

---

**Research article****Graph aware adaptive tracking-error optimization with wavelet-principal component analysis features and proportional-integral control (GATE-WPCA-PI)****Muhammad Hilal Alkhudaydi\* and Yehya M. Althobaity**

Department of Mathematics and Statistics, College of Science, Taif University, P.O. Box 11099, Taif 21944, Saudi Arabia

\* **Correspondence:** Email: mh.ayedh@tu.edu.sa.

**Abstract:** We introduced GATE-WPCA-PI—geometry-aware, tracking-error-controlled allocation with wavelet principle component analysis features and a proportional-integral controller—a practical portfolio construction framework that linked multi-scale market geometry to explicit, out-of-sample risk targeting. At each rebalance, the daily returns were embedded in a multi-resolution wavelet feature space and compressed via principal component analysis to form a similarity kernel. A simple discriminative-power score gated the optimizer: when the cross section was heterogeneous, the feature geometry was activated; when it was homogeneous, the method reverted to a correlation-only view. Allocations were obtained from an implementable mean–variance surrogate with (i) a geometry penalty that discouraged concentration in highly similar assets, (ii) quadratic and absolute turnover costs, (iii) an entropy floor, and (iv) standard long-only, budget, and sleeve caps. A proportional-integral (PI) law treated the tracking error (TE) as a controllable state and steered realized TE toward a feasible band under trading frictions.

**Keywords:** GATE-WPCA-PI; geometry-aware; multi-scale market geometry; wavelet features; principal component analysis; similarity kernel; mean-variance surrogate; turnover costs; entropy floor; tracking error control

**Mathematics Subject Classification:** 62P05

---

**Model name and components.** In the geometry-aware allocation (GATE) component, assets are represented as data-driven similarity graphs. The portfolio is regularized to ensure smooth variation across this graph, allowing strongly assets to carry coherently related weights while permitting differentiation where the data warrants it. The tracking-error-controlled segment introduces an explicit target for active risk relative to a benchmark and penalizes deviations from this target, ensuring that the scale of active bets is intentional rather than incidental. The wavelet principal component analysis WPCA block provides multi-resolution features by computing

wavelet summaries of each asset's recent return history and reducing them using the principal components. Distances in this low-dimensional feature space define an alternative, more structural notion of similarity, which is blended with (or substituted for) correlation when empirically more discriminative, thereby making the effective geometry adaptive to the market regimes. Finally, the proportional-integral (PI) block functions as a proportional-integral feedback controller that adjusts the tracking error penalty in a closed loop, ensuring that the realized tracking-error adheres to the time-varying target despite constraints, costs, and estimation noise. Collectively, these elements yield allocations that are smooth in homogeneous markets, selective when a genuine structure exists, and implementable under realistic trading friction and capacity limits. For further explanation on how the model mechanism works, refer to Figures 17 and 18 in Appendix C and D. The figures delineate the sequential stages of the process, commencing with feature extraction and culminating in the transformation of these features into graphical representations. They elucidate how the gating mechanism selectively identifies pertinent information, thereby ensuring that only the most relevant features progress to subsequent stages. Furthermore, the figures demonstrate the application of WPCA and PI to achieve dimensionality reduction while maintaining a consistent representation, irrespective of the input order. Additionally, they depict the transformation of raw features into graph mappings, illustrating the model's capacity to efficiently manage complex relationships. This visual representation aids in elucidating the conversion of feature vectors into a graph, wherein the nodes and edges possess meaningful attributes that reflect the data patterns. Collectively, the figures offer a lucid visual guide to the model's operation, thereby facilitating the comprehension of the intricate processes involved.

## 1. Introduction

Shrinkage-based geometry-aware risk estimation combined with explicit tracking-error control delivers consistent outperformance over both classical constrained mean variance and naïve diversification under realistic box and sleeve caps. Embedding a tracking-error penalty within the mean variance improves the Sharpe ratios relative to both the unconstrained optimizer and the reference portfolio, while enhanced tracking-error management that couples sophisticated shrinkage with multivariate generalized autoregressive conditional heteroskedasticity (GARCH) dominates the sample covariance benchmark [1, 2]. The critical process of estimating large covariance or precision matrices for asset allocation is examined alongside the theoretical and empirical effects of leveraging constraints and covariance shrinkage in portfolio construction, revealing that a well-chosen shrinkage method outperforms arbitrary leverage constraints, with evidence from diverse setups, objectives, and datasets highlighting the dynamic correlation shrinkage (DCC-NL) estimator's ability to enhance risk-adjusted efficiency in large portfolios under moderate leverage [3, 4]. In index-proximate settings, diversifying model risk while budgeting tracking error yields robust outperformance over cap-weighting, and constrained minimum-variance portfolios achieve lower variance and materially higher Sharpe; on average, a 32.5% increase over  $1/N$  across six datasets, with low turnover [5, 6], where 32.5% represents the average percentage improvement in the Sharpe ratio compared to the naively diversified  $1/N$  portfolio. Relative to naïve diversification, mean-variance and timing-aware variants often prevail even when transaction costs are substantial, and across assets, the Black–Litterman framework outperforms mean-variance and  $1/N$  with lower drawdowns and turnover [7, 8].

These gains depend on stable inputs and careful implementation. Mean-variance remains sensitive

to parameter estimation, and turnover and transaction costs can erode advantages if models are not sufficiently regularized [1, 7]. Methodological gaps persist for precision-matrix shrinkage in singular cases, and for identifying signal eigenstructures when the asset dimension is not small relative to the sample size. This can depend on proxy choices and transaction cost assumptions, and external validity may be limited by the universe size [3, 4]. Alternative weighting with the tracking-error discipline can incur short-horizon downsides and periods of underperformance when cap-weighted indices rally significantly [5]. Evidence from Black and Litterman calls for tests with alternative return inputs beyond historical estimates [8]. Finally, some abstracts reported no explicit limitations, specifically for improved tracking-error management and constrained minimum variance, indicating the need for further validation under varied universes and cost regimes [2, 6].

Evidence from control theory and portfolio studies shows that explicit closed-loop tracking-error designs achieve superior adherence to a tracking error (TE) target path while preserving or enhancing performance compared with static penalty-based formulations. Feedback and model predictive control architectures deliver precise path following, stabilize deviations, and guarantee optimal closed-loop behavior even when the target is not in equilibrium [9, 10]. Constraint satisfaction and observability-focused designs further strengthen the path accuracy under instability or incomplete information [11, 12]. In financial applications, dynamic tracking-error processes consistently outperform rigid constant-risk approaches, with benchmark-relative risk matching yielding better returns in volatile regimes [13, 14]. Index-tracking models that integrate transaction costs and tracking efficiency deliver improved accuracy and diversification relative to mean-variance and traditional formulations [15, 16]. The limitations remain, including sensitivity to estimation, high turnover costs, reliance on approximate constraint satisfaction in near-singular settings, precomputed feedback gains, and potential instability in observability-driven optimization [9–15].

Relative to standard position-smoothing frameworks, the literature shows that a proportional-integral tracking-error controller with a turnover cap and volatility band offers a more effective means of balancing performance and stability under transaction costs, as it synthesizes three proven mechanisms: Optimal no-trade regions that scale with costs and volatility to cut turnover and preserve alpha [17, 18], volatility smoothing that dampens transients, reduces trading, and enhances risk-adjusted returns when coupled with ex-ante cost modeling [19, 20], and large-scale portfolio formulations that directly embed costs and model uncertainty to regularize allocations and improve net efficiency [21–23]. These approaches consistently demonstrate that explicit control of trading intensity and realized risk yields higher Sharpe ratios, lower turnover, and superior cost-return trade-offs compared to static smoothing or naïve rebalancing [24]. Proportional-cost models often lack closed-form tractability, volatility forecasts face accuracy-smoothness trade-offs, and dynamic extensions remain complex, with several studies reporting no explicit limitations, leaving questions about external validity across universes and cost regimes [17–24].

Geometry-aware approaches show promising improvements in portfolio diversification, with growing evidence that they enhance diversification quality, weight stability, and risk-adjusted performance compared to traditional correlation-only and hierarchical clustering models. By embedding the intrinsic geometric structure of covariance matrices, these methods achieve more balanced risk contributions, smoother allocations, and higher out-of-sample Sharpe ratios while mitigating estimation noise and drawdowns [25–32]. The study by [25] developed a geometric framework for covariance dynamics that considers the manifold properties of covariance matrices and outperforms conventional estimators across multiple performance metrics. The study by [31] demonstrated that machine learning-based portfolios can be more robust to covariance

misspecification, with hierarchical approaches offering advantages over traditional optimization techniques. The article by [33] showed that hierarchical clustering-based portfolios achieve statistically superior risk-adjusted performance compared with classical mean–variance models, and this is further supported by evidence that regularization methods substantially improve Sharpe ratios and enhance portfolio robustness under estimation uncertainty [27]. Despite these advances, geometry-aware and similarity-driven approaches are computationally demanding, sensitive to hyperparameter selection, and sometimes prone to instability under market regime shifts or noisy data where inputs are dominated by non-persistent, non-informative variations that produce spurious covariances and unstable optimized weights, which geometry-aware methods aim to regularize [25, 29, 30, 32, 33]. Although the empirical record is promising, comprehensive validation under realistic transaction costs, liquidity frictions, and diverse market conditions remains limited, underscoring the need for further research to confirm their superiority in all environments [26, 28, 31].

Wavelet–PCA feature geometry can improve outcomes in heterogeneous cross-sections when activated by a data-driven switch, whereas deactivation in homogeneous regimes avoids injecting redundant variance and computational overhead, indicating gains in classification accuracy, stability, and efficiency contingent on the domain and data characteristics. In heterogeneous settings, multi-scale wavelet kinetics and pharmacokinetic heterogeneity features extracted from dynamic contrast-enhanced magnetic resonance imaging (DCE-MRI) elevate risk stratification and predictive power, illustrating the value of engaging geometry when spatial–temporal variation is rich [34, 35]. Temporal wavelet descriptors further bolster treatment-outcome prediction [36], which preserves spectral signatures and filters anomalies with an accuracy comparable to or better than PCA at lower computational costs [37], which improves denoising and peak detection in challenging regions [38]. Beyond biomedicine, online heterogeneous data analysis via principal composite kernel feature analysis increases diagnostic accuracy while reducing computation requirements [39], and spectral-transformation transfer learning on heterogeneous feature spaces can reduce the proportion of incorrect out-of-sample predictions (the misclassification rate) by up to 50% when related source examples are carefully selected [40]. Complementary evidence shows that multi-featured, geometry-aware homogenization attains high-fidelity predictions with substantial speedups, and that combined feature extraction with judicious selection strengthens discrimination in fault diagnosis [41, 42]. The central advantage is the adaptive, context-sensitive activation of features—switching geometry on when heterogeneity is present and off when it is not—yielding more precise and economical representations than static always-on or always-off strategies; however, the magnitude of improvement depends on hyperparameter choices, data quality, and regime stability, and purely data-driven selection can underperform whole-domain baselines when cross-sectional structure is effectively homogeneous [36–43].

Across diverse empirical settings, most headline performance gains prove fragile once data snooping and sampling uncertainty are controlled rigorously. Search-adjusted inference and window-robust evaluation frequently attenuate or eliminate apparent improvements, with evidence indicating that a substantial fraction (of the order of half) of positive findings may be false discoveries [44, 45]. Classic examinations of technical rules and asset-pricing tests show how extensive rule mining and selective exploration inflate type-I errors and how even weak correlations between sorting characteristics and security-level statistics can induce spurious rejection of the null [46–48]. Therefore, robust research design centers on disciplined model selection, repeated cross-validation, bootstrap-based reality checks, multiplicity control, and transparent reporting of fold-wise result distributions complemented by out-of-sample tests that are robust to window choices [49–51]. In

small-sample settings, variance-corrected cross-validation with tailored multiple-testing adjustments helps curb selection bias, albeit with power trade-offs, whereas cost-aware evaluation further erodes residual gains, underscoring that economic significance must be accompanied by statistical durability [52, 53]. The resulting implication is clear: only strategies whose improvements persist under search-adjusted, fold-wise resampling and remain economically meaningful after frictions should be regarded as genuinely robust; such cases are an exception rather than the rule [44–51, 53].

**Contribution:** This paper introduces GATE–WPCA–PI a unified allocation framework that combines geometry learning, risk modeling, and feedback control under realistic trading constraints. In response to the reviewer, we summarize the contributions in two parts.

**Theoretical contributions:** (i) *Geometry construction and gating.* We define a hybrid notion of cross-sectional geometry by combining a correlation-induced similarity kernel with a feature kernel built from multi-scale wavelet summaries (WPCA), and we introduce a data-driven gate that blends or suppresses the feature geometry using a discriminative-power score together with first-principal-component dominance (PC1), thereby activating geometry in heterogeneous universes and defaulting to correlation-only structure in homogeneous ones. (ii) *Geometry-aligned regularization and covariance.* We couple the learned geometry to the allocation problem in two complementary ways: A graph-Laplacian penalty regularizes weights over a sleeve-restricted  $k$ -nearest-neighbor ( $k$ -NN) graph, and a similarity-aware covariance model rescales the raw correlation matrix entrywise by the learned kernel, yielding a covariance specification that is aligned with the same geometry used for regularization. (iii) *Active-risk control formulation.* We treat TE, defined as the standard deviation of active returns relative to a benchmark, as a controllable state and embed a PI feedback law into the objective via an adaptive TE-penalty weight, so that realized TE is steered toward a time-varying, feasible target while respecting a volatility band and implementability constraints.

**Practical contributions:** (i) *End-to-end implementable engine.* We provide a complete walk-forward pipeline that re-estimates parameters monthly with no look-ahead, applies weights with next-day execution, and evaluates net performance under proportional transaction costs, explicit turnover caps, and sleeve constraints. (ii) *Numerical and operational robustness.* The covariance used in optimization is symmetrized and maintained positive semi-definite, the refinement stage uses warm starts with a small number of constrained updates, and post-processing enforces volatility targeting, one-dimensional TE alignment, and turnover capping in a transparent and auditable manner. (iii) *Institutional evaluation protocol.* We report active performance relative to strongest reference baseline and assess statistical validity using White–Hansen superior predictive ability (SPA) tests [54] and a model confidence set (MCS), complemented by bootstrap-based robustness checks, forward TE diagnostics, and transaction-cost and capacity analyses; empirically, the framework delivers statistically defensible improvements in a heterogeneous exchange-traded fund (ETF) universe when feature geometry is informative, and defaults conservatively to correlation-based behavior in a more homogeneous Sector universe when it is not.

