asset weights without affecting the total portfolio sum. The final term, $\widehat{\boldsymbol{w}}_{funding}$ combines $\boldsymbol{w}_{funding}$ as defined in (50) with the corrective term in (51), which accounts for replacing $\boldsymbol{w}_{\mathrm{SAA}\,replication}$ with $\boldsymbol{w}_{\mathrm{SAA}\,minTE}$, as well as all corrective terms in (51) arising from normalizing $\boldsymbol{w}_{\alpha}$ and $\boldsymbol{w}_{constraints}$. The sum of weights of the terms in (58) is summarized in Table 2.

Table 2: TAA weights decomposed between relevant sub-portfolios.

Portfolio	SAA $Min\ TE$	$Views$	$Alpha$	$Constraints$	$Funding$
Sum weights	100%	x%	0	0	-x%

The decomposition of the weights in (58) provides a clearer insight into the adjustments made by the optimizer to implement the investment views while satisfying the funding constraint, as we shall illustrate with the numerical examples.

II.F.5 Dispatching offset contributions from weight constraints

While equation (58) offers a clear decomposition of the optimized portfolio into contributions from SAA replication, investment views, alphas, constraints, and the funding adjustment, further refinements may be helpful in practice. This is especially relevant when the effect of a constraint nearly cancels out the impact of another effect with a large contribution, which can make it difficult to discern the true drivers of the final allocation.

For example, consider a strong negative view on an asset that would, in the absence of constraints, result in a substantial negative weight. If a long-only constraint is imposed, then the corresponding sub-portfolio in (58) will very likely show a large positive weight on this same asset to bring it up to zero or higher, effectively neutralizing the intended effect of the view (Green and Hollifield (1992)). More generally, weight bounds and linear constraints can produce large contributions that offset other effects, particularly when the unconstrained solution would otherwise significantly violate these bounds.

When such offsetting occurs, the decomposition may be dominated by large, opposing terms, reducing its interpretability. If this arises because of individual constraints on the maximum and minimum value of weights, then it can be useful to refine the decomposition by reducing the number of sub-portfolios in (58) by redistributing at least some of these most extreme offsetting contributions back on the criteria that caused it in first place. For instance, one may reallocate the portion of the constraint sub-portfolio that neutralizes a view because it violates weight bounds back on the view itself.

Here, we introduce a method for redistributing the impact of a given minimum or maximum constraint on asset weights by reallocating its offsetting effect to the most relevant sources of the allocation that violated it. This refinement ensures that the decomposition remains transparent and meaningful, even when binding constraints significantly alter the unconstrained solution.*

Let $\boldsymbol{D}$ be a matrix obtained from $\widehat{\boldsymbol{w}}_{constraints}$ by keeping only the n_D columns that are to be dispatched onto the other sub-portfolios of (58). All columns of $\boldsymbol{D}$ must be associated with a weight bound constraint and have at least one nonzero weight.

Note that not all columns in $\widehat{\boldsymbol{w}}_{constraints}$ with large values should be dispatched. If a sub-portfolio in $\widehat{\boldsymbol{w}}_{constraints}$ simply comes about because of a regulatory constraint or some other specific reason not related to replication of the SAA, views, alpha, or funding, then trying to dispatch may serve no purpose.

Now, let $\boldsymbol{R}$ be the concatenation into columns of all the n_R sub-portfolios in (58) that have not been chosen for $\boldsymbol{D}$. Stacking both $\boldsymbol{R}$ and $\boldsymbol{D}$ into columns of a matrix $\boldsymbol{M}$:

$$\boldsymbol{M} = [\boldsymbol{R}|\boldsymbol{D}] \tag{59}$$

The matrix $\boldsymbol{M}$ has all the sub-portfolios in the decomposition (58) and thus verifies:

$$\boldsymbol{M}\boldsymbol{1}_{n_R+n_D} = \boldsymbol{w}_{\text{TAA}} \tag{60}$$

where $\boldsymbol{1}_{n_R+n_D}$ is the vector of size $n_R + n_D$ with all coefficients equal to 1. Our goal is to design a $n_D \times n_R$ matrix $\boldsymbol{P}$ that will project the portfolios $\boldsymbol{D}$ onto the same space as the $\boldsymbol{R}$ portfolios, while staying relevant for the allocation explanation. We shall then compute a new $n \times n_R$ matrix $\widehat{\boldsymbol{M}} = \boldsymbol{R} + \boldsymbol{D}\boldsymbol{P}$ that will explain the optimizer choices using only the criteria that have been deemed relevant. Because we still want $\widehat{\boldsymbol{M}}$ columns sum to equal $\boldsymbol{w}_{\text{TAA}}$, $\boldsymbol{P}$ must verify:

$$\boldsymbol{P}\boldsymbol{1}_{n_R} = \boldsymbol{1}_{n_D} \tag{61}$$

Recall that for all i, $\delta_{min,i} \geq 0$ and $\delta_{max,i} \geq 0$. Moreover, the complementary slackness condition implies that the only way for both these Lagrangian multipliers to be nonzero is $w_{min,i} = w_{max,i}$, which happens when the *i-th* asset weight has been set to a given value prior to the optimization. There are three cases:

* Jagannathan and Ma (2003) proposed a different approach to deal with the long only constraints, based on modifying the covariance matrix Σ in mean-variance optimization. However, we find their method less adapted for the purposes of this paper.

- If $\delta_{min,i} > \delta_{max,i}$: the unconstrained weight for asset *i* wants to be below the minimum bound. The optimizer increases the asset's effective return. The constraint pushes the weight upward.
- If $\delta_{min,i} < \delta_{max,i}$: the unconstrained weight for asset *i* wants to be above the maximum bound. The optimizer decreases the asset's effective return. The constraint pushes the weight downwards.
- If $\delta_{min,i} = \delta_{max,i}$: the constraint does not affect the optimizer and the column can be dropped.

Based on these observations, the projection matrix $\boldsymbol{P}$ should increase weights of assets for which $\delta_{min,i} > \delta_{max,i}$ and lower the weights if $\delta_{min,i} < \delta_{max,i}$ in the relevant portfolios $\boldsymbol{R}$**.** In both cases, $\boldsymbol{P}$ should push the weights of the concerned assets towards the bound imposed on the portfolio it must respect.

Let the $n_D \times (n_R + n_D)$ matrix $\boldsymbol{A}$ be constructed row-by-row, where each row corresponds to an asset whose constraint-related sub-portfolio is being dispatched, i.e., the assets appearing in matrix $\boldsymbol{D}$. Its purpose is to encode how the weights of the constraint portfolios $\boldsymbol{D}$ should be redistributed across all sub-portfolios, both $\boldsymbol{R}$ and $\boldsymbol{D}$, based on which constraints are binding and in which direction.

For a given asset belonging to matrix $\boldsymbol{D}$, we look at the corresponding row of $\boldsymbol{M}$ and apply the following selection rules. For each asset i, we compare $\delta_{\min,i}$ for the min-weight constraint with the $\delta_{\max,i}$ for the max-weight constraint. Then,

- For the case $\delta_{\min,i} > \delta_{\max,i}$ we keep only the negative coefficients from row i of $\boldsymbol{M}$. These are the sub-portfolios that try to decrease the weight.
- For the case $\delta_{\min,i} < \delta_{\max,i}$ we keep only the positive coefficients from row i of $\boldsymbol{M}$. These are the sub-portfolios that try to increase the weight.
- For the case $\delta_{\min,i} = \delta_{\max,i}$ then column of $\boldsymbol{D}$ can be ignored and is not used for $\boldsymbol{A}$. We drop this row.

After this, the $\boldsymbol{A}$ rows with nonzero sums are then divided by their sums, while the coefficients of rows with zero sums are all set to $\frac{1}{n_R+n_D}$, so that we have:

$$\boldsymbol{A} = [\boldsymbol{A}_R | \boldsymbol{A}_D] \tag{62}$$

In this way, each row in $\boldsymbol{A}$ is a vector of percentages, summing to 1, indicating how the corresponding asset in $\boldsymbol{D}$ should be projected onto all the other columns of $\boldsymbol{M} = [\boldsymbol{R} \mid \boldsymbol{D}]$.

$$\boldsymbol{A}\boldsymbol{1}_{n_R+n_D} = \boldsymbol{1}_{n_R} \tag{63}$$

Recall that multiplying $\boldsymbol{D}$ by $\boldsymbol{A}$ produces $\boldsymbol{DA} = [\boldsymbol{DA}_R | \boldsymbol{DA}_D]$. This means that part of the weights in $\boldsymbol{D}$ is dispatched to the relevant portfolios $\boldsymbol{R}$ (through $\boldsymbol{DA}_R$), while another part is redistributed back into the constraint portfolios $\boldsymbol{D}$ themselves (through $\boldsymbol{DA}_D$). Because $\boldsymbol{DA}_D$ feeds weights back into $\boldsymbol{D}$, the redistribution process becomes recursive and each time weights loop back into $\boldsymbol{D}$, they must again be projected onto $\boldsymbol{R}$.

To fully dispatch all the weights from $\boldsymbol{D}$ into $\boldsymbol{R}$, we must therefore accumulate this entire recursive sequence $\boldsymbol{D}\boldsymbol{A}_R + \boldsymbol{D}\boldsymbol{A}_D\boldsymbol{A}_R + \boldsymbol{D}\boldsymbol{A}_D^2\boldsymbol{A}_R + \cdots$. This infinite series has a closed-form solution, which leads us to define the projection matrix:

$$\boldsymbol{P} = \left(\mathbf{I}_{n_D} - \boldsymbol{A}_D\right)^{-1}\boldsymbol{A}_R \tag{64}$$

Here, $\mathbf{I}_{n_D}$ is the identity matrix of size n_D. The inverse is well-defined because $\boldsymbol{A}_{\boldsymbol{D}}^k \to 0$ as $k \to \infty$. This holds in practice since each row of $\boldsymbol{A}_{\boldsymbol{D}}$ contains percentages whose row-sum is strictly less than 1 as some weight has already been allocated to $\boldsymbol{A}_{\boldsymbol{R}}$.