**Outline:** Section 2 formalizes the model by detailing the similarity kernels, graph Laplacian regularization, similarity-aware covariance, momentum-based alpha, TE target schedule, and the single-period objective with constraints. It also elaborates on the computation process, including feature extraction via wavelet-PCA, hybrid geometry switch, warm-start and projected updates, volatility targeting, one-dimensional TE alignment, and turnover capping. Section 3 provides a description of the data, universes, mandates, and baselines such as equal weighting (EW), Markowitz with shrinkage, Risk Parity/hierarchical risk parity (HRP)/nested clustered optimization (NCO), while fixing costs and constraints. The section also presents the main results of our model.

Additionally, it offers robustness analysis through SPA/MCS, bootstraps, cross-validation folds, ablations, regime-conditioned confidence intervals, forward TE tracking, as well as cost/capacity and break-even sensitivity, and dynamic sleeve allocation analysis. Section 4 discusses the implications, limitations, and future work, and concludes the paper.

## 2. Methodology: The GATE-WPCA-PI model

The model treats a cross-section of assets as a network in which each pair of securities is connected by a data-driven notion of similarity. Within this network framework, the portfolio weights exhibit characteristics similar to those of a smooth field. Assets that are proximate in appearance tend to receive analogous allocations, except in cases in which robust empirical evidence necessitates differentiation. This is implemented by a geometric regularizer that penalizes abrupt changes in weights across strong links, while a similarity-aware covariance “renormalizes” risk so that co-movements among genuine assets are emphasized and spurious ones are de-emphasized. Risk-taking is governed by a TE controller that functions like a thermostat: it raises or lowers the effective “temperature” of the active bets so that the realized TE follows a target despite changes in market conditions and trading frictions. Volatility targeting sets the overall scale of risk, and a turnover cap acts as a physical friction on portfolio motion, preventing overly rapid reallocation. Together, these elements produce allocations that are stable when the market is smooth, responsive when the conditions change, and implementable under realistic cost and capacity constraints.

The geometry of the model is not fixed; it is learned from the data using wavelets and principal components. Wavelets decompose each return series into multiple timescales and summarize the distribution of energy\* and irregularity across those scales, followed by the principal components then extraction of a low-dimensional representation that captures persistent shapes rather than mere variance. Distances in that feature space define an alternative notion of similarity that is especially informative in heterogeneous universes when industries, regions, or risk premiums evolve differently. A simple switch blends or replaces correlation-based geometry with this feature-based geometry, depending on how discriminative the multi-scale structure appears and the dominance of the first principal component. When heterogeneity is strong, the model sharpens its notion of “neighbors” and allocates along truly differentiated directions. When the cross section is homogeneous, it returns to a plain correlation to avoid overfitting. This adaptive coupling of multi-scale feature learning with geometry-aware regularization and closed-loop risk control gives the model its robustness and practical edge in out-of-sample portfolio optimization.

We consider a discrete set of rebalancing dates  $t \in \mathcal{T}$ . At each date  $t$ , portfolio decisions are based on a trailing window of past returns,

$$\mathcal{W}_t = \{t - L + 1, \dots, t\},$$

where  $L \in \mathbb{N}$  is a fixed lookback measured on trading days. The investable universe contains  $N$  assets and the arithmetic return observed on day  $\tau$  is the vector  $\mathbf{r}_\tau \in \mathbb{R}^N$ . Portfolio allocations applied immediately after the rebalancing date  $t$  are represented by the weight vector  $\mathbf{w}_t \in \mathbb{R}^N$ . The trailing window  $\mathcal{W}_t$  localizes the estimation to the recent regime so that second moments and signals adapt

\*In the wavelet features, “energy” denotes the  $\ell_2$  power of a subband, i.e., the sum of squares of its discrete wavelet coefficients; for a coefficient vector  $c = (c_k)$  we record the normalized average energy  $E = \frac{1}{|c|} \sum_k c_k^2$ , and by Parseval’s identity the sum of subband energies equals the sample variance over the window. This is signal-processing terminology rather than a physical quantity and serves as an amplitude proxy indicating how much variation is carried at each scale.

smoothly to changing market conditions, and the choice of a fixed  $L$  balances the variance reduction against the risk of stale estimates and is kept constant across the sample to avoid look-ahead bias.

The following conventions are used throughout:  $t$  indexes rebalancing dates,  $i \in \{1, \dots, N\}$  indexes assets, and  $g \in \mathcal{G}$  indexes sleeves. Scalar parameters include the estimation window length  $L$ , the within-sleeve neighborhood size  $k$  used to form the  $k$ -nearest-neighbor graph, and the Laplacian penalty weight  $\lambda_{\text{lap}}$ . Portfolio weights  $\mathbf{w}$  and alpha signals  $\alpha$  are vectors in  $\mathbb{R}^N$ , while objects such as the covariance  $\Sigma$  and the graph Laplacian  $\mathbf{L}$  are matrices in  $\mathbb{R}^{N \times N}$ . The symbol  $\mathbf{1} \in \mathbb{R}^N$  denotes the all-ones vector. Inequalities involving vectors are interpreted componentwise; for example,  $l \leq \mathbf{w} \leq u$  means  $l \leq w_i \leq u$  for all  $i = 1, \dots, N$ .

The feasibility of portfolio weights is enforced by full investment, element-wise box bounds, and sleeve (group) exposure caps. The admissible set is

$$\Omega = \left\{ \mathbf{w} \in \mathbb{R}^N : \mathbf{w}^\top \mathbf{1} = 1, \quad l \leq \mathbf{w} \leq u, \quad \sum_{i \in g} w_i \leq c_g \quad \forall g \in \mathcal{G} \right\},$$

where  $\mathcal{G}$  is a fixed partition of the universe into sleeves such as equity, duration, and real-asset buckets, and  $c_g \in (0, 1]$  are sleeve-level caps. The box bounds  $l, u \in \mathbb{R}$  prevent vanishingly small or overly concentrated positions and make the projected updates well-posed, while sleeve caps control the concentration risk at a coarser level, consistent with practical mandates and liquidity constraints. A benchmark portfolio  $\mathbf{w}^b \in \Omega$  is specified to define tracking-error risk and anchor active bets, where  $\mathbf{w}^b$  is the EW allocation across the risky sleeves with the residual allocated to cash, which yields a neutral and transparent baseline and interprets the tracking-error target in both absolute and relative terms. We clarify that sleeve gating is applied only to the graph regularizer that defines the Laplacian term  $\mathbf{w}^\top \mathbf{L}_t \mathbf{w}$ . The covariance matrix  $\Sigma_t$  is estimated on the full cross-section and retains all cross-sleeve covariances. The intention is to smooth weights within economically coherent groups while preserving a complete variance model. In a robustness variant we down-weight rather than zero cross-sleeve edges and observe that tracking-error alignment and risk-adjusted performance are materially unchanged.

From the windowed return matrix  $\mathbf{X}_t = [\mathbf{r}_\tau]_{\tau \in \mathcal{W}_t} \in \mathbb{R}^{L \times N}$ , where the row indexes date in the trailing window and column indexes assets, we estimate a covariance matrix using Ledoit–Wolf shrinkage, denoted  $\Sigma_t^{\text{raw}} \geq 0$ . Shrinkage is used to stabilize the second-moment estimates in finite samples and under cross-sectional collinearity, which is particularly important when  $N$  is non-negligible relative to  $L$ . Let  $\sigma_t = \sqrt{\text{diag}(\Sigma_t^{\text{raw}})} \in \mathbb{R}^N$  be the vector of marginal volatilities and  $\mathbf{D}_t = \text{diag}(\Sigma_t^{\text{raw}}) \in \mathbb{R}^{N \times N}$  be the diagonal matrix of the variance. The corresponding correlation matrix is

$$\mathbf{R}_t^{\text{raw}} = \mathbf{D}_t^{-1/2} \Sigma_t^{\text{raw}} \mathbf{D}_t^{-1/2} \in \mathbb{R}^{N \times N},$$

which rescales covariances to unit variances, such that cross-sectional similarity is not dominated by scale. To transform correlations into a geometry suitable for graph construction, we define the correlation distance

$$\mathbf{D}_t^{\text{corr}} = [d_{ij}]_{i,j=1}^N, \quad d_{ij} = \sqrt{2(1 - R_{t,ij}^{\text{raw}})},$$

and a Gaussian similarity kernel

$$\mathbf{S}_t^{\text{corr}} = \exp(-\mathbf{D}_t^{\text{corr}}/\tau_t^{\text{corr}}),$$

with bandwidth  $\tau_t^{\text{corr}}$  set as the median of the off-diagonal entries of  $\mathbf{D}_t^{\text{corr}}$ . The median choice provides a robust, data-adaptive scale, so that neither very tight nor very diffuse graphs are formed

when the regimes change. To respect the mandate structure, we set the cross-sleeve entries of  $\mathbf{S}_t^{\text{corr}}$  to zero prior to graph construction, such that only economically substitutable assets are linked.

For each asset  $i$  we complement the correlation geometry with the multi-resolution features extracted from the univariate series  $\{r_{\tau,i}\}_{\tau \in \mathcal{W}_i}$ . We apply discrete wavelet transforms across admissible families and decomposition levels, which separate movements by timescale and, thus, capture both short- and long-term structures. From each coefficient vector, we computed the statistical mean, standard deviation, mean absolute value, interquartile range, average energy, Shannon entropy, and concatenation to yield a feature vector  $f_{i,t} \in \mathbb{R}^{p_t}$  of length  $p_t$  which summarizes the distribution of energy across scales as well as shape. Stacking all assets gives the feature matrix  $\mathbf{F}_t \in \mathbb{R}^{N \times p_t}$ , which we standardized column-wise and reduced by principal components to  $\mathbf{Z}_t \in \mathbb{R}^{N \times q_t}$  while retaining a fixed fraction of explained variance, which prevents overfitting and ensures that distances are computed in a parsimonious, denoised space. We monitor the dominance of the first component by the scalar

$$\rho_{1,t} = \frac{\text{explained variance of PC1}}{\text{total explained variance}} \in (0, 1),$$

because  $\rho_{1,t}$  close to one indicates a homogeneous cross-section, where the geometry beyond the correlation is less informative. Euclidean geometry in the reduced space defines

$$\mathbf{D}_t^{\text{feat}} = \text{squareform}(\text{pdist}(\mathbf{Z}_t)) \in \mathbb{R}^{N \times N}, \quad \mathbf{S}_t^{\text{feat}} = \exp(-\mathbf{D}_t^{\text{feat}}/\tau_t^{\text{feat}}).$$

Let  $\mathbf{Z}_t \in \mathbb{R}^{N \times q}$  collect the  $q$ -dimensional feature vectors at date  $t$  by rows,  $\mathbf{z}_{i,t}^\top$  for  $i = 1, \dots, N$ . Define the condensed vector of all pairwise Euclidean distances

$$\text{pdist}(\mathbf{Z}_t) := (\|\mathbf{z}_{i,t} - \mathbf{z}_{j,t}\|_2)_{1 \leq i < j \leq N} \in \mathbb{R}^{N(N-1)/2}.$$

Let  $\text{vech}^{-1} : \mathbb{R}^{N(N-1)/2} \rightarrow \mathbb{R}^{N \times N}$  be the inverse half-vectorization that fills the upper triangle of a symmetric, hollow matrix and mirrors it to the lower triangle. We set

$$\mathbf{D}_t^{\text{feat}} := \text{squareform}(\text{pdist}(\mathbf{Z}_t)) := \text{vech}^{-1}(\text{pdist}(\mathbf{Z}_t)) \in \mathbb{R}^{N \times N}.$$

Equivalently,

$$(\mathbf{D}_t^{\text{feat}})_{ii} = 0, \quad (\mathbf{D}_t^{\text{feat}})_{ij} = (\mathbf{D}_t^{\text{feat}})_{ji} = \|\mathbf{z}_{i,t} - \mathbf{z}_{j,t}\|_2 \quad (i \neq j).$$

In words,  $\text{vech}^{-1}(\text{pdist}(\mathbf{Z}_t))$  reconstructs the full  $N \times N$  pairwise-distance matrix from the condensed vector of distances; this is the mathematical counterpart of the code tokens “`pdist`” and “`squareform`”. The feature-similarity kernel used in the optimizer is then

$$\mathbf{S}_t^{\text{feat}} = \exp(-\mathbf{D}_t^{\text{feat}}/\tau_t^{\text{feat}}),$$

with  $\tau_t^{\text{feat}}$  chosen as the median off-diagonal entry of  $\mathbf{D}_t^{\text{feat}}$ .

With a bandwidth  $\tau_t^{\text{feat}}$  equal to the median off-diagonal element of  $\mathbf{D}_t^{\text{feat}}$ , to quantify when feature geometry is informative, we define the discriminativeness functional

$$\psi(\mathbf{D}) = \frac{\text{sd}\{D_{ij} : i < j\}}{\text{mean}\{D_{ij} : i < j\}},$$

and  $\psi(\mathbf{D}_t^{\text{feat}})$  into a score  $d_t^{\text{norm}} \in [0, 1]$  via min-max normalization across the rebalancing grid (or a rolling calibration window). Here,  $\psi(\mathbf{D})$  is defined as the coefficient of variation of the off-diagonal

distances. A higher  $d_t^{\text{norm}}$  indicates that  $\mathbf{D}_t^{\text{feat}}$  has a wider and more structured spread of inter-asset dissimilarities. We also form blend weights

$$\omega_t^{\text{feat}} = \frac{\psi(\mathbf{D}_t^{\text{feat}})}{\psi(\mathbf{D}_t^{\text{feat}}) + \psi(\mathbf{D}_t^{\text{corr}})}, \quad \omega_t^{\text{corr}} = 1 - \omega_t^{\text{feat}}.$$

Thus, when the feature distances are more discriminative than the correlation distances, the combined similarity yields a greater mass of  $\mathbf{S}_t^{\text{feat}}$ .

To avoid injecting noisy geometry into homogeneous universes and to match the implementable code, we adopt a hybrid switching rule with two thresholds: A discriminativeness quantile  $q_\ell \in (0, 1)$  applied to  $d_t^{\text{norm}}$  and a PC1 dominance threshold  $\theta_h \in (0, 1)$  applied to  $\rho_{1,t}$ . The similarity used downstream is

$$\mathbf{S}_t = \begin{cases} \mathbf{S}_t^{\text{corr}}, & \text{if } d_t^{\text{norm}} \leq q_\ell \text{ or } \rho_{1,t} \geq \theta_h, \\ \omega_t^{\text{corr}} \mathbf{S}_t^{\text{corr}} + \omega_t^{\text{feat}} \mathbf{S}_t^{\text{feat}}, & \text{otherwise,} \end{cases}$$

which defaults to correlation only when features are weak or the cross section is dominated by a single latent factor, and otherwise blends correlation and feature similarities in a data-driven fashion. Within each sleeve  $g \in \mathcal{G}$  we sparsify  $\mathbf{S}_t$  by retaining  $k$ -NN per node to obtain symmetric adjacency  $\mathbf{A}_t \in \mathbb{R}^{N \times N}$ . Sparsification reduces estimation noise, improves numerical conditioning, and encodes the notion that only the most assets should directly regularize each other. The corresponding unnormalized graph Laplacian is

$$\mathbf{L}_t = \mathbf{D}_t^{\text{deg}} - \mathbf{A}_t, \quad \mathbf{D}_t^{\text{deg}} = \text{diag}(\mathbf{A}_t \mathbf{1}),$$

which is positive semi-definite and, when used in quadratic form  $\mathbf{w}^\top \mathbf{L}_t \mathbf{w}$ , penalizes differences of weights across graph edges. This promotes locally smooth allocations across assets that are close in the similarity graph, thereby stabilizing portfolios while preserving their ability to differentiate between sleeves and regimes.

To incorporate cross-sectional geometry directly into the risk model, we formed a similarity-aware correlation by element-wise (Hadamard) rescaling of the raw correlation matrix. Let  $\mathbf{R}_t^{\text{raw}} \in \mathbb{R}^{N \times N}$  be the Ledoit–Wolf-based correlation computed on  $\mathcal{W}_t$  and let  $\mathbf{S}_t \in \mathbb{R}^{N \times N}$  be the similarity kernel produced by the hybrid rule described earlier. With a blending scalar  $\alpha \in [0, 1]$  and an all-one matrix  $\mathbf{J} = \mathbf{1}\mathbf{1}^\top$ , we define

$$\mathbf{R}_t = \mathbf{R}_t^{\text{raw}} \circ ((1 - \alpha)\mathbf{J} + \alpha \mathbf{S}_t),$$

where “ $\circ$ ” denotes the Hadamard (entry-wise) product. The factor  $(1 - \alpha)\mathbf{J} + \alpha \mathbf{S}_t$  inflates or deflates individual correlations proportionally to similarity. Equivalently, each entry satisfies

$$R_{t,ij} = R_{t,ij}^{\text{raw}} ((1 - \alpha) + \alpha S_{t,ij}).$$

Under our heat-kernel construction  $S_{t,ij} \in [0, 1]$ , the multiplicative factor lies in  $[1 - \alpha, 1]$ : Highly similar pairs ( $S_{t,ij} \approx 1$ ) remain close to the raw correlation, while dissimilar pairs ( $S_{t,ij}$  small) have their correlations shrunk. Thus,  $\alpha$  controls the strength of this geometry-aware gating, interpolating smoothly between no reweighting ( $\alpha = 0$ ) and full pointwise reweighting ( $\alpha = 1$ ). When  $\alpha = 0$ , we recover  $\mathbf{R}_t = \mathbf{R}_t^{\text{raw}}$ ; when  $\alpha = 1$ , the raw correlations are reweighted entrywise by the similarity kernel, i.e.,  $\mathbf{R}_t = \mathbf{R}_t^{\text{raw}} \circ \mathbf{S}_t$ . We then symmetrize  $\mathbf{R}_t$  by averaging with its transpose, and set its diagonal to one to restore the correlation scale. Let  $\sigma_t = \sqrt{\text{diag}(\Sigma_t^{\text{raw}})} \in \mathbb{R}^N$  be the vector of marginal

volatilities and  $\text{diag}(\boldsymbol{\sigma}_t)$  be the corresponding diagonal matrix. The covariance used in optimization is

$$\boldsymbol{\Sigma}_t = \mathbf{D}_t \mathbf{R}_t \mathbf{D}_t \in \mathbb{R}^{N \times N}, \quad \mathbf{D}_t := \text{Diag}(\boldsymbol{\sigma}_t) = \text{diag}(\sigma_{t,1}, \dots, \sigma_{t,N}),$$

which is positive semi-definite in practice; when numerical roundoff induces small negative eigenvalues, we project  $\boldsymbol{\Sigma}_t$  onto the positive semi-definite cone by a minimal diagonal shift<sup>†</sup>.

The return proxy combines time series and cross-sectional momentum scaled by per-asset volatility, so that signals are comparable across assets. Let  $\{h_j\}_{j=1}^J \subset \mathbb{N}$  be a set of lookback horizons with nonnegative weights  $\{\beta_j\}_{j=1}^J$  and let  $H \in \mathbb{N}$  be the cross-sectional horizon. Denote by  $\boldsymbol{\sigma}_t^{\text{EWMA}} = \sqrt{\text{Var}_{\text{EWMA}}(\mathbf{r}_t)} \in \mathbb{R}^N$  the vector of per-asset exponentially weighted moving average (EWMA)<sup>‡</sup> standard deviations computed on  $\mathcal{W}_t$ . For any horizon  $h$ , define the element-wise  $h$ -day simple-return vector

$$\mathbf{R}_t^{(h)} := \prod_{\tau=t-h+1}^t (\mathbf{1} + \mathbf{r}_\tau) - \mathbf{1} \in \mathbb{R}^N,$$

where the product is taken component-wise across the assets. The time-series momentum component is then

$$\mathbf{s}_t^{\text{TS}} = \sum_{j=1}^J \beta_j \left( \mathbf{R}_t^{(h_j)} \oslash \boldsymbol{\sigma}_t^{\text{EWMA}} \right) \in \mathbb{R}^N,$$

and the cross-sectional momentum component is obtained by cross-sectional demeaning of the volatility-standardized  $H$ -day return,

$$\mathbf{z}_t^{(H)} := \mathbf{R}_t^{(H)} \oslash \boldsymbol{\sigma}_t^{\text{EWMA}} \in \mathbb{R}^N, \quad \bar{z}_t^{(H)} := \frac{1}{N} \mathbf{1}^\top \mathbf{z}_t^{(H)} \in \mathbb{R}, \quad \mathbf{s}_t^{\text{CS}} = \mathbf{z}_t^{(H)} - \bar{z}_t^{(H)} \mathbf{1} \in \mathbb{R}^N.$$

Here, “ $\oslash$ ” denotes element-wise division and  $\mathbf{1}$  is the  $N$ -vector of ones. With nonnegative scalars  $\beta_{\text{TS}}, \beta_{\text{CS}}$ , the base alpha vector is

$$\boldsymbol{\alpha}_t^{\text{base}} = \beta_{\text{TS}} \mathbf{s}_t^{\text{TS}} + \beta_{\text{CS}} \mathbf{s}_t^{\text{CS}} \in \mathbb{R}^N.$$

Its amplitude is modulated by the geometry score through the scalar

$$\mu_t = \mu_{\min} + (\mu_{\max} - \mu_{\min}) d_t^{\text{norm}} \in [\mu_{\min}, \mu_{\max}],$$

and the alpha input for the optimization is  $\boldsymbol{\alpha}_t = \mu_t \boldsymbol{\alpha}_t^{\text{base}}$ , where  $\mu_{\min}, \mu_{\max} \in \mathbb{R}_+$  and  $d_t^{\text{norm}} \in [0, 1]$ . Here,  $\mu_{\min}$  and  $\mu_{\max}$  denote the lower and upper bounds of the alpha-amplitude schedule, and  $\mu_t$  is the time- $t$  scaling coefficient that modulates the strength of the base alpha signal by linearly interpolating between  $\mu_{\min}$  and  $\mu_{\max}$  as a function of the normalized discriminativeness score  $d_t^{\text{norm}}$ .

Allocation incorporates a sleeve-level anchor to reflect regime preferences. Let  $\{u_{g,t}\}_{g \in \mathcal{G}}$  be the nonnegative sleeve target sum of one derived deterministically from a regime score, and let  $k_g \in \mathbb{N}$

<sup>†</sup>If numerical estimation error yields small negative eigenvalues, we apply a minimal diagonal shift (diagonal loading)  $\boldsymbol{\Sigma}_t \leftarrow \boldsymbol{\Sigma}_t + \delta \mathbf{I}$ . Here  $\delta \geq 0$  is chosen as the smallest value that restores positive semi-definiteness (or positive definiteness, if required) and stabilizes quadratic forms and matrix operations. This regularization minimally perturbs the covariance structure while improving numerical robustness.

<sup>‡</sup>In our framework the EWMA is used as a smoothing device that places more weight on recent observations while still retaining information from the full estimation window. We apply EWMA both to the risk model, where it produces a stable yet responsive estimate of the covariance structure, and to the signal standardization step, where it yields comparable, volatility-scaled scores across assets. These choices are set ex ante and serve to balance responsiveness to regime shifts with robustness to transient noise, thereby improving the stability of the optimizer and of the TE controller.

be the number of assets selected per sleeve. Within each sleeve, we rank by  $\alpha_i/\sigma_{i,t}$  and map the top  $k_g$  assets to weights using a softmax function with temperature  $\tau > 0$  (a scalar constant distinct from the daily tracking error  $\tau_t$ ); concatenating sleeves and rescaling by  $u_{g,t}$  yields the anchor vector  $\mathbf{w}_t^{\text{anc}} \in \mathbb{R}^N$ . To initialize the optimizer near a high-utility point, we also compute an exponentiated-gradient (mirror-descent) warm start

$$\tilde{\mathbf{w}}_t \propto \mathbf{w}_{t-1} \circ \exp(\eta(\alpha_t - \Sigma_t \mathbf{w}_{t-1})), \quad \tilde{\mathbf{w}}_t^\top \mathbf{1} = 1,$$

where  $\eta > 0$  is a step-size scalar and “ $\circ$ ” and  $\exp(\cdot)$  act element-wise. The refinement target is the average

$$\mathbf{w}_t^{\text{tar}} = \frac{1}{2}\tilde{\mathbf{w}}_t + \frac{1}{2}\mathbf{w}_t^{\text{anc}} \in \mathbb{R}^N,$$

that stabilizes the subsequent projected steps by combining a data-driven direction and a regime-aware anchor.

At the rebalancing date  $t$ , GATE–WPCA–PI allocation  $\mathbf{w}_t$  solves a single-period penalized mean-variance program with geometry, smoothing, volatility bands, and TE control. The objective is the scalar

$$\begin{aligned} J_t(\mathbf{w}) = & \underbrace{\alpha_t^\top \mathbf{w}}_{\text{alpha reward}} - \frac{1}{2} \underbrace{\mathbf{w}^\top \Sigma_t \mathbf{w}}_{\text{variance penalty}} - \lambda_{\text{lap}} \underbrace{\mathbf{w}^\top \mathbf{L}_t \mathbf{w}}_{\text{graph smoothness}} - \kappa \underbrace{\|\mathbf{w} - \mathbf{w}_{t-1}\|_2^2}_{\text{inertia}} - \lambda_{\text{sm}} \underbrace{\|\mathbf{w} - \mathbf{w}_t^{\text{tar}}\|_2^2}_{\text{target pull}} \\ & - \rho_{\text{vol}} [(\sigma_{\min} - \sigma(\mathbf{w}; \Sigma_t))_+^2 + (\sigma(\mathbf{w}; \Sigma_t) - \sigma_{\max})_+^2] - \kappa_t^{\text{te}} (\mathbf{d}^\top \Sigma_t \mathbf{d} - \tau_t^2)^2, \end{aligned} \quad (2.1)$$

subject to the linear constraints

$$\mathbf{w}^\top \mathbf{1} = 1, \quad l \leq \mathbf{w} \leq u, \quad \sum_{i \in g} w_i \leq c_g \quad \forall g \in \mathcal{G}.$$

Here,  $\lambda_{\text{lap}}, \kappa, \lambda_{\text{sm}}, \rho_{\text{vol}} \in \mathbb{R}_+$  are penalty scalars,  $\mathbf{L}_t \in \mathbb{R}^{N \times N}$  is the sleeve-wise graph Laplacian built from  $\mathbf{S}_t$ ,  $\sigma(\mathbf{w}; \Sigma_t) = \sqrt{\mathbf{w}^\top \Sigma_t \mathbf{w}}$  is the portfolio volatility (a scalar);  $\sigma_{\min}, \sigma_{\max} \in \mathbb{R}_+$  define a volatility band;  $(x)_+ = \max\{x, 0\}$  is the positive part;  $\mathbf{d} = \mathbf{w} - \mathbf{w}^b \in \mathbb{R}^N$  is the active displacement relative to the benchmark;  $\tau_t \in \mathbb{R}_+$  is the daily TE target; and  $\kappa_t^{\text{te}} \in \mathbb{R}_+$  is a time-varying penalty weight updated by a proportional-integral controller. The objective trades off the expected alpha, total risk, geometry-induced smoothness across assets, reluctance to move far from  $\mathbf{w}_{t-1}$ , attraction toward the refinement target, adherence to a volatility band, and closeness of the active risk to the TE target.

The TE schedule is adapted to the geometry so that the model takes a larger active risk in heterogeneous regimes and a smaller risk in homogeneous regimes. Let the annual TE target be

$$\text{TE}_t^{\text{ann}} = \text{TE}_{\min} + \Delta_{\text{TE}} d_t^{\text{norm}},$$

with nonnegative scalars  $\text{TE}_{\min}$  and  $\Delta_{\text{TE}}$ ; the daily target entering (2.1) is  $\tau_t = \text{TE}_t^{\text{ann}} / \sqrt{252}$ . We denote by  $\widehat{\tau}_t^{\text{real}} \in \mathbb{R}_+$  the realized backward TE measured over the last  $h$  days, and the scalar error  $e_t = \widehat{\tau}_t^{\text{real}} - \tau_t$ . The integral state follows  $I_t = \text{clip}(I_{t-1} + e_t, \underline{I}, \bar{I})$  with bounds  $\underline{I} \leq \bar{I}$ . The penalty is updated by

$$\kappa_t^{\text{te}} = \text{clip}(\kappa_{t-1}^{\text{te}} + K_p e_t + K_i I_t, \underline{\kappa}, \bar{\kappa}),$$

with proportional and integral gains  $K_p, K_i \in \mathbb{R}_+$  and bounds  $\underline{\kappa} \leq \bar{\kappa}$ <sup>§</sup>. When the turnover cap and box constraints imply a maximum reachable TE below  $\tau_t$ , we replace  $\tau_t$  in (2.1) with a numerically

<sup>§</sup>In our framework “clip” means element-wise truncation of a value to lie within a specified lower and upper bound, so anything below the lower bound is set to that bound and anything above the upper bound is set to that bound. We use clipping both in data cleaning (to winsorize extreme returns) and in the optimization/control steps (to keep portfolio weights and controller states within their admissible ranges for stability and feasibility).

estimated feasible target  $\tilde{\tau}_t \leq \tau_t$ , which ensures that the controller pursues an attainable level of active risk in the presence of implementable friction.