Finally, having calculated $\boldsymbol{P}$, we can construct the new matrix $\widehat{\boldsymbol{M}}$:

$$\widehat{\boldsymbol{M}} = \boldsymbol{R} + \boldsymbol{D}\boldsymbol{P} \tag{65}$$

Compared with the original decomposition $\boldsymbol{M}$ in (58), the new matrix $\widehat{\boldsymbol{M}}$ contains fewer columns, since the sub-portfolios in $\boldsymbol{D}$ have now been dispatched onto the relevant components. Each column of $\widehat{\boldsymbol{M}}$ represents a cleaner and more interpretable sub-portfolio. Together, these columns provide a redefined and more transparent breakdown of the final portfolio $\boldsymbol{w}_{\mathrm{TAA}}$. They still correspond to the same drivers of the allocation (replication of the SAA, tactical views, fund alphas, the funding constraint, and any remaining constraints that were not dispatched) while removing the noise created by the offsetting constraint effects.

III. Results

In this section, we illustrate the implementation of the portfolio decomposition framework using as an example the construction of a TAA portfolio benchmarked against an SAA portfolio. The SAA allocates exclusively to core asset classes represented by core indices, while the TAA portfolio is implemented using a mix of active and passive funds. This setting represents the realities of portfolio management, where investment views tend to be formulated at the level of broad asset-class indices, but implementation uses a set of imperfectly aligned investable funds.

All data sources and calculation details are documented in the Appendix. Table A1 lists the core indices used to construct the SAA portfolio and to define the TAA views. The table also reports the SAA allocation itself, as well as the target active unconstrained TAA portfolio derived from these views using the methodology described in Section II.C.

Table A2 provides the characteristics of the funds selected for implementation, including the minimum allocation to Sustainable Investments specified in their respective prospectus. This information is required for the case in which we impose a constraint on the minimum allocation to Sustainable Investments in the optimization.

In Table A3, we report each fund's exposure to the core asset classes, estimated using Lasso regressions as also explained in the Appendix. The table includes the associated R-squared and the volatility of the regression residuals, which quantifies the degree of specific risk of each fund.

Finally, we assume that the active funds were selected based on the expectation that they generate positive alpha beyond their systematic exposures to the core indices, with a target information ratio of +0.5. For each fund, the expected alpha is thus calculated as the product of this information ratio with the fund specific volatility (the residual volatility from the Lasso regression), minus the ongoing charges (OCR). These expected alphas net of OCR can also be found in Table A3.

The risk model is based on a principal components analysis (PCA) approach. The PCA risk factors are constructed from the time series of returns of the core indices. We retain the first six eigenvectors as risk factors. Together they explain 89% of the total return variance. The weights of the core assets in each of these eigenvector portfolios are shown in Table A4 in the Appendix.

The optimization is performed as described in Section II.D, using equation (45), where the implied returns are calculated using (34), with, as inputs, the risk model and the unconstrained tactical views constructed as in equation (25) and shown in Table A1 in the Appendix.

In the Appendix, we provide full details on the construction of the risk model $\mathbf{\Sigma}$ and the uncertainty matrix $\mathbf{\Omega}$**,** both defined over the combined universe of core assets and funds.

The other inputs required for the optimization are the risk aversion parameter, λ, the aversion to uncertainty parameter, κ, and the risk budget, RB, as well as the overall confidence in the expected fund alphas, γ. We have set RB = 2% and γ = 15 and used $\lambda = (1/2) * (0.4/RB)$ and $\kappa = 0.23 * min\left(1, \sum_i \left|\boldsymbol{S}_{dir}^{i}\right|\right)$ as proposed by Mallouli et al. (2025).

III.A No tactical views

We first consider the example in which the tactical views are muted, so that the optimizer simply maximizes the net expected alpha of the fund allocation, tilting in favor of the funds with the strongest positive net alpha and away from funds that have smaller or negative net alpha. Increasing the parameter γ increases the extent to which the portfolio tilts more in favor of funds with positive net alpha. We apply only standard constraints: each asset weight must be non-negative, must not exceed 100%, and the portfolio must be fully invested.

Table 3. Optimal TAA portfolio decomposition in the absence of tactical views

	Portfolio Weights					Fund Exposure to Sustainable Investments	Expected Net Alpha
	TAA	SAA Replication	Alpha	Constraints	Funding		
Equity Europe Mid-large Active Fundamental	5.3%	3.2%	2.1%	0.0%	0.0%	30%	0.6%
Equity Europe Mid-large Passive Index	17.0%	19.7%	-2.7%	0.0%	0.0%	40%	-0.2%
Equity USA Growth Active Fundamental	5.2%	4.2%	0.9%	0.0%	0.0%	25%	4.0%
Equity USA Mid-large Passive Index	5.5%	6.2%	-0.7%	0.0%	0.0%	0%	-0.1%
Equity Japan Mid-large Active Fundamental	5.1%	4.1%	1.0%	0.0%	0.0%	30%	3.3%
Equity Japan Mid-large Passive Index	4.1%	4.8%	-0.7%	0.0%	0.0%	40%	-0.2%
Equity Emerging Mid-large Active Fundamental	4.4%	3.4%	1.0%	0.0%	0.0%	20%	3.0%
Equity Emerging Mid-large Passive Index	4.7%	5.7%	-0.9%	0.0%	0.0%	20%	-0.3%
Bonds Global Aggregate Active Fundamental	4.0%	3.1%	0.9%	0.0%	0.0%	20%	0.1%
Bonds EUR Aggregate Active Fundamental	0.0%	2.2%	-2.2%	0.0%	0.0%	20%	-0.2%
Bonds EUR Sovereign Active Fundamental	7.1%	1.3%	5.7%	0.0%	0.0%	20%	0.0%
Bonds EUR Sovereign Passive Index	19.8%	24.8%	-5.1%	0.0%	0.0%	0%	-0.2%
Bonds EUR IG Active Fundamental	3.9%	0.1%	3.8%	0.0%	0.0%	15%	0.8%
Bonds EUR IG Passive Index	0.0%	2.5%	-2.5%	0.0%	0.0%	30%	-0.2%
Bonds USD IG Passive Index	13.9%	14.6%	-0.7%	0.0%	0.0%	25%	-0.2%
Portfolio Weight Sum	100.0%	100.0%	0.0%	0.0%	0.0%		
Tracking error	1.1%	1.1%	0.2%	0.0%	0.0%		
Expected Net Fund Alpha	0.5%	0.3%	0.2%	0.0%	0.0%		
Allocation to Sustainable Investments	21%	21%	0%	0%	0%		

Notes: IG: Investment Grade. Sustainable Investment allocation is calculated from each fund minimum exposure in Table A2. The tracking error of the TAA portfolio and the SAA Replication sub-portfolio is measured relative to the SAA portfolio. For the Alpha sub-portfolios it is simply its volatility.

Table 3 reports the decomposition of the TAA allocation according to equation (46), breaking the final portfolio into its constituent sub-portfolios. The optimized TAA portfolio holds most funds in the universe, with the largest positions in the Equity Europe Mid–Large Passive Index and the Bonds EUR Sovereign Passive Index. These allocations are primarily driven by the SAA replication component, which seeks to match the strategic exposures to core indices. However, because both passive funds exhibit slightly negative net alpha, their final weights are reduced relative to what a pure replication objective would prescribe. This adjustment is captured in the Alpha sub-portfolio, which tilts toward funds with positive net alpha.

In this example, the SAA Replication portfolio is fully invested, and the Alpha portfolio naturally emerges as zero-sum even without explicit normalization. This follows directly from the optimization: none of the imposed constraints bind, so the solution meets all requirements without additional constraint-driven adjustments.

The subsequent examples introduce non-zero implied returns derived from the tactical views. As before, we impose only standard constraints: non-negative weights, no position above 100%, and full investment. In the final example, we also add a minimum allocation of 30% to Sustainable Investments to illustrate how such constraints can be captured by the decomposition.

III.B Tactical views

Table 4 reports the results using the same decomposition as in equation (46), but now with tactical views. In this case, several constraints are binding. The Constraints column aggregates the adjustments required to enforce weight bounds. These are triggered by negative views that would otherwise imply short positions. For example, the negative tactical view on the Bond EUR Sovereign core index would push the optimizer to short the corresponding fund (Bonds EUR Sovereign Passive Index), which would violate the long-only bounds; the constraint sub-portfolio introduces the compensating correction.

Relative to Table 1, the influence of tactical views is now evident and explains additional deviations from the SAA Replication portfolio. A positive view on Equity EMU generates an overweight in the Equity Europe Mid-Large Passive Index fund; a negative view on Bond EUR Sovereign produces a strong underweight in the Bonds EUR Sovereign Passive Index fund; and a positive view on Bond EUR IG results in an overweight in the Bonds EUR IG Passive Index fund. Two other views concern core assets without a direct fund proxy in the selected universe: the positive view on Bonds USD HY ultimately tilts towards the Bonds USD IG Passive Index and the Bonds EUR IG Passive Index funds, while the mildly negative view on Bond EMD HC induces modest underweights in the same pair, reflecting their nearest-proxy roles.

In this example, the SAA Replication portfolio is only slightly short of full investment, and the Alpha portfolio is close to but not exactly zero-sum. That will not always hold. When these components deviate more materially from full-investment or zero-sum, respectively, the decomposition can become harder to interpret, which motivates the normalizations introduced in Sections II.F.1 and II.F.3.

Finally, while the aggregate Constraints column in Table 4 remains interpretable, a further split into the individual non-zero weight-bound corrections, as shown in Table 5, makes it less transparent: although only three bounds are active, their separate contributions are not straightforward to read. This is precisely the situation addressed by the dispatching procedure in

Section II.F.5, which reallocates the most extreme weight-bound effects back to their underlying drivers, preserving a cleaner explanation.