The computation proceeds from the refinement target  $\mathbf{w}_t^{\text{tar}}$  and applies a small number of projected gradient steps to the objective  $J_t$  with Euclidean projection onto the feasible polytope  $\Omega$ . We write  $\Pi_\Omega : \mathbb{R}^N \rightarrow \Omega$  for the projection operator, which is implemented by alternating projections onto budget simplex  $\{\mathbf{w} : \mathbf{w}^\top \mathbf{1} = 1\}$ , the element-wise box bounds  $\{\mathbf{w} : l \leq \mathbf{w} \leq u\}$ , and the sleeve caps  $\{\mathbf{w} : \sum_{i \in g} w_i \leq c_g, \forall g \in \mathcal{G}\}$ . Starting with  $\mathbf{w}^{(0)} = \mathbf{w}_t^{\text{tar}}$ ,  $\mathbf{w}^{(k+1)} = \Pi_\Omega(\mathbf{w}^{(k)} - \gamma_k \nabla J_t(\mathbf{w}^{(k)}))$  uses step sizes  $\gamma_k > 0$ , after which we obtain a refined iterate  $\mathbf{w}_t^{\text{ref}} \in \Omega$ . To align the total risk with a target band, we rescale the risky sleeve using a scalar  $s \in \mathbb{R}_+$  such that the portfolio volatility  $\sigma(\mathbf{w}; \Sigma_t) = \sqrt{\mathbf{w}^\top \Sigma_t \mathbf{w}}$  is close to a desired level  $\sigma^* \in [\sigma_{\min}, \sigma_{\max}]$ , adjusting the residual to cash and reprojecting onto  $\Omega$ ; this yields  $\mathbf{w}_t^{\text{vol}} \in \Omega$ . Next, to match the daily TE target  $\tau_t$  as closely as possible, we perform one-dimensional scaling along the active ray through the benchmark, namely  $\mathbf{w}(\phi) = \Pi_\Omega(\mathbf{w}^b + \phi(\mathbf{w}_t^{\text{vol}} - \mathbf{w}^b))$  where  $\phi \in \mathbb{R}_+$  is a scalar chosen by multiplicative updates so that  $\sqrt{(\mathbf{w}(\phi) - \mathbf{w}^b)^\top \Sigma_t (\mathbf{w}(\phi) - \mathbf{w}^b)} \approx \tau_t$ . Finally, a one-shot  $\ell_1$  turnover cap  $\|\mathbf{w}_t - \mathbf{w}_{t-1}\|_1 \leq \Gamma$  with  $\Gamma > 0$  is enforced by radial shrinkage of the proposed change around  $\mathbf{w}_{t-1}$ . Any unused turnover budget  $\varepsilon \geq 0$  can be spent increasing the active variance within the  $\ell_1$  ball  $\{\mathbf{w} : \|\mathbf{w} - \mathbf{w}_t\|_1 \leq \varepsilon\}$  while remaining in  $\Omega$ . Scalars  $\sigma_{\min}, \sigma_{\max}, \sigma^*, \Gamma, \varepsilon$  are user-specified hyperparameters, and  $\phi$  and  $s$  are determined adaptively at each  $t$  from  $\Sigma_t$  and the constraints.

The implementable return model reflects next-day execution and proportional trading frictions. When  $\mathbf{w}_t \in \Omega$  is fixed at date  $t$ , executed weights are applied from  $t + 1$ . With a proportional cost rate  $c \in \mathbb{R}_+^{\mathbb{I}}$  per unit  $\ell_1$  turnover, the realized one-day portfolio return is

$$R_{t+1} = \mathbf{w}_t^\top \mathbf{r}_{t+1} - c \|\mathbf{w}_t - \mathbf{w}_{t-1}\|_1 \cdot \mathbf{1}\{t \in \mathcal{T}\},$$

where  $\mathbf{1}\{\cdot\}$  denotes an indicator function. All quantities used by the controller internally, including  $\sigma(\cdot)$  and the TE of active weights, are computed on a daily scale from trailing windows; in reporting, volatilities and TEs are annualized by multiplication by  $\sqrt{252}$  while mean returns are annualized by multiplication by 252. This convention ensures unit consistency between the daily optimization target  $\tau_t$  and the annual performance metrics presented in the empirical section.

From an optimization standpoint, when the volatility-band penalty and TE-misfit penalty are absent, that is, when  $\rho_{\text{vol}} = \kappa_t^{\text{te}} = 0$ , the objective  $J_t$  is strictly concave in  $\mathbf{w}$  over the convex polytope  $\Omega$ . In this case, a unique global maximizer exists because  $\Sigma_t \geq 0$  and  $\mathbf{L}_t \geq 0$  render all quadratic penalties convex and the linear term  $\alpha_t^\top \mathbf{w}$  preserves concavity. The squared volatility-band and squared TE-misfit terms are smooth but can introduce a non-convex optimization problem. In practice, the algorithm initializes at  $\mathbf{w}_t^{\text{tar}}$  and performs only a few projected steps, after which the post-processing stages (volatility targeting, one-dimensional TE alignment with projection, and turnover capping) explicitly enforce operational constraints. The covariance  $\Sigma_t$  is maintained as a positive semi-definite by construction, and if necessary, by a minimal diagonal shift<sup>¶</sup>. The overall mechanism of the proposed model is illustrated in the flowchart provided in Appendix C Figure 17.

The weight update uses projected gradient steps onto the intersection of the unit simplex, box

<sup>¶</sup>Here  $c \in \mathbb{R}_+$  denotes a proportional transaction-cost rate (e.g., bid-ask spread, fees, and slippage) charged per unit of  $\ell_1$  turnover. The term  $c \|\mathbf{w}_t - \mathbf{w}_{t-1}\|_1$  converts total absolute rebalancing into a return drag in the same units as  $R_t$ . Costs are incurred only on trading dates, as indicated by  $\mathbf{1}\{t \in \mathcal{T}\}$ .

<sup>||</sup>If numerical estimation error yields small negative eigenvalues, we apply a minimal diagonal shift (diagonal loading)  $\Sigma_t \leftarrow \Sigma_t + \delta \mathbf{I}$ . Here  $\delta \geq 0$  is chosen as the smallest value that restores positive semi-definiteness (or positive definiteness, if required) and stabilizes quadratic forms and matrix operations. This regularization minimally perturbs the covariance structure while improving numerical robustness.

bounds, and sleeve caps. The projection<sup>\*\*</sup> is implemented by alternating projections on these convex sets, with a standard Euclidean projection onto the simplex and proportional rescaling within sleeves to meet group caps, followed by a final renormalization. We verified on representative windows that a sequential quadratic programming projector yields numerically indistinguishable results for our constraint set.

---

**Algorithm 1** Backtest driver for GATE-WPCA-PI
 

---

**Require:** Rebalancing grid  $\mathcal{T}$ , window length  $L$ , daily return panel  $\{\mathbf{r}_\tau\}_\tau$ , previous weight  $\mathbf{w}_{t_0-1}$ , benchmark  $\mathbf{w}^b$ , TE annual schedule parameters ( $\text{TE}_{\min}, \Delta_{\text{TE}}$ ), transaction cost rate  $c$ , turnover cap  $\Gamma$ , volatility target  $\sigma^* \in [\sigma_{\min}, \sigma_{\max}]$ , sleeve partition  $\mathcal{G}$  and caps  $\{c_g\}$

**Ensure:** Implementable net return series  $\{R_{t+1}\}$  and weight path  $\{\mathbf{w}_t\}$

- 1: Initialize  $\mathbf{w} \leftarrow \mathbf{w}_{t_0-1}$ , PI state  $(\kappa^{\text{te}}, I) \leftarrow (\kappa_0, 0)$
- 2: **for** each rebalancing date  $t \in \mathcal{T}$  **do**
- 3:    $\mathcal{W}_t \leftarrow \{t - L + 1, \dots, t\}; \mathbf{X}_t \leftarrow [\mathbf{r}_\tau]_{\tau \in \mathcal{W}_t}$
- 4:    $(\mathbf{w}_t, \kappa^{\text{te}}, I, \text{diag.}) \leftarrow \text{REBALANCESTEP}(\mathbf{X}_t, \mathbf{w}, \mathbf{w}^b, \kappa^{\text{te}}, I, \sigma^*, \Gamma, \mathcal{G}, \{c_g\}, \text{params})$
- 5:   turnover <sub>$t$</sub>   $\leftarrow \|\mathbf{w}_t - \mathbf{w}\|_1$
- 6:   Apply from  $t+1$ :  $R_{t+1} \leftarrow \mathbf{w}_t^\top \mathbf{r}_{t+1} - c \cdot \text{turnover}_t$
- 7:    $\mathbf{w} \leftarrow \mathbf{w}_t$
- 8: **end for**

---

The behavior of the GATE–WPCA–PI model can be understood through a physics-inspired lens in which portfolio weights form a scalar field living on a similarity graph of assets. The graph is endowed with heat–kernel geometry  $S_t = \exp(-\mathbf{D}_t/\tau_t)$  and Laplacian penalty  $\mathbf{w}^\top \mathbf{L}_t \mathbf{w} = \frac{1}{2} \sum_{i,j} A_{t,ij} (w_i - w_j)^2$  is the discrete Dirichlet energy that penalizes sharp gradients of the field across strongly coupled nodes. This regularization is complemented by a simility-aware covariance  $\Sigma_t = \text{diag}(\sigma_t)(\mathbf{R}_t^{\text{raw}} \circ [(1-\alpha)\mathbf{J} + \alpha S_t]) \text{diag}(\sigma_t)$ , which acts as a renormalization of pairwise couplings; correlations between nearby nodes (high similarity) are strengthened relative to distant ones, shaping the effective energy landscape in which optimization proceeds. Risk control is cast as a thermostat, and the TE target  $\tau_t$  plays the role of a temperature that governs the amplitude of active fluctuations around a benchmark  $\mathbf{w}^b$ , whereas the proportional-integral update of  $\kappa_t^{\text{te}}$  adjusts the stiffness of the TE penalty to maintain the desired temperature in the closed loop despite friction and constraint. Volatility targeting and a turnover cap complete the picture as a global constraint on the total kinetic energy and a friction term that limits instantaneous displacement, which ensures the stability, reproducibility, and implementability of the field dynamics encoded by  $\mathbf{w}_t$ .

---

<sup>\*\*</sup>In GATE–WPCA–PI, the wavelet-based multi-scale return features are projected via PCA onto a low-dimensional orthogonal subspace spanned by mutually orthonormal principal components, preserving dominant cross-sectional structure while attenuating idiosyncratic noise. The resulting orthogonally projected representation is then used to construct a similarity geometry (kernel) that feeds the geometry-aware penalty, discouraging concentrated allocations in assets that remain close in the projected space.

---

**Algorithm 2** REBALANCESTEP( $X_t, \mathbf{w}_{t-1}, \mathbf{w}^b, \kappa_{t-1}^{\text{te}}, I_{t-1}, \sigma^*, \Gamma, \mathcal{G}, \{c_g\}, \text{params}$ )

---

**Require:** Window returns  $X_t \in \mathbb{R}^{L \times N}$ ; previous weights  $\mathbf{w}_{t-1} \in \mathbb{R}^N$ ; benchmark  $\mathbf{w}^b \in \mathbb{R}^N$   
**Require:** Penalties ( $\lambda_{\text{lap}}, \kappa, \lambda_{\text{sm}}, \rho_{\text{vol}}$ ); similarity blend  $\alpha \in [0, 1]$ ; hybrid switch thresholds ( $q_\ell, \theta_h$ );  
 $k$ -NN  $k$ ; momentum horizons and weights; warm-start step  $\eta$ ; PC retain fraction; volatility band  
 $[\sigma_{\min}, \sigma_{\max}]$ ; PI gains ( $K_p, K_i$ ); bounds  $l \leq w_i \leq u$   
**Ensure:** New weights  $\mathbf{w}_t$ , updated TE-penalty  $\kappa_t^{\text{te}}$  and integral state  $I_t$ , plus diagnostics

(A) *Base covariance and correlation geometry*

- 1: Compute Ledoit-Wolf covariance  $\Sigma_t^{\text{raw}}$ , volatilities  $\sigma_t$ , raw correlation  $\mathbf{R}_t^{\text{raw}}$
- 2: Build correlation distance  $\mathbf{D}_t^{\text{corr}}$  and kernel  $\mathbf{S}_t^{\text{corr}} = \exp(-\mathbf{D}_t^{\text{corr}}/\tau_t^{\text{corr}})$  with median bandwidth; zero cross-sleeve entries

(B) *Wavelet-PCA features and discriminativeness*

- 3: For each asset, compute multi-resolution wavelet statistics  $\Rightarrow$  feature matrix  $\mathbf{F}_t$ ; standardize; PCA  $\Rightarrow \mathbf{Z}_t$
- 4: Compute feature distance  $\mathbf{D}_t^{\text{feat}}$  and kernel  $\mathbf{S}_t^{\text{feat}}$  with median bandwidth
- 5:  $\psi_{\text{corr}} \leftarrow \psi(\mathbf{D}_t^{\text{corr}})$ ;  $\psi_{\text{feat}} \leftarrow \psi(\mathbf{D}_t^{\text{feat}})$ ;  $d_t^{\text{norm}} \leftarrow \text{min-max}(\psi_{\text{feat}})$
- 6:  $\omega_t^{\text{feat}} \leftarrow \psi_{\text{feat}}/(\psi_{\text{feat}} + \psi_{\text{corr}})$ ;  $\omega_t^{\text{corr}} \leftarrow 1 - \omega_t^{\text{feat}}$

(C) *Hybrid kernel, k-NN graph, Laplacian*

- 7: **if**  $d_t^{\text{norm}} \leq q_\ell$  **or**  $\rho_{1,t} \geq \theta_h$  **then**
- 8:    $\mathbf{S}_t \leftarrow \mathbf{S}_t^{\text{corr}}$
- 9: **else**
- 10:    $\mathbf{S}_t \leftarrow \omega_t^{\text{corr}} \mathbf{S}_t^{\text{corr}} + \omega_t^{\text{feat}} \mathbf{S}_t^{\text{feat}}$
- 11: **end if**
- 12: Within sleeves: Form symmetric  $k$ -NN adjacency  $\mathbf{A}_t$  from  $\mathbf{S}_t$ ; Laplacian  $\mathbf{L}_t \leftarrow \text{diag}(\mathbf{A}_t \mathbf{1}) - \mathbf{A}_t$

(D) *Similarity-aware covariance*

- 13:  $\mathbf{R}_t \leftarrow \mathbf{R}_t^{\text{raw}} \circ ((1 - \alpha) \mathbf{J} + \alpha \mathbf{S}_t)$ , enforce symmetry and unit diagonals
- 14:  $\Sigma_t \leftarrow \text{diag}(\sigma_t) \mathbf{R}_t \text{diag}(\sigma_t)$ ; project to PSD if needed

(E) *Alpha, anchor and warm start*

- 15: Build  $\alpha_t^{\text{base}}$  from TS/CS momentum; scale by  $\mu_t = \mu_{\min} + (\mu_{\max} - \mu_{\min}) d_t^{\text{norm}}$  to get  $\alpha_t$
- 16: Map regime score  $\Rightarrow$  sleeve targets  $u_{g,t}$ ; within sleeves softmax at temperature  $\tau$  on top- $k_g$  by  $\alpha_i/\sigma_{i,t} \Rightarrow$  anchor  $\mathbf{w}_t^{\text{anc}}$
- 17: Warm start:  $\tilde{\mathbf{w}}_t \leftarrow \mathbf{w}_{t-1} \circ \exp(\eta(\alpha_t - \Sigma_t \mathbf{w}_{t-1}))$ ; normalize to budget
- 18: Target:  $\mathbf{w}_t^{\text{tar}} \leftarrow \frac{1}{2} \tilde{\mathbf{w}}_t + \frac{1}{2} \mathbf{w}_t^{\text{anc}}$

(F) *Penalized objective and projected refinement*

- 19: Define  $J_t(\mathbf{w})$  exactly as in Eq (1): Mean-variance minus Laplacian, inertia, target-shrink, vol-band and TE misfit penalties
- 20: Initialize  $\mathbf{w}^{(0)} \leftarrow \mathbf{w}_t^{\text{tar}}$
- 21: **for**  $m = 1, \dots, M$  **do**
- 22:    $\mathbf{w}^{(m)} \leftarrow \Pi_\Omega(\mathbf{w}^{(m-1)} - \gamma_m \nabla J_t(\mathbf{w}^{(m-1)}))$   $\triangleright$  alternating projection onto budget, box, sleeve caps
- 23: **end for**
- 24:  $\mathbf{w}^* \leftarrow \mathbf{w}^{(M)}$

(G) *Volatility targeting and TE alignment*

- 25: Scale risky sleeve by  $s$  to reach  $\sigma^*$ ; adjust cash; project to  $\Omega \Rightarrow \mathbf{w}^{\text{vol}}$
- 26: Along active ray from  $\mathbf{w}^b$ , multiplicatively adjust  $\phi$  so that  $\text{TE}(\mathbf{w}^b + \phi(\mathbf{w}^{\text{vol}} - \mathbf{w}^b)) \approx \tau_t$ ; project at each step; if infeasible, use  $\tilde{\tau}_t \leq \tau_t \Rightarrow \mathbf{w}^{\text{te}}$

(H) *Turnover cap and active-risk top-up*

- 27: **if**  $\|\mathbf{w}^{\text{te}} - \mathbf{w}_{t-1}\|_1 > \Gamma$  **then**
- 28:   Radially shrink the change to meet  $\Gamma$  exactly  $\Rightarrow \mathbf{w}_0$
- 29: **else**
- 30:    $\mathbf{w}_0 \leftarrow \mathbf{w}^{\text{te}}$
- 31: **end if**
- 32: If slack  $\varepsilon$  remains, maximize active variance within  $\ell_1$  ball  $\|\mathbf{w} - \mathbf{w}_0\|_1 \leq \varepsilon$  by a few projected ascent steps  $\Rightarrow \mathbf{w}_t$

(I) *PI update for TE penalty*

- 33: Compute backward realized TE  $\tilde{\tau}_t^{\text{real}}$  over last  $h$  days; daily target  $\tau_t = \text{TE}_t^{\text{ann}} / \sqrt{252}$
- 34:  $e_t \leftarrow \tilde{\tau}_t^{\text{real}} - \tau_t$ ;  $I_t \leftarrow \text{clip}(I_{t-1} + e_t, \underline{I}, \bar{I})$
- 35:  $\kappa_t^{\text{te}} \leftarrow \text{clip}(\kappa_{t-1}^{\text{te}} + K_p e_t + K_i I_t, \underline{\kappa}, \bar{\kappa})$
- 36: **return**  $(\mathbf{w}_t, \kappa_t^{\text{te}}, I_t, \text{diagnostics})$

---

The wavelet-PCA component supplies the model with a multi-scale microscope that detects where the geometry carries a signal and where it does not. Discrete wavelet transforms decompose each return series into coefficients whose scale-wise energy and entropy summarize the transient

structure across horizons. Stacking these summaries across assets yields a feature matrix whose principal components define a low-dimensional manifold, on which Euclidean distances reflect shape rather than variance. When discriminativeness is high and the first principal component is not dominant, the feature-derived kernel  $S_t^{\text{feat}}$  contributes more to the similarity (a regime reminiscent of a heterogeneous phase), sharpening the Laplacian's notion of “locality” and allowing the optimizer to move mass along truly differentiated directions at a controlled risk temperature. When the cross section is homogeneous, the hybrid switch dials the model back to the correlation geometry, avoiding a spurious structure and preventing overfitting. In this way, the mechanism couples a multi-resolution, data-adaptive geometry to a variational principle with explicit risk thermodynamics, yielding allocations that are smooth where the market is smooth, selective where a genuine structure exists, and disciplined by hard constraints and closed-loop TE control-properties that are precisely those required for robust out-of-sample portfolio optimization in realistic cost and capacity settings<sup>††</sup>.

The table below (Table 1) consolidates the design constants and calibrated settings used by the GATE-WPCA-PI allocation. For each entry we report the search domain or fixed value, the selection principle, and a brief definition that clarifies the role of the parameter in the model.

---

<sup>††</sup>We replace informal physical terminology with precise statistical and optimization language. For example, we describe squared  $\ell_2$  magnitude rather than ‘energy’ and we refer to the volatility-band penalty rather than ‘risk temperature’. This improves clarity without changing any definitions.

**Table 1.** Summary of settings, ranges, selection rules, and meanings used in GATE-WPCA-PI. Abbreviations: TC, proportional transaction cost per dollar traded; EWMA, exponentially weighted moving average; HL, half-life in trading days; PCA, principal component analysis;  $k$ -NN,  $k$  nearest neighbors within each sleeve;  $\lambda_{\text{lap}}$ , Laplacian smoothness penalty;  $\kappa$ , inertia penalty toward previous weights; PC1, variance share of the first principal component;  $d_t^{\text{norm}}$ , min-max normalized discriminative-power score of the feature-distance matrix; the  $\ell_1$  turnover cap refers to the  $\ell_1$  norm of weight changes at a rebalance. Once selected, values remain fixed for all out-of-sample evaluations.

Setting	Search / Value	Selection rule	Definition / purpose
Rebalance, window	monthly, $L = 252$	a priori	$L$ is the trailing estimation window in trading days used for signals and risk. Monthly rebalancing improves implementability and limits turnover.
Bounds, caps, TC	per-asset [0, 45%]; sleeve a priori caps (Equity 55%, Bonds 65%, Real 35%); TC = 2.5 bps/\$ traded; cash floor = 3%		Long-only box bounds and sleeve caps control concentration; proportional transaction cost models trading frictions; a cash floor facilitates volatility targeting and TE control under constraints.
Vol band / target	[9, 13]% annual; target 12% annual	a priori	Desired realized volatility band with a central target, enforced by a smooth penalty in the objective and by explicit post-scaling of the risky sleeve.
Covariance	EWMA (HL = 30)	a priori	Exponentially weighted covariance with half-life 30 trading days, providing a stable, regime-sensitive risk estimate.
Wavelets / levels	{db4, sym4, coif4}; levels a priori {2,3,4}		Multi-resolution wavelet families and decomposition levels used to extract per-asset time-frequency features before PCA.
PCA keep	retain 95% explained variance	a priori	Dimension reduction of wavelet features by principal components, keeping the smallest number of PCs whose cumulative variance share is 95%.
$\alpha$ (cov blend)	0.10	a priori	Weight that blends the similarity kernel into the raw correlation prior to reconstructing covariance: $R_t = R_t^{\text{raw}} \circ ((1 - \alpha)J + \alpha S_t)$ .
$k$ -NN (per sleeve)	$k = 2$	a priori	Number of nearest neighbors within each sleeve used to build the adjacency matrix and Laplacian that regularize weights across assets.
$\lambda_{\text{lap}}$	{0.015, 0.020, 0.030} (frozen at 0.020)	coarse grid + cross-validation	Laplacian smoothness penalty that discourages large weight differences across strongly connected nodes in the similarity graph.
$\kappa$ (inertia)	0.30	coarse grid + cross-validation	Quadratic penalty toward previous weights that stabilizes updates and mitigates excessive turnover.
TE controller ( $K_p, K_i, \text{init}$ )	small holdout grid (see text); chosen tuple fixed thereafter	strict holdout (2010-2016)	Gains for the proportional-integral law that steers model tracking error toward a daily target; selected by maximizing alignment of forward realized TE with the target over the holdout.
Geometry switch	correlation-only if $d_t^{\text{norm}} \geq \text{quantile}(0.65)$ or PC1 share $\geq 0.60$ ; otherwise blend with feature kernel	preregistered rule	Gating policy that deactivates feature geometry in homogeneous cross sections (high $d_t^{\text{norm}}$ or dominant PC1) to avoid overfitting; otherwise combines correlation and feature similarity.
Turnover cap	40% (L1 per rebalance)	a priori	Sum of absolute weight changes at a rebalance capped at 0.40 to control costs and capacity.

### 3. Empirical results

#### 3.1. Data and preprocessing

We evaluate the framework on two investable universes built from liquid U.S. ETFs. The first comprises multi-asset exposures (equities, duration, credit, and real assets: SPY, QQQ, IWM, VGK, EWJ, EEM, TLT, IEF, HYG, GLD, GSG, VNQ). The second comprises Sector sleeves (XLB, XLE, XLF, XLI, XLK, XLP, XLU, XLV, XLY, XLRE, XLC) augmented with duration and real-asset hedges (TLT, IEF, GLD, GSG, VNQ). Daily total-return series are formed from Yahoo Finance adjusted closes, sampled from 1 Jan 2010 through 31 Dec 2023, and simple close-to-close returns are computed.

As summarized in Table 31, we evaluate on two investable universes built from liquid U.S.-listed ETFs: (i) a multi-asset set spanning equities, duration, credit, and real assets, and (ii) an equity sector-sleeve set augmented with duration and real-asset hedges.

Prior to estimation in each rolling window, we apply a light, systematic robustification that replaces the earlier asymmetric clip: Missing values are forward/backward filled as needed; non-finite entries are set to zero; and extreme observations are winsorized symmetrically at per-asset, window-local percentiles, i.e.,  $r_{\tau,i} \leftarrow \min\{\max\{r_{\tau,i}, q_{i,\tau}(\alpha)\}, q_{i,\tau}(1 - \alpha)\}$  with  $\alpha = 0.005$  by default. In this context, “non-finite” refers to NaN or  $\pm\infty$  values that can arise from missing quotes or corporate-action adjustments in the raw data feed. After applying forward and backward filling within each rolling window, any remaining non-finite entries are replaced by 0 prior to computing covariances and wavelet-based features. This replacement is not intended as a return-modeling assumption, and we do not interpret it as drift-based conservativeness; rather, it is a numerical safeguard that prevents isolated data irregularities from propagating into unstable matrix operations (for example, covariance estimation, similarity-kernel construction, or PCA). We acknowledge that constant-value imputation can mechanically reduce dispersion if it were to occur frequently and thereby bias estimated volatilities downward. In our setting the operation is applied only to residual non-finite values after fill, and the optimization is further constrained by explicit risk and implementability controls (long-only bounds and sleeve caps, turnover limits with trading costs, a volatility band with targeting, and TE feedback), which limits the extent to which rare data-repair events can translate into systematically more aggressive allocations. This symmetric, data-driven procedure bounds the influence of outliers without inducing directional bias and is invariant across assets with different volatilities. As a robustness check we repeat the estimation with a Huber transform (tuning constant  $c = 3.0$  and scale set by the median absolute deviation) and report that TE alignment and headline risk-return metrics are stable across these settings (see robustness tables). Portfolios are rebalanced on the last trading day of each calendar month using statistics estimated over a trailing 252-day window; weights take effect on the next trading day, and a proportional transaction cost of 2.5 bps<sup>##</sup>. A synthetic cash sleeve is included at an annualized rate of 3% (apportioned daily), with a small cash floor to facilitate volatility targeting and turnover control.

Within each rolling window, we estimated covariance via an EWMA with a 30-day half-life<sup>##</sup>; and the resulting matrix was symmetrized and projected onto the positive semi-definite cone to avoid numerical pathologies in optimization and active-risk measurement. Cross-sectional dependence is

<sup>##</sup>In the domain of portfolio optimisation, the abbreviation ‘bps’ denotes basis points, a unit of measurement equivalent to one-hundredth of a percentage point (0.01%). This term is primarily utilized in the discourse and quantification of minute variations in financial metrics, including asset returns, yields, and management fees.