Table 4. Optimal TAA portfolio decomposition

	Portfolio Weights									
	TAA	SAA Replication	Views					Alpha	Constraints	Funding
			Equity EMU	Bond EUR Sovereign	Bond EUR IG	Bond USD HY	Bond EMD HC			
			0.50	-1.00	0.50	0.50	-0.25			
Equity Europe Mid-large Active Fundamental	5.2%	2.9%	0.8%	0.0%	0.0%	0.1%	0.0%	1.4%	0.0%	-0.1%
Equity Europe Mid-large Passive Index	21.9%	17.7%	4.9%	0.0%	0.1%	0.9%	-0.3%	-1.6%	0.2%	0.0%
Equity USA Growth Active Fundamental	5.6%	5.2%	0.1%	0.0%	0.0%	0.3%	-0.1%	0.5%	-0.1%	-0.2%
Equity USA Mid-large Passive Index	7.3%	7.5%	0.1%	0.0%	0.0%	0.4%	-0.1%	-0.6%	-0.1%	0.2%
Equity Japan Mid-large Active Fundamental	5.7%	4.4%	0.1%	0.0%	0.0%	0.2%	0.0%	0.7%	0.1%	0.2%
Equity Japan Mid-large Passive Index	5.4%	5.2%	0.1%	0.0%	0.0%	0.3%	0.0%	-0.5%	0.1%	0.2%
Equity Emerging Mid-large Active Fundamental	4.5%	3.2%	0.0%	0.1%	0.0%	0.4%	-0.3%	0.7%	0.3%	0.1%
Equity Emerging Mid-large Passive Index	5.6%	5.3%	0.1%	0.1%	0.0%	0.7%	-0.5%	-0.6%	0.4%	0.0%
Bonds Global Aggregate Active Fundamental	3.2%	6.8%	-0.1%	-0.7%	0.0%	-1.8%	-0.6%	0.8%	-3.6%	2.4%
Bonds EUR Aggregate Active Fundamental	0.0%	1.6%	0.0%	-2.5%	0.8%	0.1%	-0.1%	-5.1%	-4.3%	9.6%
Bonds EUR Sovereign Active Fundamental	0.0%	1.0%	0.0%	-1.8%	0.0%	-0.1%	0.0%	4.0%	-3.6%	0.4%
Bonds EUR Sovereign Passive Index	0.0%	21.3%	-0.1%	-37.0%	0.0%	-1.8%	-0.2%	-0.5%	22.7%	-4.5%
Bonds EUR IG Active Fundamental	5.4%	0.0%	0.0%	0.0%	0.1%	0.0%	0.0%	3.1%	0.0%	2.1%
Bonds EUR IG Passive Index	19.4%	3.9%	0.3%	-0.2%	18.9%	5.8%	-1.6%	-1.7%	-3.0%	-2.9%
Bonds USD IG Passive Index	10.7%	12.6%	0.1%	-0.5%	0.4%	3.4%	-2.2%	-0.4%	-2.2%	-0.6%
Portfolio Weight Sum	100.0%	98.4%	6.3%	-42.4%	20.4%	9.1%	-5.9%	0.3%	6.8%	7.0%
Tracking error	2.2%	1.1%	0.1%	0.1%	0.1%	0.5%	0.2%	0.2%	0.5%	0.3%
Expected Net Fund Alpha	0.5%	0.3%	0.0%	0.1%	0.0%	0.0%	0.0%	0.1%	0.0%	0.0%
Allocation to Sustainable Investments	28%	21%	2%	-1%	6%	3%	-1%	0%	-4%	2%

Notes: IG: Investment Grade, HY: High Yield, HC: Hard Currency. Sustainable Investment allocation is calculated from each fund minimum exposure in Table A2. The tracking error of the TAA portfolio and the SAA Replication sub-portfolio is measured relative to the SAA portfolio. For each view sub-portfolio, the tracking error is measured against the portfolio that expresses the underlying view using core indices only. For the Alpha, Constraints, and Funding sub-portfolios, the relevant risk measure is simply their volatility.

The sub-portfolios in Table 5 that correct for violations of weight bounds appear extreme because they correspond to the unconstrained minimum-variance portfolios with unit exposure to asset i, as can be seen from the structure of equation (49). In other words, when a weight bound is breached, the optimizer effectively introduces the hedging portfolio that offsets the unwanted exposure as efficiently as possible in variance terms, an inherently aggressive adjustment. Using the robustified matrix $\boldsymbol{Q}$ instead of the traditional covariance matrix $\boldsymbol{\Sigma}$ is not enough adequate to attenuate this effect sufficiently and the $\boldsymbol{Q}$ based minimum-variance hedging portfolios are still rather sensitive to correlations of assets.

Even so, the *sum* of these constraint-related sub-portfolios is much easier to interpret. The complete decomposition must add up exactly to the TAA portfolio, and the other components (SAA Min-TE, Alpha, and Funding) tend to be far less extreme by construction. As a result, while the individual constraint-correction portfolios can be difficult to interpret on their own, their aggregate effect fits naturally into the overall decomposition and remains consistent with the economic drivers of the final allocation.

Table 6 reports the decomposition of the same portfolio as in Table 4, but now (i) forcing the SAA Replication portfolio to be fully invested (per equation (51)), (ii) normalizing the Alpha portfolio to be zero-sum (per equation (57)), and (iii) dispatching the Constraint sub-portfolios (per equation (65)).

Table 5. Decomposition of constraints sub-portfolio in Table 3

	Portfolio Weights			
	Constraints	Constraints		
		Bonds EUR Aggregate Active Fundamental	Bonds EUR Sovereign Active Fundamental	Bonds EUR Sovereign Passive Index
Equity Europe Mid-large Active Fundamental	0.0%	0.0%	0.0%	0.0%
Equity Europe Mid-large Passive Index	0.2%	0.0%	0.0%	0.2%
Equity USA Growth Active Fundamental	-0.1%	0.0%	0.0%	-0.1%
Equity USA Mid-large Passive Index	-0.1%	0.0%	0.0%	-0.1%
Equity Japan Mid-large Active Fundamental	0.1%	0.0%	0.0%	0.1%
Equity Japan Mid-large Passive Index	0.1%	0.0%	0.0%	0.1%
Equity Emerging Mid-large Active Fundamental	0.3%	0.0%	0.0%	0.2%
Equity Emerging Mid-large Passive Index	0.4%	0.0%	0.0%	0.4%
Bonds Global Aggregate Active Fundamental	-3.6%	-0.4%	-0.1%	-3.1%
Bonds EUR Aggregate Active Fundamental	-4.3%	484.4%	-21.7%	-467.0%
Bonds EUR Sovereign Active Fundamental	-3.6%	-15.2%	345.3%	-333.8%
Bonds EUR Sovereign Passive Index	22.7%	-320.3%	-327.9%	671.0%
Bonds EUR IG Active Fundamental	0.0%	-0.8%	0.0%	0.7%
Bonds EUR IG Passive Index	-3.0%	-129.7%	5.6%	121.1%
Bonds USD IG Passive Index	-2.2%	-0.1%	-0.1%	-2.0%
Portfolio Weight Sum	6.8%	17.9%	1.1%	-12.2%
Tracking error	0.5%	2.7%	2.8%	3.9%
Expected Net Fund Alpha	-1.4%	-0.4%	0.4%	-0.1%
Allocation to Sustainable Investments	0%	55%	66%	-125%

Notes: IG: Investment Grade. Sustainable Investment allocation is calculated from each fund minimum exposure in Table A2. The tracking error of each view sub-portfolio is measured against the portfolio that expresses the underlying view using core indices only. For the Constraints sub-portfolio, created as their sum, it is just its volatility.

Table 6. Optimal TAA portfolio decomposition with normalization and dispatching of offsetting constraints

	Portfolio Weights								
	TAA	SAA Min TE	Views					Alpha	Funding
			Equity EMU	Bond EUR Sovereign	Bond EUR IG	Bond USD HY	Bond EMD HC		
			0.50	-1.00	0.50	0.50	-0.25		
Equity Europe Mid-large Active Fundamental	5.2%	2.9%	0.8%	0.0%	0.0%	0.1%	0.0%	1.5%	-0.1%
Equity Europe Mid-large Passive Index	21.9%	17.7%	4.9%	0.1%	0.1%	0.9%	-0.3%	-1.5%	0.0%
Equity USA Growth Active Fundamental	5.6%	5.1%	0.1%	0.1%	0.0%	0.3%	-0.1%	0.7%	-0.4%
Equity USA Mid-large Passive Index	7.3%	7.5%	0.1%	-0.1%	0.0%	0.4%	-0.1%	-0.7%	0.3%
Equity Japan Mid-large Active Fundamental	5.7%	4.4%	0.1%	0.0%	0.0%	0.2%	0.0%	0.6%	0.4%
Equity Japan Mid-large Passive Index	5.4%	5.2%	0.1%	0.0%	0.0%	0.3%	0.0%	-0.5%	0.4%
Equity Emerging Mid-large Active Fundamental	4.5%	3.2%	0.0%	0.1%	0.0%	0.4%	-0.3%	0.8%	0.2%
Equity Emerging Mid-large Passive Index	5.6%	5.3%	0.1%	0.3%	0.0%	0.7%	-0.5%	-0.4%	0.0%
Bonds Global Aggregate Active Fundamental	3.2%	7.2%	-0.1%	-3.1%	0.0%	-1.9%	-0.6%	-2.3%	3.9%
Bonds EUR Aggregate Active Fundamental	0.0%	3.1%	0.0%	-14.5%	0.8%	-0.6%	-0.1%	-3.0%	14.3%
Bonds EUR Sovereign Active Fundamental	0.0%	0.9%	0.0%	-5.9%	0.0%	-0.3%	0.0%	5.7%	-0.4%
Bonds EUR Sovereign Passive Index	0.0%	21.2%	0.0%	-18.8%	0.0%	-0.9%	-0.1%	2.4%	-3.8%
Bonds EUR IG Active Fundamental	5.4%	0.5%	0.0%	-0.8%	0.1%	0.0%	0.0%	1.9%	3.7%
Bonds EUR IG Passive Index	19.4%	3.3%	0.3%	1.4%	18.9%	5.9%	-1.6%	-4.0%	-4.7%
Bonds USD IG Passive Index	10.7%	12.4%	0.1%	-1.1%	0.4%	3.4%	-2.2%	-1.1%	-1.3%
Portfolio Weight Sum	100.0%	100.0%	6.3%	-42.4%	20.4%	9.1%	-5.9%	0.0%	12.5%
Tracking error	2.2%	1.1%	0.1%	0.3%	0.1%	0.5%	0.2%	0.2%	0.5%
Expected Net Fund Alpha	0.5%	0.3%	0.0%	0.1%	0.0%	0.0%	0.0%	0.1%	0.0%
Allocation to Sustainable Investments	28%	21%	2%	-5%	6%	3%	-1%	-1%	3%

Notes: IG: Investment Grade, HY: High Yield, HC: Hard Currency. Sustainable Investment allocation is calculated from each fund minimum exposure in Table A2. The tracking error of the TAA portfolio and the SAA Min TE sub-portfolio is measured relative to the SAA portfolio. For each view sub-portfolio, the tracking error is measured against the portfolio that expresses the underlying view using core indices only. For the Alpha and Funding sub-portfolios, the relevant risk measure is simply their volatility.

Under this normalization, the results for SAA Min-TE in Table 6 are very close to the SAA Replication in Table 4, and the Alpha sup portfolios are also broadly consistent in the two tables. The main differences are: first, the explicit Constraints columns are no longer present (their effects have been dispatched to the relevant drivers); second, the Tactical Views sub-portfolio adjusts accordingly. In particular, the previously negative view on Bond EUR Sovereign, which had generated a violation of the long-only constraint, now appears less negative on the Bonds EUR Sovereign Passive Index fund. Instead, part of the adjustment is reallocated as an additional underweight in the Bonds EUR Aggregate Active Fundamental fund.

Overall, the zero allocation across these two bond funds is decomposed differently than in Table 4 yet remains economically consistent with those earlier results but now expressed without resorting to the separate Constraint sub-portfolios and with clearer attribution to the view that caused the adjustment.

III.C Additional linear constraints

In the final case, in Table 7 we present the decomposition for the same example, now with an additional linear constraint requiring a minimum allocation of 30% to Sustainable Investments in the TAA portfolio. The breakdown follows the approach in Table 6: the SAA Min TE portfolio is fully invested, the Alpha portfolio is normalized to be zero-sum, and weight-bound Constraint portfolios are dispatched. However, we do not dispatch the Constraint sub-portfolio associated with the sustainability requirement, as we wish to assess its magnitude and impact explicitly.

Table 7. Optimal TAA portfolio decomposition with normalization, dispatching of offsetting weight constraints and contribution from a constraint on minimum allocation to sustainable investments

	Portfolio Weights											
	TAA	SAA Min TE	Views					Alpha	Constraints	Funding		
			Equity EMU	Bond EUR Sovereign	Bond EUR IG	Bond USD HY	Bond EMD HC		30% Minimum Allocation to Sustainable Investments		Fund Exposure to Sustainable Investments	Expected Net Alpha
			0.50	-1.00	0.50	0.50	-0.25					
Equity Europe Mid-large Active Fundamental	2.1%	2.9%	0.8%	0.0%	0.0%	0.2%	0.0%	1.5%	-3.1%	-0.1%	30%	0.6%
Equity Europe Mid-large Passive Index	25.9%	17.8%	5.1%	0.0%	0.0%	0.9%	-0.3%	-1.6%	3.9%	0.0%	40%	-0.2%
Equity USA Growth Active Fundamental	6.4%	5.1%	0.0%	0.0%	0.0%	0.3%	-0.1%	0.6%	0.9%	-0.5%	25%	4.0%
Equity USA Mid-large Passive Index	4.8%	7.5%	0.1%	0.0%	0.0%	0.4%	-0.1%	-0.6%	-2.6%	0.3%	0%	-0.1%
Equity Japan Mid-large Active Fundamental	5.5%	4.4%	0.0%	0.0%	0.0%	0.2%	0.0%	0.6%	-0.4%	0.5%	30%	3.3%
Equity Japan Mid-large Passive Index	6.1%	5.2%	0.0%	0.0%	0.0%	0.3%	0.0%	-0.5%	0.6%	0.4%	40%	-0.2%
Equity Emerging Mid-large Active Fundamental	4.4%	3.2%	0.0%	0.1%	0.0%	0.4%	-0.3%	0.8%	0.0%	0.2%	20%	3.0%
Equity Emerging Mid-large Passive Index	5.4%	5.3%	0.0%	0.1%	0.0%	0.7%	-0.5%	-0.5%	0.2%	0.0%	20%	-0.3%
Bonds Global Aggregate Active Fundamental	0.0%	7.2%	-0.1%	-1.0%	0.0%	-1.7%	-0.6%	0.0%	-8.3%	4.4%	20%	0.1%
Bonds EUR Aggregate Active Fundamental	0.0%	2.7%	0.0%	-22.8%	0.7%	-1.0%	-0.1%	-13.2%	20.5%	13.3%	20%	-0.2%
Bonds EUR Sovereign Active Fundamental	0.0%	0.3%	0.0%	-19.5%	-0.1%	-1.0%	-0.1%	-9.1%	32.6%	-3.1%	20%	0.0%
Bonds EUR Sovereign Passive Index	0.0%	22.0%	0.0%	-4.2%	0.0%	-0.2%	0.0%	20.2%	-36.8%	-0.9%	0%	-0.2%
Bonds EUR IG Active Fundamental	0.0%	0.5%	0.0%	0.2%	0.1%	0.0%	0.0%	3.4%	-8.2%	4.0%	15%	0.8%
Bonds EUR IG Passive Index	28.4%	3.5%	0.2%	4.7%	19.7%	6.2%	-1.6%	-0.8%	1.0%	-4.3%	30%	-0.2%
Bonds USD IG Passive Index	11.0%	12.5%	0.1%	-0.6%	0.2%	3.4%	-2.2%	-0.7%	-0.3%	-1.3%	25%	-0.2%
Portfolio Weight Sum	100.0%	100.0%	6.3%	-43.1%	20.7%	9.2%	-6.0%	0.0%	0.0%	12.9%		
Tracking error	2.2%	1.1%	0.1%	0.2%	0.0%	0.5%	0.2%	0.2%	0.5%	0.5%		
Expected Net Fund Alpha	0.4%	0.3%	0.0%	0.1%	0.0%	0.0%	0.0%	0.1%	-0.1%	0.0%		
Allocation to Sustainable Investments	30%	21%	2%	-7%	6%	3%	-1%	-4%	9%	2%		

Notes: IG: Investment Grade, HY: High Yield, HC: Hard Currency. Sustainable Investment allocation is calculated from each fund minimum exposure in Table A2. The tracking error of the TAA portfolio and the SAA Min TE sub-portfolio is measured relative to the SAA portfolio. For each view sub-portfolio, the tracking error is measured against the portfolio that expresses the underlying view using core indices only. For the Alpha, Constraints, and Funding sub-portfolios, the relevant risk measure is simply their volatility.

Imposing the 30% floor materially reshapes the optimized TAA. The optimizer reallocates toward funds with higher Sustainable Investments exposure and away from those with limited or no sustainable content. Funds with stronger sustainability profiles absorb a larger share, as they are the most efficient vehicles to meet the requirement with minimal disruption to tracking error and expected performance. Conversely, lower-Sustainable Investments funds are reduced, sometimes to their lowest feasible levels, with the sustainability-constraint sub-portfolio recording the necessary offsetting adjustments.

Through dispatching, the constraint interacts with the Views and Alpha components. View-driven positions remain directionally consistent but are mildly attenuated to accommodate the sustainability tilt. As with weight-bound corrections in Table 5, however, keeping the sustainability Constraint sub-portfolio explicit can complicate interpretation; dispatching the other constraints makes this interaction more visible in the Alpha term. The complexity is amplified for funds in which the tactical views already push towards zero in Tables 3 and 5.

Two practical remedies are available. One is to revert to the simplified presentation used in Table 4, aggregating all constraint effects, including the sustainability floor, into a single Constraints portfolio. The other, if the sustainability Constraint portfolio is to be shown explicitly, is to reapply the framework to a reduced investable universe that excludes funds effectively set to zero by tactical views. We report the results of the latter approach in Table 8 after removing the funds Bonds EUR Aggregate Active Fundamental, Bonds EUR Sovereign Active Fundamental and Bonds EUR Sovereign Passive Index.

Table 8. Optimal TAA portfolio decomposition with normalization, dispatching of offsetting weight constraints, contribution from a constraint on sustainable investments on restricted fund universe.