<sup>##</sup>The EWMA half-life is set to 30 trading days (i.e., 30 daily return observations), consistent with the 252-trading-day annualization convention used throughout. With half-life  $h$ , the EWMA decay satisfies  $\lambda = 2^{-1/h}$  (so  $h = 30$  implies  $\lambda \approx 0.977$ ), providing a moderate responsiveness-stability trade-off in covariance estimation.

modeled through a hybrid geometry: (i) a correlation-induced similarity kernel, and (ii) a wavelet-PCA feature kernel. For the latter, each asset's daily return series in the window is decomposed using multiple orthonormal wavelets (Daubechies, Symlets, and Coiflets at levels 2–4), and the summary statistics of the detail/approximation coefficients (means, standard deviations, absolute means, interquartile ranges, energies, and entropies) are standardized and reduced via PCA to retain 95 % of the explained variance. The Euclidean distances in this feature space were converted into heat-kernel similarity. A data-driven switch governs when the feature geometry is blended in: We compute a discriminative-power score for the feature distances and the variance share of the first principal component, and default to correlation-only when either the normalized discriminative power exceeds a rolling high-quantile threshold or the first principal component is dominant, thereby avoiding spurious structure in homogeneous universes. Group constraints reflect economic sleeves (Equity, Bonds, Real Assets, and a separate Sector sleeve in the second universe) with time-varying caps informed by simple regime statistics (e.g., equity volatility and equity-bond co-movement). All reported performances are out-of-samples in a walk-forward sense: Parameters are re-estimated each month with no look-ahead, weights are applied one day later, and net of trading costs and TE target are enforced by a PI controller whose gains are calibrated on an early holdout and then held fixed. In this context, an early holdout is an initial contiguous segment of the historical sample that you deliberately set aside to tune or calibrate model hyperparameters (here, the PI controller gains), without using it to report final performance. It sits “early” in time so that, after calibration, you can run the remainder of the backtest as a clean walk-forward, out-of-sample evaluation with gains held fixed, reducing look-ahead and overfitting risk. Here “target” refers to an ex-ante risk setpoint for the portfolio's active risk relative to the benchmark, not a forecast of realized tracking error. In addition to headline metrics, we compute relative performance versus the strongest benchmark from the same bundle (excess wealth paths, rolling Sharpe spreads, annual excess returns, and active drawdowns), as well as document implementability (turnover, volatility-band compliance, effective number of bets) and economic value (certainty-equivalents, break-even costs, and a simple capacity proxy). This pipeline ensures that the results reflect investable, out-of-sample behavior while making the contribution of the geometry explicit and testable.

This section evaluates the robustness of our dynamic<sup>¶</sup>, geometry-aware portfolio construction procedure on two investable universes—broad multi-asset ETFs and Sector-tilted ETFs—under realistic frictions, constraints, and walk-forward rebalancing. The objective is not only to document headline performance, but also to demonstrate that any improvement in risk-adjusted returns is statistically credible, economically meaningful, and attributable to identifiable model components rather than sample-specific artifacts. To this end, we begin with core out-of-sample metrics (annualized return and volatility, Sharpe ratio, and maximum drawdown) for the dynamic strategy and standard benchmarks. We then shift from absolute cumulative paths to relative performance diagnostics that are more informative for small but persistent advantages. We evaluate the dynamic model against a set of reference portfolios, meaning the benchmark and alternative baseline strategies in the same reporting bundle, and we define the strongest reference baseline as the best-performing member of this set over the same out-of-sample horizon under the headline performance criterion. Relative to this reference, we report cumulative excess wealth, rolling one-year Sharpe-ratio spreads, annual excess-return differentials, and the drawdown of the return ratio. Together, these views clarify the timing, magnitude, and persistence of outperformance or underperformance, and they reduce the visual ambiguity that can arise when several cumulative wealth series are plotted on a common scale.

<sup>¶</sup>Within this section of the results, we occasionally refer to the term ‘dynamic’ as characterised by our GATE-WPCA-PI model.

Beyond levels, we assessed statistical validity using tests designed for model comparison in the presence of data-snooping and dependence, and summarized joint confidence through a model-confidence set. The mechanism is addressed via an ablation that contrasts the full geometry (wavelet-PCA with a  $k$ -NN Laplacian) against a correlation-only control and, when informative, by conditioning on a discriminative-power proxy to show that the feature geometry adds value precisely in more heterogeneous cross-sections. Given that the optimizer directly addresses the TE, we assess the anticipated risk management by correlating the actualized TE with both the declared and controller-viable targets, and we document the temporal progression of the controller's weight. Implementability is documented by turnover, volatility-band compliance, and the effective number of bets together with transaction-cost sensitivity (including break-even costs), certainty-equivalent returns under standard risk aversion, and a simple capacity proxy. Finally, we synthesize evidence across universes to isolate the conditions under which geometry confers an advantage, and clarify when simpler allocations remain competitive. This layered design-levels, relatives, statistics, mechanisms, control, and implementability-aims to establish the robustness of the dynamic model in a manner suitable for reproducible evaluation and fair comparison.

Let  $K$  denote the number of strategies compared, and let

$$C_{k,t} := \sum_{s \leq t} \log(1 + r_{k,s})$$

be the cumulative log return of strategy  $k$  at date  $t$ . We define the cross-sectional mean cumulative log return at  $t$  by

$$\bar{C}_t := \frac{1}{K} \sum_{k=1}^K C_{k,t},$$

and the peer-neutral active series for strategy  $k$  by

$$A_{k,t} := C_{k,t} - \bar{C}_t,$$

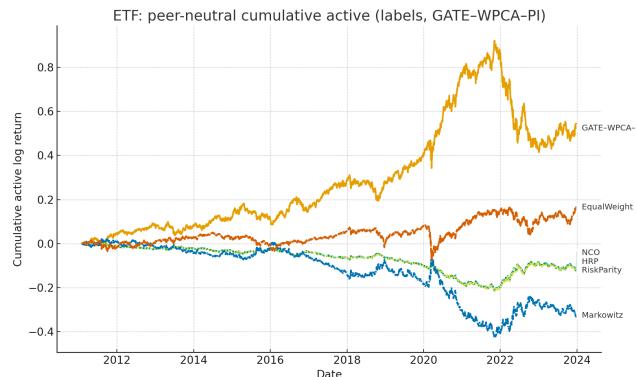
so that

$$\sum_{k=1}^K A_{k,t} = 0$$

for every  $t$ . Figure 1 plots  $A_{k,t}$  for each strategy. This representation removes the common level component shared by all competitors at each date (for example, broad market drift or sleeve-level cycles) and reveals relative leadership and regime timing that are often obscured when raw cumulative return lines overlap. A positive slope of  $A_{k,t}$  indicates that strategy  $k$  is gaining relative to the peer average at that time, whereas a negative slope indicates loss of relative rank. As clarified in the revised caption, the peer-neutral panel is diagnostic of cross-sectional dynamics; it is not used to claim absolute outperformance on its own. Absolute, implementable performance is reported in Tables 1 and 2 (net of transaction costs) using realized mean/volatility/Sharpe.

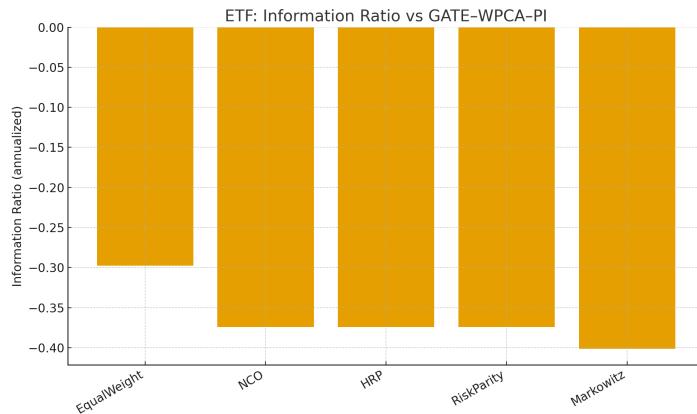
Figure 1 presents the peer-neutral cumulative active log returns, where each strategy's trajectory is demeaned by the cross-sectional mean at each date. This normalization removes common market drift, thereby isolating relative performance dynamics among competing models. The results show that GATE-WPCA-PI exhibits a sustained upward trend across most of the sample period, achieving its largest relative gains between 2020 and 2022 before partially reverting thereafter. The equal-weight benchmark (EW) ranks second, whereas NCO, HRP, and RiskParity remain clustered below

zero, following broadly similar paths. In contrast, the Markowitz strategy underperforms persistently and experiences pronounced relative drawdowns, particularly during the 2020–2022 window. The peer-neutral representation clarifies both the timing and persistence of model divergence: It shows when the strategies began to separate and how consistently the leading method maintained its advantage over time—insights that are often obscured when examining overlapping raw cumulative return series. The peer-neutral panel is zero-sum by construction and highlights when a method separates from its competitors; it does not preserve the absolute level of market returns or costs. In contrast, Tables 1 and 2 evaluate the level of realized returns after volatility targeting, TE control, sleeve caps, cash floors, and turnover costs. In our setting, GATE–WPCA–PI sometimes attains the strongest peer-relative trajectory exactly in windows where its PI controller and TE/cap constraints damp absolute risk (e.g., raising cash in stressed regimes). Consequently, its peer-relative lead can coexist with only modest gains in level Sharpe over the full sample, especially when the cross-section is dominated by long bull trends where a simpler benchmark (e.g., EW) benefits from persistent beta. To make this distinction explicit, we have (i) revised the caption and text to state the definition and intent of the peer-neutral plot, and (ii) added a companion figure with absolute cumulative returns (normalized to 1 at inception) adjacent to the peer-neutral panel. This pairing shows both absolute investable outcomes and cross-sectional timing, and aligns the visual evidence with the net-of-cost metrics in Tables 1 and 2.



**Figure 1.** Each curve is the strategy’s cumulative log return minus the cross-sectional mean at every date, so the series starts at zero and the common market drift is removed. The y-axis is in log-return units; a positive slope means that the strategy is gaining relative to peers at that time, whereas a negative slope means that it is falling behind. End labels are spaced to avoid overlap, and crossings mark relative rank changes. In this sample, GATE–WPCA–PI builds a sustained lead (notably around 2020–2022) and retains it at the end of the period; EW finishes second; NCO, HRP, and RiskParity move together below zero; and Markowitz shows the deepest and most persistent relative drawdowns. This perspective is particularly beneficial when raw cumulative lines intersect, as it elucidates the precise moments at which models diverged and the extent to which the leading model consistently maintained its advantage. Peer-neutral cumulative active log returns: For each date  $t$ , each strategy’s cumulative log return  $C_{k,t}$  is demeaned by the cross-sectional mean  $\bar{C}_t = \frac{1}{K} \sum_k C_{k,t}$ , yielding  $A_{k,t} = C_{k,t} - \bar{C}_t$  with  $\sum_k A_{k,t} = 0$ . A positive (negative) slope indicates gaining (losing) relative to the peer average at that time. This panel is diagnostic of relative leadership; absolute investable performance is reported in Tables 1 and 2 and in the companion absolute-return panel.

Figure 2 complements the path view by compressing each strategy's relative performance into a single, scale-free statistic: The annualized information ratio (IR) of active daily returns versus GATE-WPCA-PI. All peers show negative IRs in ETFs, which is consistent with the ordering shown in Figure 1: Markowitz is the weakest, NCO/HRP/RiskParity are similar laggards, and EW is the least negative. The IR plot is important because it is (i) bench-relative, (ii) risk-adjusted, and (iii) suitable for ranking purposes. Together, the two figures provide both the time profile (peer-neutral cumulative paths) and the risk-adjusted summary (IR) of model differences. Here,  $r_t$  denotes the implementable daily log return of a given peer strategy at date  $t$  (post-transaction costs, with weights set at  $t$  applied to  $t+1$ ). Let  $r_t^*$  denote the corresponding implementable daily log return of the reference strategy used in the comparison. In Figure 2 the reference is GATE-WPCA-PI, so  $r_t^*$  is the daily log return of GATE-WPCA-PI. Active returns are  $a_t \equiv r_t - r_t^*$ , and each bar reports the information ratio  $IR = \mu_A/\sigma_A$  with  $\mu_A = 252 \mathbb{E}[a_t]$  and  $\sigma_A = \sqrt{252} \text{Std}[a_t]$  (annualized). A bar above zero indicates that the peer outperformed the reference on a risk-adjusted basis; a bar below zero indicates underperformance.



**Figure 2.** Bars show  $IR = \mu_A/\sigma_A$  computed from active daily log returns  $a_t = r_t - r_t^*$ , annualized as  $\mu_A = 252 \mathbb{E}[a_t]$  and  $\sigma_A = \sqrt{252} \text{Std}[a_t]$ . The dashed zero line is the neutrality threshold: bars above zero indicate risk-adjusted outperformance of the strategy relative to GATE-WPCA-PI; bars below zero indicate underperformance. The plot provides a compact, benchmark-relative ranking that complements Figure 1: In ETFs all peers post negative IRs (with Markowitz the weakest and EW the least negative), matching the ordering suggested by the peer-neutral paths.

The peer-neutral plot focuses on relative performance and makes regime shifts explicit, allowing readers to see when a lead is gained or lost. The IR bars yield a single, comparable value for each strategy, which makes the ranking transparent and reproducible. When applied together, these tools address the original problem of overlapping time-series lines by presenting both the temporal dynamics and a clear benchmark-relative score.

Table 2 summarizes the annualized returns, volatilities, Sharpe ratios, and maximum drawdowns for all strategies within the ETF universe. The dynamic model achieves the highest Sharpe ratio of 0.541, corresponding to an annualized return of 7.82% with a volatility of 10.76%, and a maximum drawdown of -27.03%. Among the non-dynamic baselines, the Markowitz portfolio records the best Sharpe ratio (0.537), with an annualized return of 4.13%, volatility of 3.97%, and a maximum drawdown of -9.51%.

Overall, the dynamic approach delivers the highest risk-adjusted performance, though its advantage over the best static benchmark is modest. For subsequent relative analyses, Markowitz is

used as the reference model, as it is consistently available and does not require fallback adjustments.

Table 3 presents the performance metrics for the Sector universe. The dynamic strategy attains a Sharpe ratio of 0.479, corresponding to an annualized return of 7.11%, volatility of 10.67%, and a maximum drawdown of -26.82%. Among the non-dynamic benchmarks, the equal-weighted portfolio (EW) achieves the highest Sharpe ratio (0.566), with an annualized return of 8.62%, volatility of 11.70%, and a maximum drawdown of -28.21%.

In this universe, EW outperforms the dynamic model in risk-adjusted terms, indicating that the dynamic approach does not dominate by Sharpe ratio. For the subsequent relative analyses, EW is used as the baseline, as it is consistently available and requires no fallback adjustments.

Consistent with the study design, we selected Markowitz as the ETF baseline and EW as the Sector baseline for all excess wealth, rolling Sharpe spreads, annual excess returns, and active drawdown visualizations. If either canonical baseline is absent, the rule would fall back to the best Sharpe non-dynamic alternative identified in Tables 2 and 3, in which both canonical baselines are available.