	Portfolio Weights											
	TAA	SAA Min TE	Views					Alpha	Constraints	Funding		
			Equity EMU	Bond EUR Sovereign	Bond EUR IG	Bond USD HY	Bond EMD HC		30% Minimum Allocation to Sustainable Investments		Fund Exposure to Sustainable Investments	Expected Net Alpha
			0.50	-1.00	0.50	0.50	-0.25					
Equity Europe Mid-large Active Fundamental	2.1%	2.7%	0.8%	0.1%	0.0%	0.2%	0.0%	1.5%	-3.2%	0.1%	30%	0.6%
Equity Europe Mid-large Passive Index	25.9%	17.5%	5.1%	0.5%	0.0%	0.9%	-0.3%	-1.6%	3.8%	0.0%	40%	-0.2%
Equity USA Growth Active Fundamental	6.4%	4.7%	0.0%	-0.2%	0.0%	0.3%	-0.1%	0.6%	0.7%	0.3%	25%	4.0%
Equity USA Mid-large Passive Index	4.8%	8.0%	0.1%	-0.3%	0.0%	0.4%	-0.1%	-0.6%	-2.3%	-0.2%	0%	-0.1%
Equity Japan Mid-large Active Fundamental	5.5%	4.7%	0.0%	0.3%	0.0%	0.2%	0.0%	0.6%	-0.2%	-0.2%	30%	3.3%
Equity Japan Mid-large Passive Index	6.1%	5.4%	0.0%	0.3%	0.0%	0.3%	0.0%	-0.5%	0.7%	-0.2%	40%	-0.2%
Equity Emerging Mid-large Active Fundamental	4.4%	2.9%	0.0%	0.7%	0.0%	0.5%	-0.3%	0.7%	-0.1%	0.0%	20%	3.0%
Equity Emerging Mid-large Passive Index	5.4%	4.5%	0.0%	1.2%	0.0%	0.8%	-0.5%	-0.6%	-0.2%	0.1%	20%	-0.3%
Bonds Global Aggregate Active Fundamental	0.0%	18.4%	-0.1%	-9.4%	0.1%	-2.1%	-0.6%	0.6%	-3.8%	-3.1%	20%	0.1%
Bonds EUR IG Active Fundamental	0.0%	4.6%	0.0%	-0.1%	0.1%	0.0%	0.0%	3.3%	-6.8%	-1.2%	15%	0.8%
Bonds EUR IG Passive Index	28.4%	11.3%	0.2%	-12.0%	19.9%	5.4%	-1.7%	-3.6%	10.3%	-1.5%	30%	-0.2%
Bonds USD IG Passive Index	11.0%	15.4%	0.1%	-6.5%	0.3%	3.2%	-2.3%	-0.4%	1.1%	0.2%	25%	-0.2%
Portfolio Weight Sum	100.0%	100.0%	6.3%	-25.2%	20.5%	10.0%	-5.9%	0.0%	0.0%	-5.7%		
Tracking error	2.2%	1.4%	0.1%	1.2%	0.1%	0.5%	0.2%	0.2%	0.3%	0.2%		
Expected Net Fund Alpha	0.4%	0.4%	0.0%	0.0%	0.0%	0.0%	0.0%	0.1%	-0.1%	0.0%		
Allocation to Sustainable Investments	30%	26%	2%	-6%	6%	3%	-1%	-1%	2%	-1%		

Notes: IG: Investment Grade, HY: High Yield, HC: Hard Currency. Sustainable Investment allocation is calculated from each fund minimum exposure in Table A2. The tracking error of the TAA portfolio and the SAA Min TE sub-portfolio is measured relative to the SAA portfolio. For each view sub-portfolio, the tracking error is measured against the portfolio that expresses the underlying view using core indices only. For the Alpha, Constraints, and Funding sub-portfolios, the relevant risk measure is simply their volatility.

Table 8 addresses the interpretability issues observed in Table 6. In this version, the Alpha sub-portfolio once again behaves as expected by tilting in favor of funds with higher net alpha and away from those with lower or negative net alpha. The Constraints sub-portfolio cleanly reflects the sustainability requirement, reallocating weight away from funds with limited Sustainable Investments exposure and towards those with stronger commitments.

The Tactical-views sub-portfolios remain easy to interpret. They continue to tilt allocations in line with the underlying directional views, favoring funds that provide exposure to positively viewed asset classes and reducing allocations to those linked to negatively viewed ones. The most negative view, on Bond EUR Sovereign, now produces an underweight in the Bonds EUR IG Passive Index and the Bonds Global Aggregate Active Fundamental funds. Both funds exhibit meaningful exposure to Bond EUR Sovereign risk, as captured by the systematic variance–covariance structure, $\boldsymbol{\Sigma}_{systematic}^{fund,fund}$ in (A6), and the corresponding elements of the uncertainty matrix, $\boldsymbol{\Omega}_{systemic}^{fund,fund}$ in (A9).

While this refinement improves interpretability, it comes at a cost: the SAA Min-TE sub-portfolio must adapt to the reduced investable universe after excluding certain funds, which leads to an increase in its tracking error relative to the original baseline in Table 6.

Overall, Table 8 restores a clean, coherent decomposition in which Alpha, Views, and Constraints once again reflect their intended economic drivers. The sustainability requirement yields a transparent, interpretable re-tilt of the TAA portfolio, achieved through this second iteration of the framework. The decomposition remains fully additive: each fund's final weight can still be traced back to contributions from SAA replication, tactical views, fund net alphas, the explicit sustainability constraint, and the funding adjustment.

IV. Discussion

We propose a practical approach for real-world applications that makes the allocation of portfolio optimization for a benchmarked, constrained TAA portfolio transparent and traceable to its drivers. The framework builds on the well-known result that an MVO portfolio can be written as an unconstrained solution plus correction terms and extend this decomposition to RPO. We adapt the approach to state-of-the-art TAA portfolio construction using RPO so that each optimized weight can be decomposed into contributions of sub-portfolios that reflect a replication of the SAA using a list of funds selected for the implementation, the tactical views, the expected net alpha of the selected funds, the binding constraints, and the funding adjustment required to meet the fully invested constraint. This framework allows stakeholders to see precisely why the optimizer chose a given allocation. A case study with real funds and a minimum sustainable-investment constraint illustrates these elements clearly in the decomposition.

V. Appendix A

In Table A1, we list the core assets that compose the SAA portfolio, together with their indices, associated volatilities, and SAA weights. The core assets form the set *N* of non-investable assets in the analytical framework of this paper. The table also presents the directional tactical bets used in numerical examples, as well as the corresponding unconstrained active tactical portfolio constructed from these directional scores and asset volatilities using equation (25).

In Table A2, we present the characteristics of the funds selected for implementing the TAA portfolio. The data set used in this paper is based on existing, investable funds. For fixed-income funds denominated in non-EUR currencies, we use the EUR-hedged share class to ensure consistency with the EUR-based risk model. For equity funds, currency returns are simply converted into EUR. The minimum allocation to Sustainable Investments for each fund, defined according to the Sustainable Finance Disclosure Regulation (SFDR) framework of the European Union, is obtained directly from the respective fund prospectus.

Table A1. SAA portfolio and the active unconstrained TAA portfolio

Core Assets	Tickers	Index provider	Currency hedging into EUR	Volatility	SAA portfolio allocation	Directional Tactical View Score S_i	Unconstrained Tactical Active Weights
Equity Europe EMU	NDDLEURO Index	1	-	15.5%	19%	0.50	6%
Equity Europe EMU SC	NCLDEMU Index	1	-	16.5%			
Equity Europe UK	NDDLUK Index	1	No	11.7%			
Equity USA	NDDUUS Index	1	No	15.4%	15%		
Equity USA SC	RU20INTR Index	2	No	20.8%			
Equity Japan	NDDLJN Index	1	No	14.1%	9%		
Equity Emerging Global	NDUEEGF Index	1	No	16.5%	7%		
Bond EUR Sovereign	LEATTREU Index	3	-	5.4%	20%	-1.00	-43%
Bond EUR Investment Grade	LECPTREU Index	3	-	4.7%	5%	0.50	24%
Bond EUR High Yield	LF88TREU Index	3	-	7.0%			
Bond USD Sovereign	LUATTRUU Index	3	Yes	4.8%	13%		
Bond USD IG	LUACTRUU Index	3	Yes	6.9%	12%		
Bond USD HY	LF89TRUU Index	3	Yes	7.3%		0.50	11%
Bond EMD HC Sov Global	JPGCCOMP Index	4	Yes	9.0%		-0.25	-6%
Bond EMD LC Sov Global	JGENVUUG Index	4	Yes	10.7%			
Diversification Real Estate Pan-Europe	TRNHUE Index	5	Yes	18.9%			
Diversification Commodity Global	BCOMXALT Index	3	Yes	16.7%			

Notes: SC: Small Caps, IG: Investment Grade, HY: High Yield, HC: Hard Currency, LC: Local Currency. Source: 1) MSCI, 2) Bloomberg, 3) J.P. Morgan, 4) FTSE EPRA. Authors' calculations.

Table A2. Characteristics of the funds selected for portfolio construction

Designation	Asset Class	Coverage	Style	Approch	Philosophy	Hedging into EUR	Sustainable Investment	Inception Date
Equity Europe Mid-large Active Fundamental	Equity	Eurozone	Mid-large	Active	Fundamental	No	30%	23-Oct-03
Equity Europe Mid-large Passive Index	Equity	Eurozone	Mid-large	Passive	Index	No	40%	1-Dec-10
Equity USA Growth Active Fundamental	Equity	USA	Growth	Active	Fundamental	No	25%	3-Jan-95
Equity USA Mid-large Passive Index	Equity	USA	Mid-large	Passive	Index	No	0%	10-Jun-08
Equity Japan Mid-large Active Fundamental	Equity	Japan	Mid-large	Active	Fundamental	No	30%	31-Dec-90
Equity Japan Mid-large Passive Index	Equity	Japan	Mid-large	Passive	Index	No	40%	2-Aug-23
Equity Emerging Mid-large Active Fundamental	Equity	Emerging Markets	Mid-large	Active	Fundamental	No	20%	20-Oct-97
Equity Emerging Mid-large Passive Index	Equity	Emerging Markets	Mid-large	Passive	Index	No	20%	3-Sep-12
Bonds Global Aggregate Active Fundamental	Bond	Global	Aggregate	Active	Fundamental	Yes	20%	5-Nov-01
Bonds EUR Aggregate Active Fundamental	Bond	EUR	Aggregate	Active	Fundamental	-	20%	4-Apr-00
Bonds EUR Sovereign Active Fundamental	Bond	EUR	Sovereign	Active	Fundamental	-	20%	27-Jun-01
Bonds EUR Sovereign Passive Index	Bond	EUR	Sovereign	Passive	Index	-	0%	31-May-17
Bonds EUR IG Active Fundamental	Bond	EUR	Investment Grade	Active	Fundamental	-	15%	1-Feb-22
Bonds EUR IG Passive Index	Bond	EUR	Investment Grade	Passive	Index	-	30%	15-Jan-19
Bonds USD IG Passive Index	Bond	USD	Investment Grade	Passive	Index	Yes	25%	12-Sep-23

Notes: SC: Small Caps, IG: Investment Grade, HY: High Yield, HC: Hard Currency, LC: Local Currency.