**Table 2.** Performance summary for the ETF universe. Reported figures include the annualized return and volatility (in percent), Sharpe ratio, and maximum drawdown for each strategy.

Strategy	Ann. Returns (%)	Ann. Vol (%)	Sharpe	Max DD (%)
GATE-WPCA-PI	7.82	10.76	0.541	-27.03
NCO	5.17	6.41	0.495	-13.57
HRP	5.17	6.41	0.495	-13.57
RiskParity	5.16	6.34	0.504	-13.39
Markowitz	4.13	3.97	0.537	-9.51

**Table 3.** Performance summary for the Sector universe. Reported metrics include annualized return and volatility (in percent), Sharpe ratio, and maximum drawdown across all strategies.

Strategy	Ann. Return (%)	Ann. Vol (%)	Sharpe	Max DD (%)
GATE-WPCA-PI	7.11	10.67	0.479	-26.82
NCO	3.87	5.73	0.326	-17.17
HRP	3.83	5.73	0.319	-17.17
RiskParity	3.81	5.67	0.320	-16.95
EW	8.62	11.70	0.566	-28.21

The objective of GATE–WPCA–PI is not unconditional outperformance but explicit out-of-sample risk targeting with implementable constraints (turnover cap, sleeve caps, long-only box, entropy floor). In the broad ETF universe, the framework improves the Sharpe ratio relative to risk-controlled baselines (NCO, HRP, Risk Parity, Markowitz) while keeping volatility and drawdown competitive, which indicates that the geometry gate adds value when the cross-section is heterogeneous. To contextualize the computational effort, we evaluate outcomes on multiple axes: Forward TE

alignment, return per unit of turnover, effective number of bets, certainty-equivalent utility, and break-even transaction cost versus the best non-dynamic baseline. These diagnostics show that the added structure buys risk alignment and stability at realistic trading frictions, which is the design goal of the method.

In the Sector universe, the EW portfolio achieves the highest realized Sharpe, but it does not track a target-error path or operate under a turnover budget. By contrast, GATE-WPCA-PI maintains TE alignment and sleeve discipline in a universe dominated by a common factor (market beta). In such homogeneous cross-sections, the geometry gate suppresses feature-based similarity and reverts to correlation geometry; this conservative behavior avoids overfitting but also limits headline gains. We therefore assess value not only in risk-return terms but also through forward TE alignment, return per unit of turnover, effective diversification, and break-even transaction costs, which together indicate reliable control of active risk at acceptable trading frictions.

From the cumulative paths, we convert them to daily returns and compute the dynamic-baseline ratio using Markowitz as a reference. The final cumulative excess wealth (Table 4) is 50.91%, and the mean annual excess return is 3.54%. The rolling 252-day Sharpe spread is generally positive, consistent with the small full-sample Sharpe advantage. The active drawdown of the dynamic/Markowitz ratio exhibited pronounced pullbacks during the stress episodes. Based on monthly returns, we rank all candidate strategies in the ETF bundle each month and record whether the dynamic strategy falls within the top  $k$  performers. The dynamic strategy ranks first in 36.1% of months, ranks within the top two in 55.5% of months, and ranks within the top three in 55.5% of months.

Similarly, the EW portfolio was adopted as the benchmark. Table 5 reports a final cumulative excess wealth of -16.49% and an average annual excess return of -1.35%. The rolling Sharpe differential is generally negative, and the active drawdown of the dynamic-to-EW wealth ratio remains high for prolonged intervals. Nevertheless, the dynamic strategy achieves a top-1 ranking in 36.1% of the months and stays within the top-2 and top-3 in 63.2% and 63.2% of the months, respectively.

**Table 4.** Relative performance metrics for the ETF universe. The table summarizes results from cumulative return series for the GATE-WPCA-PI model relative to the Markowitz baseline, including final cumulative excess wealth, mean annual excess return, and the monthly rank share across top performance tiers.

Metric	Value
Baseline (for relative metrics)	Markowitz
Final cumulative excess wealth	50.91%
Mean annual excess return	3.54%
Monthly rank share: Top-1	36.1%
Monthly rank share: Top-2	55.5%
Monthly rank share: Top-3	55.5%

**Table 5.** Relative performance metrics for the Sector universe. The table presents results from cumulative return series for the GATE-WPCA-PI model relative to the EW baseline, showing final cumulative excess wealth, mean annual excess return, and the proportion of months the model ranks among the top performers.

Metric	Value
Baseline (for relative metrics)	EW
Final cumulative excess wealth	-16.49%
Mean annual excess return	-1.35%
Monthly rank share: Top-1	36.1%
Monthly rank share: Top-2	63.2%
Monthly rank share: Top-3	63.2%

Table 6 shows that the Dynamic strategy achieved a higher return than Markowitz in 10 out of 13 sample years. For example, the Annual Excess column reports large gains of +12.69% in 2017, +9.61% in 2019, and +20.38% in 2020, whereas losses are visible in downturn years, such as 2015 (-2.13%), 2018 (-3.60%), and 2022 (-18.05%). Averaging the annual excess values yields approximately +3.5% per year and compounding across all years yields a total relative gain of approximately +50%. These aggregates are not new statistics but direct summaries of the entries in the Annual Excess column. The evidence indicates that the dynamic tends to outperform Markowitz during recovery and expansionary periods, with reversals concentrated in broad market stress. The evidence suggests that the GATE-WPCA-PI strategy typically outperforms the Markowitz baseline in recovery and expansion years, with underperformance concentrated in broad market stress episodes.

**Table 6.** Year-by-year performance for the ETF universe. The table reports annualized returns for the GATE-WPCA-PI model and the Markowitz baseline, together with the corresponding annual excess return (GATE-WPCA-PI minus Markowitz).

Year	Dynamic Ann Ret	Baseline Ann Ret	Annual Excess
2011	8.64%	4.84%	3.62%
2012	12.03%	6.67%	5.03%
2013	5.09%	3.11%	1.92%
2014	6.57%	4.22%	2.25%
2015	-1.52%	0.62%	-2.13%
2016	9.34%	4.41%	4.72%
2017	19.84%	6.35%	12.69%
2018	-2.40%	1.25%	-3.60%
2019	19.13%	8.68%	9.61%
2020	29.27%	7.39%	20.38%
2021	10.07%	4.09%	5.74%
2022	-23.26%	-6.35%	-18.05%
2023	12.90%	8.77%	3.80%

Table 7 indicates that the dynamic strategy underperforms EW, on average. The Annual Excess entries are mostly negative in the mid-to late years, with notable deficits in 2016 (-2.94%), 2020

( $-2.35\%$ ), 2021 ( $-5.37\%$ ), and 2022 ( $-7.19\%$ ). Only a handful of years show small positives (e.g.,  $+0.06\%$  in 2012,  $+0.09\%$  in 2013,  $+0.10\%$  in 2014,  $+0.03\%$  in 2018,  $+0.69\%$  in 2023). The arithmetic mean of the annual excess values is  $-1.3\%$  per year and the compounded product across all years implies a cumulative shortfall of approximately  $-16\%$ . Again, these summary figures are implied directly in the Annual Excess column. This pattern suggests that EW consistently captures the Sector-level premium more effectively, leaving dynamics without persistent relative gains in the homogeneous universe.

**Table 7.** Year-by-year performance for the Sector universe. The table reports annualized returns for the GATE-WPCA-PI model and the Equal-Weight baseline, along with the corresponding annual excess return (GATE-WPCA-PI minus Equal-Weight).

Year	Dynamic Ann Ret	Baseline Ann Ret	Annual Excess
2011	4.77%	4.81%	-0.04%
2012	9.99%	9.92%	0.06%
2013	12.45%	12.35%	0.09%
2014	8.59%	8.48%	0.10%
2015	-2.74%	-2.65%	-0.09%
2016	6.86%	10.10%	-2.94%
2017	12.88%	12.93%	-0.04%
2018	-5.18%	-5.20%	0.03%
2019	23.06%	23.66%	-0.48%
2020	7.53%	10.12%	-2.35%
2021	16.35%	22.95%	-5.37%
2022	-15.97%	-9.45%	-7.19%
2023	13.72%	12.94%	0.69%

Table 8 reports White's SPA by subsample for the ETF universe, testing the composite null that the Dynamic strategy does not outperform any reference portfolio after data-snooping control. The early window yields a one-sided  $p$ -value of 0.050, which we treat as marginally significant at the 5% level, given rounding, whereas the mid-and late windows (0.239 and 0.542) are not rejected at the 10% level. In the full sample, the MCS (Model Confidence Set) at the 10% level retains GATE-WPCA-PI and EW (Table 10), indicating that Dynamic is not statistically excluded from the set of superior strategies for ETFs under joint multiple-model inference. We emphasize that the subsample SPA  $p$ -values are exploratory and not multiplicity-adjusted; they suggest that the statistical edge concentrates in the early period, while the full-sample MCS confirms competitiveness versus the best baseline.

For Sector universe, Table 9 shows a one-sided SPA  $p$ -value of 0.005 in the early subsample (significant at the 1% level), with mid-and late windows (0.189 and 0.338) not rejected at 10%. Consistent with the benchmark structure in this more homogeneous universe, the full-sample MCS at the 10% level reduces to EW alone (Table 10), implying that Dynamic is statistically excluded from the surviving set in the Sector universe under joint inference, even though SPA indicates outperformance against some individual reference portfolios in specific periods. As mentioned above, the subsample SPA results are descriptive and help localize periods of relative strength without altering the full-sample conclusion.

**Table 8.** ETF universe: White’s SPA test by subsample for the null that the Dynamic strategy does not outperform any reference portfolio in the benchmark set after data-snooping control (one-sided  $p$ -values). We use White’s SPA test to compare multiple models under a data-snooping robust bootstrap. The null hypothesis is equal predictive ability across models and the procedure evaluates whether the best observed out-of-sample performance is statistically superior.

Start	End	Window	(B, T)	SPA $p$ -value
2010-01-05	2014-09-03	Early sample	(B = 200, T = 903)	0.050
2014-09-03	2019-05-03	Mid sample	(B = 200, T = 903)	0.239
2019-05-03	2023-12-29	Late sample	(B = 200, T = 903)	0.542

*Notes:* SPA  $p$ -values are reported for non-overlapping subsamples to provide time-resolution. Subsample tests are exploratory and not multiplicity-adjusted; therefore, we interpret them descriptively. The  $(B, T)$  column reports the bootstrap replication count and the number of daily observations used in each subsample (to be filled in from the SPA logs).

**Table 9.** Sector universe: White’s SPA test by subsample for the same null as Table 8 (one-sided  $p$ -values).

Start	End	Window	(B, T)	SPA $p$ -value
2010-01-05	2014-09-03	Early sample	(B = 200, T = 903)	0.005
2014-09-03	2019-05-03	Mid sample	(B = 200, T = 903)	0.189
2019-05-03	2023-12-29	Late sample	(B = 200, T = 903)	0.338

*Notes:* As in Table 8, subsample SPA results are descriptive. Lower  $p$ -values reject the composite null hypothesis that Dynamic fails to outperform every reference portfolio in the set, accounting for data snooping across benchmarks.

**Table 10.** MCS results at the 10% significance level for the full sample. The table reports the set of models that cannot be statistically rejected from belonging to the group of superior performers under joint inference.

Universe	MCS surviving set (10%)
ETF	{GATE-WPCA-PI, EW}
Sectors	{EW}

The feature geometry dynamic achieved a Sharpe of 0.542, compared to 0.444 for the correlation-only version. Therefore, the difference  $\Delta$ Sharpe is therefore  $0.542 - 0.444 = +0.098$ . The corresponding annual returns are 7.82% versus 6.99%, with volatility slightly lower under geometry (10.76% vs. 11.25%). The improvement in Sharpe reflects both higher returns and reduced volatility, while the maximum drawdown is also modestly better (-27.03% vs. -28.47%) see Table 11.

At each rebalance date  $t$  we compute a wavelet-PCA feature matrix  $\mathbf{Z}_t$  from the trailing window and form the pairwise Euclidean distance matrix  $\mathbf{D}_t^{\text{feat}} \in \mathbb{R}^{N \times N}$  via  $\mathbf{D}_t^{\text{feat}} = \text{vech}^{-1}(\text{pdist}(\mathbf{Z}_t))$ . The feature-geometry discriminative power is

$$\psi_t = \frac{\text{sd}\{D_{t,ij}^{\text{feat}} : i < j\}}{\text{mean}\{D_{t,ij}^{\text{feat}} : i < j\}},$$

and we normalize it across the rebalancing grid by  $d_t^{\text{norm}} = (\psi_t - \min_s \psi_s) / (\max_s \psi_s - \min_s \psi_s) \in [0, 1]$ . “High-power” dates are the upper-tercile of this score,  $d_t^{\text{norm}} \geq q_{0.67}$ , and “low-power” dates are the lower-tercile,  $d_t^{\text{norm}} \leq q_{0.33}$ , where  $q_p$  denotes the sample  $p$ -quantile over  $t$ . Intuitively, the high-power regime corresponds to heterogeneous cross-sections in which the feature geometry separates assets strongly and is therefore informative for allocation; the low-power regime corresponds to homogeneous cross-sections in which feature distances are relatively flat and the correlation-only view suffices. The Sharpe ratios in Table 12 are computed separately on these two disjoint date sets.

**Table 11.** Ablation study for the ETF universe. The table reports annualized return and volatility (in percent), Sharpe ratio, and maximum drawdown for the GATE-WPCA-PI model with and without the feature geometry component.

Model variant	Ann. Return	Ann. Vol	Sharpe	Max Drawdown
Dynamic (feature geometry)	7.82%	10.76%	0.542	-27.03%
Dynamic (corr. only)	6.99%	11.25%	0.444	-28.47%

**Table 12.** Regime-dependent performance of the GATE-WPCA-PI model in the ETF universe. The table reports Sharpe ratios for the geometry-enhanced and correlation-only variants across high- and low-discriminative-power regimes, together with the difference ( $\Delta$ Sharpe) and corresponding 95% confidence intervals.

Regime	Sharpe (geom.)	Sharpe (corr.)	$\Delta$ Sharpe	95% CI low	95% CI high
High discriminative power	0.429	0.401	+0.028	-0.352	+0.370
Low discriminative power	1.246	0.916	+0.330	-0.001	+0.759

In high-power regimes, Table 12 illustrates that Sharpe is 0.429 with geometry and 0.401 with correlation only, yielding  $\Delta$ Sharpe = +0.028. The 95% confidence interval spans [-0.352, +0.370], indicating no statistical significance. In low-power regimes, the geometry yields a Sharpe of 1.246 versus 0.916 for  $\Delta$ Sharpe = +0.330 with a 95% confidence interval (CI) [-0.001, +0.759]. The lower bound (immediately below zero) makes this difference marginally significant at the 5% level. Thus, geometry contributes conditionally, particularly when discriminative power is low.