In Table A3, we report the results of the regression analysis performed on the vector of weekly returns in excess of cash $\boldsymbol{XR}_f^i$ of each fund i listed in Table A2 against the matrix of weekly returns in excess of cash $\boldsymbol{XR}_c$ of the core indices in Table A1. The regression is performed from end of November 2020 to end of November 2025. To estimate the factor exposures, we employ Lasso regressions as proposed by Tibshirani (1996) rather than ordinary least squares. Lasso's ℓ_1-penalty allows coefficients associated with non-significant explanatory variables to shrink exactly to zero, yielding parsimonious and interpretable exposures while mitigating overfitting.

For each fund i, the vector $\boldsymbol{\beta}_{fc}^{i}$ with the betas is obtained by solving:

$$\widehat{\boldsymbol{\beta}}_{\text{LASSO}} = \arg\min_{\boldsymbol{\beta}} \left[\| \boldsymbol{XR}_f^i - \left(\boldsymbol{\beta}_{fc}^i\right)^{\top} \boldsymbol{XR}_c \|^2 + \xi \; \| \boldsymbol{\beta}_{fc}^i \|_1 \right] \tag{A1}$$

Where $\|\cdot\|^2$ denotes the squared Euclidean (i.e., ℓ_2) norm, $\|.\|_1$ is the ℓ_1 norm and ξ is the regularization parameter. The value of ξ is selected via cross-validation over a grid of 100 candidate values ranging from ξ, the smallest penalty that zeros all coefficients, down to $\xi_{\min} = \xi_{\max}/10^5$. The final choice of ξ follows the one-standard-error rule, balancing model sparsity and predictive accuracy, as recommended by Hastie et al. (2015).

Table A3 also reports the Ongoing Charges (OCR) of each fund. The alphas shown in the table are net of these charges and are used directly in equation (34). For each fund, the expected alpha is calculated as the product of its expected information ratio and its specific volatility (the residual volatility from the Lasso regression), minus the OCR. The information ratio is set to 0.5 for actively managed funds and to 0 for passive funds, reflecting the assumption that the selected active funds were included based on their managers' ability to generate alpha.

Table A3. Fund analytics

						Beta of Funds to Core Assets																
Fund Designation	Specific Volatility of fund relative to index exposures	R square of regression Fund relative to Core Assets	Expected information ratio before costs	OCR	Expected alpha	Equity Europe EMU	Equity Europe EMU SC	Equity Europe UK	Equity USA	Equity USA SC	Equity Japan	Equity Emerging Global	Bond EUR Sovereign	Bond EUR Investment Grade	Bond EUR High Yield	Bond USD Sovereign	Bond USD Investment Grade	Bond USD High Yield	Bond EMD HC Sov Global	Bond EMD LC Sov Global	Diversification Real Estate Pan-Europe	Diversification Commodity Global
Equity Europe Mid-large Active Fundamental	3.0%	0.96	0.5	0.96%	0.6%	0.99																
Equity Europe Mid-large Passive Index	1.2%	0.99	0.0	0.15%	-0.2%	0.99																
Equity USA Growth Active Fundamental	10.0%	0.79	0.5	0.96%	4.0%				1.21													
Equity USA Mid-large Passive Index	6.9%	0.81	0.0	0.14%	-0.1%				0.84													
Equity Japan Mid-large Active Fundamental	8.5%	0.72	0.5	0.98%	3.3%						0.72											
Equity Japan Mid-large Passive Index	8.0%	0.72	0.0	0.15%	-0.2%						0.76											
Equity Emerging Mid-large Active Fundamental	8.2%	0.73	0.5	1.11%	3.0%							0.77										
Equity Emerging Mid-large Passive Index	6.6%	0.83	0.0	0.27%	-0.3%							0.83										
Bonds Global Aggregate Active Fundamental	2.6%	0.79	0.5	1.15%	0.1%									0.35		0.37	0.27					
Bonds EUR Aggregate Active Fundamental	0.6%	0.99	0.5	0.49%	-0.2%								0.68	0.30								
Bonds EUR Sovereign Active Fundamental	0.8%	0.98	0.5	0.43%	0.0%								0.97									
Bonds EUR Sovereign Passive Index	0.2%	1.00	0.0	0.15%	-0.2%								0.97									
Bonds EUR IG Active Fundamental	2.6%	0.71	0.5	0.47%	0.8%									0.93								
Bonds EUR IG Passive Index	0.2%	1.00	0.0	0.15%	-0.2%									1.12								
Bonds USD IG Passive Index	0.2%	1.00	0.0	0.20%	-0.2%												1.02					

Notes: IG: Investment Grade.

Source: MSCI, Bloomberg, J.P. Morgan, FTSE EPRA. Authors' calculations.

VI. Appendix B

Following Bass, Gladstone, and Ang (2017), we construct the risk model using monthly returns in excess of cash of the core asset classes, expressed in EUR, over a 20-year period from the end of November 2005 to the end of November 2025. To capture the main sources of common variation in a parsimonious way, we apply Principal Component Analysis (PCA) to the correlation matrix of these returns. PCA decomposes the correlation matrix into orthogonal eigenvectors, interpretable as long–short portfolios, and associated eigenvalues that rank their relative importance. Denoting the correlation matrix by $\boldsymbol{C}_{core}$, its PCA representation is $\boldsymbol{C}_{core} = \boldsymbol{V}\boldsymbol{\Lambda}\boldsymbol{V}^{\top}$, where $\boldsymbol{V}$ contains the eigenvectors and $\boldsymbol{\Lambda}$ is the diagonal matrix of ordered eigenvalues. The variance–covariance matrix is then obtained by rescaling this structure with the volatilities of the core assets, $\boldsymbol{\Sigma}_{core} = \boldsymbol{\Sigma}_{diag}\, \boldsymbol{C}_{core}\, \boldsymbol{\Sigma}_{diag}$.

For the core assets, we retain the first six principal components, which together explain 89% of the total variance. These factors summarize the dominant drivers of co-movements across asset classes while filtering out high-frequency noise. The first component resembles a broad market factor, the second behaves like a duration factor loading positively on sovereign and investment-grade bonds, and the third captures risks associated with emerging markets and commodities. The remaining components explain progressively smaller shares of variance and are less economically interpretable, though they contribute to improving the overall conditioning of the risk model.

Table A4. Statistical risk model based on core asset returns

		Statistical factors						
% of variance explained	Total 89%	Factor 1 56%	Factor 2 16%	Factor 3 7%	Factor 4 4%	Factor 5 3%	Factor 6 3%	Index provider
Core Assets	Tickers	Core Asset weights in PCA orthonormal factors						
Equity Europe EMU	NDDLEURO Index	28.0%	-13.8%	-25.7%	-15.4%	52.8%	-15.3%	1
Equity Europe EMU SC	NCLDEMU Index	28.6%	-13.5%	-20.6%	-1.3%	-3.5%	-23.4%	1
Equity Europe UK	NDDLUK Index	26.3%	-14.8%	-8.7%	-17.5%	-6.7%	-25.3%	1
Equity USA	NDDUUS Index	28.5%	-10.0%	-9.8%	-22.7%	-26.4%	26.0%	1
Equity USA SC	RU20INTR Index	26.4%	-13.5%	-12.3%	-22.9%	-8.1%	23.0%	2
Equity Japan	NDDLJN Index	20.9%	-25.7%	-22.1%	-16.5%	-6.9%	48.3%	1
Equity Emerging Global	NDUEEGF Index	27.1%	-7.9%	31.5%	-19.5%	38.0%	0.0%	1
Bond EUR Sovereign	LEATTREU Index	9.2%	48.7%	-27.1%	-3.2%	18.7%	35.4%	3
Bond EUR Investment Grade	LECPTREU Index	23.5%	30.5%	-14.5%	34.3%	-15.6%	24.9%	3
Bond EUR High Yield	LF88TREU Index	27.7%	-3.2%	1.5%	52.3%	-5.9%	-4.2%	3
Bond USD Sovereign	LUATTRUU Index	1.2%	53.1%	6.4%	-39.7%	19.9%	-5.9%	3
Bond USD IG	LUACTRUU Index	22.5%	37.8%	9.7%	5.8%	-10.9%	-2.0%	3
Bond USD HY	LF89TRUU Index	29.0%	0.2%	9.4%	34.3%	9.4%	-4.6%	3
Bond EMD HC Sov Global	JPGCCOMP Index	27.2%	20.1%	22.4%	1.9%	15.2%	-13.7%	4
Bond EMD LC Sov Global	JGENVUUG Index	24.5%	9.0%	41.6%	-28.6%	18.6%	-16.1%	4
Diversification Real Estate Pan-Europe	TRNHUE Index	25.6%	4.1%	-28.8%	6.7%	25.1%	-39.1%	5
Diversification Commodity Global	BCOMXALT Index	16.8%	-18.4%	53.7%	13.5%	-32.7%	33.3%	3

Note: Risk model estimation based on a PCA model using monthly EUR returns in excess of cash from end of November 2005 through end of November 2025.
Source: 1) MSCI, 2) Russell, 3) Bloomberg, 4) J.P. Morgan, 5) FTSE EPRA. Authors' calculations.

We set the diagonal matrix of $\boldsymbol{\Sigma}_{core}$ to the variances of the core assets, with each variance estimated from weekly EUR-denominated returns over the same sample period.

$$\boldsymbol{\Sigma}_{core} = \begin{bmatrix} var_{core}^{1} & \cdots & cov_{core}^{1,n_N} \\ \vdots & \ddots & \vdots \\ cov_{core}^{n_N,1} & \cdots & var_{core}^{n_N} \end{bmatrix} \tag{A2}$$

with $var_{core}^{i} = cov_{core}^{i,i} = \left(\sigma_{core}^{i}\right)^{2}$ for core asset *i*.

The final variance–covariance matrix $\boldsymbol{\Sigma}$ for core assets and funds will have size $n_A \times n_A$ with $n_A = n_I + n_N$ where n_I is the number of investable funds and n_N the number of non-investable assets, which are exactly the core assets in this setup . This matrix can be written as the sum:

$$\mathbf{\Sigma} = \mathbf{\Sigma}_{systematic} + \mathbf{\Sigma}_{fund\ specific} \tag{A3}$$

and requires the betas of the funds, $\boldsymbol{\beta}^{i}_{fc}$, relative to core assets as well as the specific variance of funds, $var^{i}_{fc} = \left(\sigma^{i}_{fc}\right)^2$.

<u>The systematic variance-covariance</u>, $\mathbf{\Sigma}_{systematic}$, is based on $\mathbf{\Sigma}_{core}$ and on $\boldsymbol{\beta}_{all}$, a $n_A \times n_N$ matrix where the columns have the vectors of exposures i) of core assets to themselves (each vector is 1 on the row for the respective core asset and zero otherwise) and ii) of the funds to the core assets, $\boldsymbol{\beta}^{i}_{fc}$:

$$\mathbf{\Sigma}_{systematic} = \boldsymbol{\beta}_{all}\mathbf{\Sigma}_{core}\boldsymbol{\beta}^{\top}_{all} \tag{A4}$$

with:

Core assets

$$\boldsymbol{\beta}_{all} = \begin{bmatrix} 1 & \dots & 0 \\ \vdots & 1 & \vdots \\ 0 & \dots & 1 \\ \beta^{1,1}_{fc} & \dots & \beta^{1,n_N}_{fc} \\ \vdots & \dots & \vdots \\ \beta^{n_I,1}_{fc} & \dots & \beta^{n_I,n_N}_{fc} \end{bmatrix} \begin{matrix} \text{Core assets} \\ \\ \text{Funds} \end{matrix} \tag{A5}$$

and thus with:

Core assets Funds

$$\mathbf{\Sigma}_{systematic} = \begin{bmatrix} \mathbf{\Sigma}_{core} & \mathbf{\Sigma}^{core,fund}_{systematic} \\ {\mathbf{\Sigma}^{core,fund}_{systematic}}^{\top} & \mathbf{\Sigma}^{fund,fund}_{systematic} \end{bmatrix} \begin{matrix} \text{Core assets} \\ \text{Funds} \end{matrix} \tag{A6}$$

<u>The fund specific variance</u>–<u>covariance matrix</u> assumes that i) the specific risks of individual funds are uncorrelated with each other, and ii) also uncorrelated with the risks of core assets. This matrix is based only on the specific variance of funds relative to their benchmarks, $var^{i}_{fb} = \left(\sigma^{i}_{fb}\right)^2$:

Core assets Funds

$$\mathbf{\Sigma}_{fund\ specific} = \begin{bmatrix} 0 & \dots & 0 & 0 & \dots & 0 \\ \vdots & \ddots & \vdots & \vdots & 0 & \vdots \\ 0 & \dots & 0 & 0 & \dots & 0 \\ 0 & \dots & 0 & var^{1}_{fb} & \dots & 0 \\ \vdots & \ddots & \vdots & \vdots & \ddots & \vdots \\ 0 & \dots & 0 & 0 & \dots & var^{n_I}_{fb} \end{bmatrix} \begin{matrix} \text{Core assets} \\ \\ \text{Funds} \end{matrix}$$

<u>The uncertainty matrix</u> $\mathbf{\Omega}$ is a $n_A \times n_A$ matrix associated with the uncertainty in the expected returns and is an important input for the RPO process. Our choice for this matrix is:

$$\mathbf{\Omega} = \mathbf{\Omega}_{systematic} + \mathbf{\Sigma}_{fund\ specific} \tag{A7}$$

where:

$$\boldsymbol{\Omega}_{systematic} = \boldsymbol{\beta}_{all}\mathrm{diag}(\boldsymbol{\Sigma}_{core})\boldsymbol{\beta}_{all}^{\top} \tag{A8}$$

is consistent with the choice of Yin et al. (2022) to use a diagonal uncertainty matrix with the variances of the assets for which we express our tactical views. However, while a diagonal uncertainty matrix remains appropriate for core assets, we cannot ignore the correlations between funds and core assets. These correlations are essential for translating tactical views on core assets into effective allocations across funds, which is why the additional $\boldsymbol{\Sigma}_{fund\ specific}$ term in equation (A7) is needed. This extension, proposed by Mallouli et al. (2025), illustrates a setting where the uncertainty matrix should not be diagonal. The resulting uncertainty matrix is therefore:

$$\boldsymbol{\Omega}_{systematic} = \begin{array}{c} \quad\text{Core assets} \qquad\qquad \text{Funds} \\ \begin{bmatrix} \mathrm{diag}(\boldsymbol{\Sigma}_{core}) & \boldsymbol{\Omega}_{systematic}^{core,fund} \\ \boldsymbol{\Omega}_{systematic}^{core,fund\ \top} & \boldsymbol{\Omega}_{systematic}^{fund,fund} \end{bmatrix} \end{array} \begin{array}{l} \\ \text{Core assets} \\ \text{Funds} \end{array} \tag{A9}$$

VII. Appendix C

Here we show how to extend the frameworks of Sections II.A and II.B to include a term with aversion to turnover, something often used by quantitative managers in the industry. We consider first the simplest case where equation (2) includes a quadratic turnover penalty term:

$$\max_{\boldsymbol{a}} \left(\boldsymbol{\mu}^{\top}\boldsymbol{a} - \lambda(\boldsymbol{a}^{\top}\boldsymbol{\Sigma}\boldsymbol{a}) - \mathrm{TC} \parallel \boldsymbol{a} - \boldsymbol{a}_{t=0} \parallel_2^2\right) \tag{A10}$$

with TC a measure of the estimated transaction costs scaled by an aversion to turnover, $\boldsymbol{a}_{t=0}$ the current allocation and subject to constraint $\boldsymbol{B}^{\top}\boldsymbol{a} = \boldsymbol{b}$. It is not difficult to show that if we define:

$$\widetilde{\boldsymbol{\Sigma}} \equiv \boldsymbol{\Sigma} + \frac{\mathrm{TC}}{\lambda}\mathbf{I} \tag{A11}$$

$$\widetilde{\boldsymbol{\mu}} \equiv \boldsymbol{\mu} + 2\mathrm{TC}\,\boldsymbol{a}_{t=0} \tag{A12}$$

then the first order conditions and all closed form expressions in Section II.A remain valid under substitution of $\boldsymbol{\mu}$ by $\widetilde{\boldsymbol{\mu}}$ and $\boldsymbol{\Sigma}$ by $\widetilde{\boldsymbol{\Sigma}}$.

For RPO, equation (14) becomes:

$$\max_{\boldsymbol{a}} \left(\boldsymbol{\mu}^{\top}\boldsymbol{a} - \lambda(\boldsymbol{a}^{\top}\boldsymbol{\Sigma}\boldsymbol{a}) - \kappa\sqrt{\boldsymbol{a}^{\top}\boldsymbol{\Omega}\boldsymbol{a}} - \mathrm{TC} \parallel \boldsymbol{a} - \boldsymbol{a}_{t=0} \parallel_2^2\right) \tag{A13}$$

and it is easy to demonstrate that all closed equations in Section II.B remain valid under the substitution of $\boldsymbol{\mu}$ by $\widetilde{\boldsymbol{\mu}}$ and $\boldsymbol{Q}$ by $\widetilde{\boldsymbol{Q}}$ defined by:

$$\widetilde{\boldsymbol{Q}}(t) \equiv 2\lambda\,\boldsymbol{\Sigma} + \frac{\kappa}{t}\,\boldsymbol{\Omega} + 2\mathrm{TC}\,\mathbf{I} \tag{A14}$$

If we replace the quadratic turnover penalty with an ℓ_1 turnover penalty, then we obtain instead:

$$\max_{\boldsymbol{a}} \left(\boldsymbol{\mu}^{\top}\boldsymbol{a} - \lambda(\boldsymbol{a}^{\top}\boldsymbol{\Sigma}\boldsymbol{a}) - \mathrm{TC}\|\boldsymbol{a} - \boldsymbol{a}_{t=0}\|_1\right) \tag{A15}$$

subject to constraint $\boldsymbol{B}^\top \boldsymbol{a} = \boldsymbol{b}$ and where $\| \boldsymbol{x} \|_1 = \sum_i | x_i |$ measures total turnover from the current allocation $\boldsymbol{a}_{t=0}$. The problem is convex and admits a unique optimal solution when $\boldsymbol{\Sigma}$ is positive definite and $\boldsymbol{B}$ has full column rank.

Because of the kink in the norm ℓ_1 of $\boldsymbol{a} - \boldsymbol{a}_{t=0}$ at $a_i = a_{t=0,i}$ and the consequent change of sign in the derivative of $\|\boldsymbol{a} - \boldsymbol{a}_0\|_1$ with respect to $\boldsymbol{a}$ for $a_i > a_{t=0,i}$ when compared with $a_i < a_{t=0,i}$, the first order KKT condition in $\boldsymbol{a}$ becomes:

$$\boldsymbol{\mu} - 2\lambda\, \boldsymbol{\Sigma a} - \boldsymbol{B\delta} - \text{TC}\, \boldsymbol{s} = 0 \tag{A16}$$

with $\boldsymbol{s}$ defined as $s_i^\star = \text{sign}(a_i^\star - a_{t=0,i})$ and an arbitrary $s_i^\star$ if $a_i = a_{t=0,i}$, which for convenience we can set to any number in the range $[-1, +1]$, with this choice having no impact:

$$s_i = \begin{cases} +1, & a_i > a_{t=0,i}, \\ [-1, +1], & a_i = a_{t=0,i}, \\ -1, & a_i < a_{t=0,i}. \end{cases} \tag{A17}$$

If we define the ℓ_1-shrunken unconstrained portfolio:

$$\boldsymbol{a}^{(0)}(\boldsymbol{s}) \equiv \frac{1}{2\lambda}\, \boldsymbol{\Sigma}^{-1}\, (\boldsymbol{\mu} - \text{TC}\, \boldsymbol{s}) \tag{A18}$$

for any fixed sign vector $\boldsymbol{s}$, then, repeating the algebra (10) to (13) with $\boldsymbol{\mu}$ replaced by $\boldsymbol{\mu} - \text{TC}\boldsymbol{s}$ yields the same constrained solution structure:

$$\boldsymbol{a}(\boldsymbol{s}) = \boldsymbol{a}^{(0)}(\boldsymbol{s}) - \sum_{k=1}^{m} \frac{1}{2\lambda} \boldsymbol{\Sigma}^{-1} \boldsymbol{B}_k \delta_k\, (\boldsymbol{s})) \tag{A19}$$

with the multipliers $\delta_k(\boldsymbol{s})$ jointly determined by the linear equation:

$$(\boldsymbol{B}^\top \boldsymbol{\Sigma}^{-1} \boldsymbol{B})\boldsymbol{\delta}(\boldsymbol{s}) = 2\lambda(\boldsymbol{B}^\top \boldsymbol{a}^{(0)}(\boldsymbol{s}) - \boldsymbol{b}) \tag{A20}$$

Equations (A18), (A19) and (A20) are exact for the given $\boldsymbol{s}$. In its final form, the optimal sign vector $\boldsymbol{s}^\star$ is the one derived from the resulting optimization solution $\boldsymbol{a}(\boldsymbol{s}^\star)$. However, $\boldsymbol{s}^\star$ is not known ex-ante. Under ℓ_1 we must first solve for $\boldsymbol{a}^\star$, then identify $\boldsymbol{s}^\star$, and then plug $\boldsymbol{s}^\star$into (A18), (A19) and (A20) to obtain a decomposition similar to (10) to (12), i.e., an unconstrained $\boldsymbol{a}^{(0)}(\boldsymbol{s}^\star)$ plus one correction per constraint.

To solve the optimization problem numerically we can re-write it as:

$$\begin{aligned} \max_{\boldsymbol{a},\boldsymbol{\tau}} \quad & \boldsymbol{\mu}^\top \boldsymbol{a} - \lambda\, \boldsymbol{a}^\top \Sigma \boldsymbol{a} - \text{TC}\, \mathbf{1}^\top \boldsymbol{\tau} \\ \text{s.t.} \quad & \boldsymbol{a} - \boldsymbol{a}_0 \le \boldsymbol{\tau}, -(\boldsymbol{a} - \boldsymbol{a}_0) \le \boldsymbol{\tau}, \boldsymbol{B}^\top \boldsymbol{a} = \boldsymbol{b}, \boldsymbol{\tau} \ge 0 \end{aligned} \tag{A21}$$

The solver returns the unique optimal active portfolio $\boldsymbol{a}^\star$, the Lagrange multipliers $\boldsymbol{\delta}^\star$ for $\boldsymbol{B}^\top \boldsymbol{a} = \boldsymbol{b}$, and the multipliers for the turnover, which encode the sub-gradient signs of $\boldsymbol{s}^\star$. From these outputs one recovers the optimal sign vector $\boldsymbol{s}^\star$ by using $\boldsymbol{a}^\star$ in (A17).

The ℓ_1 penalty inserts a data-driven threshold TC in the returns vector. Assets with weak returns relative to TC stay at $a_{t=0,i}$ (no trade) while large returns in absolute terms flip $s_i^\star$ to $+1$ or -1 respectively depending on the sign of the returns, producing sparse and intuitive trades. The hedging portfolios for the constraints remain the same minimum variance style sub-portfolios as with the interpretation of (12).

For the RPO problem under ℓ_1 turnover constraint we have:

$$\max_{\boldsymbol{a}} \left(\tilde{\boldsymbol{\mu}}^\top \boldsymbol{a} - \frac{1}{2} \boldsymbol{a}^\top \boldsymbol{Q}(t) \boldsymbol{a} - \mathrm{TC} \parallel \boldsymbol{a} - \boldsymbol{a}_0 \parallel_1 \right) \tag{A22}$$

subject to $\boldsymbol{B}^\top \boldsymbol{a} = \boldsymbol{b}$ and with the same robust $\boldsymbol{Q}(t)$ as in (18). For a fixed $\boldsymbol{s}$, the unconstrained robust portfolio is:

$$\boldsymbol{a}^{(0)}(t, \boldsymbol{s}) = \boldsymbol{Q}(t)^{-1} (\tilde{\boldsymbol{\mu}} - \mathrm{TC}\, \boldsymbol{s}) \tag{A23}$$

and the constrained solution is equivalent to (21):

$$\boldsymbol{a}^\star(t, \boldsymbol{s}) = \boldsymbol{a}^{(0)}(t, \boldsymbol{s}) - \boldsymbol{Q}(t)^{-1} \boldsymbol{B} (\boldsymbol{B}^\top \boldsymbol{Q}(t)^{-1} \boldsymbol{B})^{-1} (\boldsymbol{B}^\top \boldsymbol{a}^{(0)}(t, \boldsymbol{s}) - \boldsymbol{b}) \tag{A24}$$

with multipliers from:

$$(\boldsymbol{B}^\top \boldsymbol{Q}(t)^{-1} \boldsymbol{B}) \boldsymbol{\delta}(t, \boldsymbol{s}) = (\boldsymbol{B}^\top \boldsymbol{a}^{(0)}(t, \boldsymbol{s}) - \boldsymbol{b}) \tag{A25}$$

As before, $\boldsymbol{s}^\star$ must be determined numerically by solving the RPO first, but now with $\boldsymbol{Q}(t)$. Then the decomposition follows by using $\boldsymbol{s}^\star$ into (A24) and (A25). The structure of the results in II.B remains unchanged replacing $\boldsymbol{\mu}$ by $\boldsymbol{\mu} - \mathrm{TC}\boldsymbol{s}^\star$.

VIII. Disclosure statement

The authors report no conflicts of interest. The authors alone are responsible for the content and writing of the paper.

IX. References

Bass, R., Gladstone, S. and Ang, A. Total Portfolio Factor, Not Just Asset, Allocation, *The Journal of Portfolio Management*, 2017, 43(5), 38–53.

Ben-Tal, A. and Nemirovski, A. Robust Convex Optimization, *Mathematics of Operations Research*, 1998, 23, 769–805.

Ben-Tal, A. and Nemirovski, A. Robust Solutions of Uncertain Linear Programs, *Operations Research Letters*, 1998, 25, 1-13.

Ben-Tal, A., El Ghaoui, L. and Nemirovski, A., *Robust Optimization*., 2009 (Princeton University Press, Princeton).

Boyd, S. and Vandenberghe, L., *Convex Optimization*, 2004 (Cambridge University Press).

Ceria, S. and Stubbs, R. A., Incorporating Estimation Errors into Portfolio Selection: Robust Portfolio Construction. *J. Asset Manage.*, 2006, 7(2), 109–127.

Grinold, R. C. and Kahn, R. N., *Active Portfolio Management*, 2000 (McGraw Hill).

Hastie, T., Tibshirani, R. and Wainwright, M.. *Statistical Learning with Sparsity: The Lasso and Generalizations*, 2015 (Chapman and Hall/CRC).

Heckel, T., Leote de Carvalho, R., Lu, X. and Perchet, R. Insights into Robust Portfolio Optimization: Decomposition into Mean-variance and Risk-based Portfolios. *Journal of Investment Strategies*, 2016, 6 (1), 1-24.

Issaoui, T., Perchet, R., Retière, O., Soupé, F., Yin, C. and Leote de Carvalho, R., Mass Customization of Asset Allocation. *J. Investing*, 2021, 31(3), 73-97.

Lobo, M. S., Vandenberghe, L., Boyd, S. and Lebret, H. Applications of Second Order Cone Programming., *Linear Algebra and Its Applications*, 1998, 284(1–3), 193–228.

Mallouli, M. A., Perchet, R., Soupé, F. and Leote de Carvalho, R., Multi-Asset Portfolios with Active and Passive Funds: A Robust Optimization Framework. *Working paper SSRN*, 2025.

Meucci, A., *Risk and Asset Allocation*, 2005 (Springer).

Somefun, K., Perchet, R., Yin, C. and Leote de Carvalho, R. Allocating to Thematic Investments. *Financial Analysts Journal*, 2022, 79(1), 18–36.

Scherer, B. Can Robust Portfolio Optimization help to Build Better Portfolios? *Journal of Asset Management*, 2006, 7(6), 374-387.

Tibshirani, R. Regression Shrinkage and Selection via the Lasso. *Journal of the Royal Statistical Society: Series B (Methodological)*, 1996, 58(1), 267–88.

Tütüncü, R.H. and König M. Robust Asset Allocation. *Annals of Operations Research*, 2004, 132(1-4), 157-187.

Yin, C., Perchet, R. and Soupé, F. A Practical Guide to Robust Portfolio Optimization. *Quantitative Finance*, 2021, 21(6), 911-928.