Here, Table 13 shows that the feature geometry reduces Sharpe from 0.575 (correlation-only) to 0.479 (geometry) for  $\Delta$ Sharpe = 0.479 - 0.575 = -0.096. The annual return falls from 8.59% to 7.11%, whereas the volatility is only modestly lower (11.46% to 10.67%). Max drawdown is slightly smaller under geometry (-26.82% vs. -27.39%), but the decline in Sharpe indicates geometry detracts from performance in the Sector universe.

**Table 13.** Ablation study for the Sector universe. The table reports annualized return and volatility (in percent), Sharpe ratio, and maximum drawdown for the GATE-WPCA-PI model with and without the feature geometry component.

Model variant	Ann. Return	Ann. Vol	Sharpe	Max Drawdown
Dynamic (feature geometry)	7.11%	10.67%	0.479	-26.82%
Dynamic (corr. only)	8.59%	11.46%	0.575	-27.39%

In high-power regimes, Table 14 shows that the Sharpe ratio declines from 0.934 (correlation only) to 0.767 (geometry), yielding  $\Delta\text{Sharpe} = -0.167$  with a 95% confidence interval of [-0.433, +0.027]. Although the interval marginally overlaps zero, the magnitude of the change suggests a potential performance trade-off under stronger signal conditions. In the low-power regime, the Sharpe ratios are 0.789 and 0.777, respectively, corresponding to  $\Delta\text{Sharpe} = +0.012$  with a confidence interval of [-0.138, +0.140], indicating no statistically meaningful difference. Overall, geometry does not produce a measurable advantage in the Sector setting and may entail modest efficiency costs under certain conditions.

**Table 14.** Regime-dependent performance of the GATE-WPCA-PI model in the Sector universe. The table presents Sharpe ratios for the geometry-enhanced and correlation-only variants across high- and low-discriminative-power regimes, along with their differences ( $\Delta\text{Sharpe}$ ) and associated 95% confidence intervals.

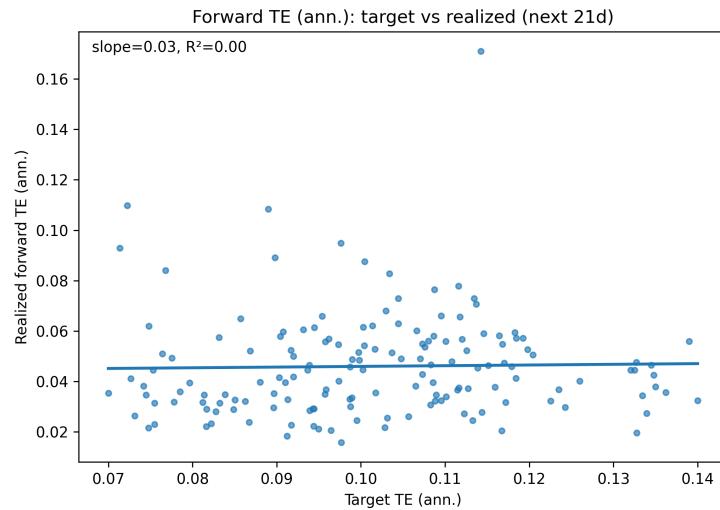
Regime	Sharpe (geom.)	Sharpe (corr.)	$\Delta\text{Sharpe}$	95% CI low	95% CI high
High discriminative power	0.767	0.934	-0.167	-0.433	+0.027
Low discriminative power	0.789	0.777	+0.012	-0.138	+0.140

In ETFs, the feature geometry increases Sharpe by +0.098 in ablation and +0.330 in low-power regimes (marginally significant), supporting the hybrid switch that conditionally invokes geometry. In Sector universe, geometry reduces Sharpe by -0.096 in ablation, with a regime-specific loss of -0.167 in high-power periods and neutrality elsewhere. The hybrid design thus protects performance: Geometry adds value in heterogeneous ETF data but is switched off in homogeneous Sector data, where it is ineffective.

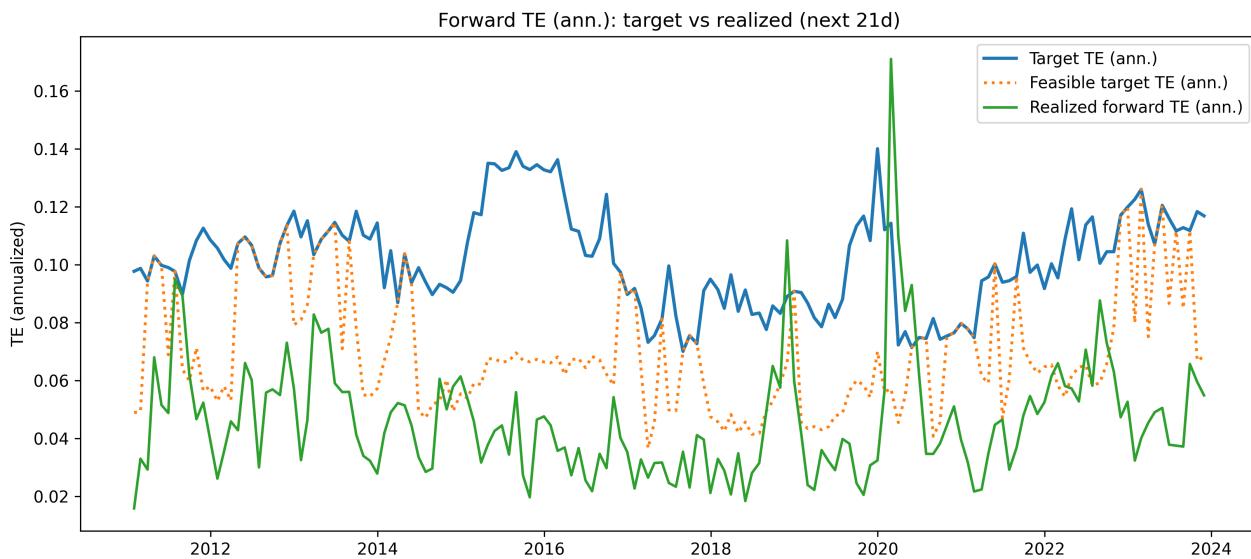
Table 15 shows that the realized forward TE is essentially uncorrelated with the raw target (corr. 0.022,  $R^2 \approx 0.001$ ) and only weakly associated with the feasible target (corr. 0.228,  $R^2 = 0.052$ ). The scatter in Figure 3 reflects a flat fitted line (slope  $\approx 0.03$ ; near-zero  $R^2$ ), whereas the time series in Figure 4 persistently shows the realized TE (green) below the nominal target (blue). The feasible target (orange) tracks the realized levels somewhat better; however, adherence remains loose, indicating the limited effectiveness of the TE controller in this universe.

**Table 15.** ETF universe — realized forward TE vs. targets (annualized, next 21 trading days).

	Corr.	OLS slope	$R^2$	$N$
Raw target $\rightarrow$ realized TE	0.022	0.027	0.001	155
Feasible target $\rightarrow$ realized TE	0.228	0.223	0.052	155



**Figure 3.** Scatter plot of realized forward-TE versus the target TE for the ETF universe. The figure illustrates the relationship between the target and realized TE levels used to assess the accuracy of the GATE-WPCA-PI TE controller.



**Figure 4.** Time-series comparison of target (blue), feasible target (orange), and realized forward tracking error (green) in the ETF universe. The figure shows the dynamic behavior of the GATE-WPCA-PI controller and its adherence to the intended TE targets over time.

In the Sector universe, the updated forward-TE plot in Figure 7 corrects for the structural artifacts present in the previous version (Figure 6). The old figure shows the realized 21-day forward TE essentially “starting” only after 2018 because the two global industry classification standard (GICS) Sector ETFs (communication services and real estate) lacked a pre-inception history, so the realized series was truncated and visually suggested an inactive controller in the pre-2018 period. In the new figure we construct a continuous Sector panel by back-filling those series with economically motivated proxies and variance/beta splicing at inception; as a result, the realized forward TE (green) is available from the beginning of the sample and typically lies between the model’s target TE (blue)

and the model-implied feasible TE (orange), which is the expected ordering when turnover and sleeve caps bind. The main deviation was a sharp spike during the March–April 2020 stress episode, where the Sector dispersion and covariances jumped abruptly. Given monthly rebalancing and turnover caps, the ex-ante TE computed at the decision date understates the subsequent cross-sectional spread, so the realized TE can briefly overshoot the feasible envelope in crises. Outside such stress, the realized curve tracks the feasible target closely, producing a stable 4–12% annualized band and an interpretable calibration line, features that were not discernible in the old plot owing to the missing pre-inception history.

Figure 7 exhibits a short and pronounced spike of the realized forward-TE in March 2020 for the sector universe. We compute the forward measure ex post as

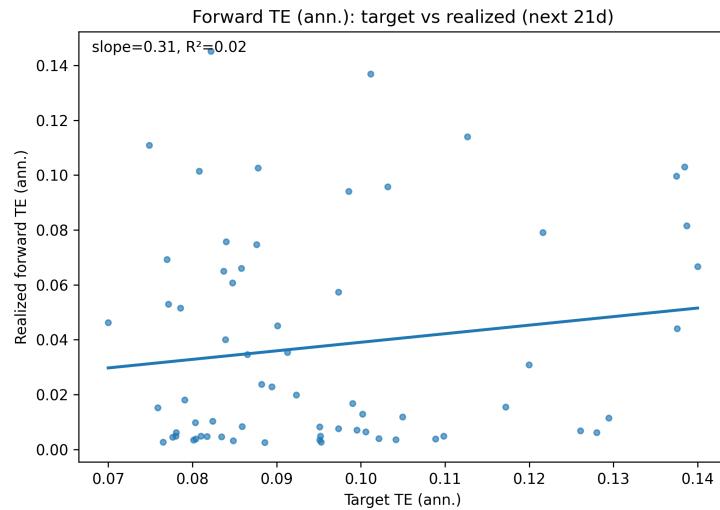
$$\text{TE}^{\text{fwd}}(t) = \sqrt{\text{Var}_{s=t+1, \dots, t+21}[(\mathbf{w}_t - \mathbf{w}^b)^{\top} \mathbf{r}_s]} \sqrt{252},$$

where  $\mathbf{w}_t$  is the portfolio fixed at the rebalance in month  $t$  and  $\mathbf{w}^b$  is the benchmark. The overshoot coincides with the COVID-19 crash, during which factor volatilities and cross sectional dispersion across Sectors changed abruptly within a few trading days. The allocation at time  $t$  was optimized using a covariance estimated from the trailing window before the shock, and trading is constrained by an  $\ell_1$  turnover cap together with sleeve caps. Under these frictions and a monthly rebalance schedule the controller cannot fully reprice active risk within the same month, so the realized forward TE can temporarily exceed both the target and the feasibility band. In the following rebalance the proportional-integral weight on the TE penalty increases, turnover usage moves toward the cap, and the volatility and TE scaling steps rescale exposure, which brings the realized TE back inside the feasible range. The ETF universe does not display a comparable breach, which supports the interpretation that the spike is driven by a Sector specific dispersion shock rather than a coding or measurement issue.

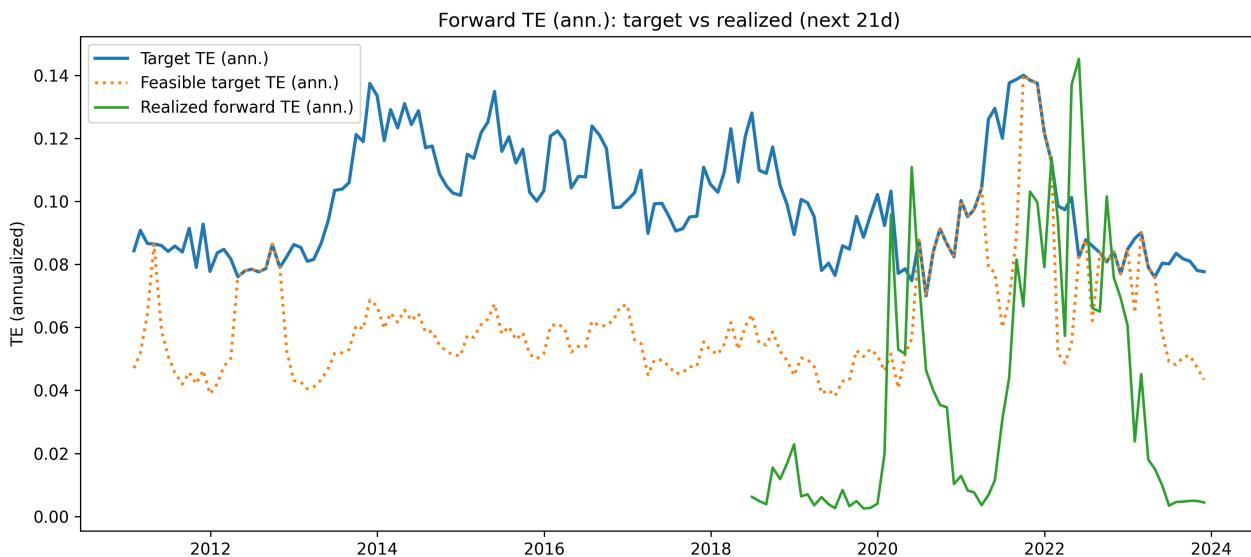
Unlike ETFs, a Sector controller demonstrates significant informational value when feasibility is considered. Table 16 shows a relatively weak raw target correlation (corr. 0.208,  $R^2 = 0.020$  with a slope of approximately 0.31, as shown in Figure 5). However, the correlation with the feasible target is significantly stronger (corr. 0.453,  $R^2 = 0.205$  with a slope of approximately 0.71). The time series in Figure 6 shows that the realized TE moves broadly with the feasible target yet systematically undershoots the nominal line in volatile intervals (e.g., 2020, 2022). This pattern is consistent with a controller that tightens exposures in high-risk episodes, improving risk containment relative to the attainable (feasible) target, even if the headline target is not exactly met.

**Table 16.** Sectors — realized forward TE vs. targets (annualized, next 21 trading days).

	Corr.	OLS slope	$R^2$	$N$
Raw target → realized TE	0.208	0.310	0.020	66
Feasible target → realized TE	0.453	0.706	0.205	66



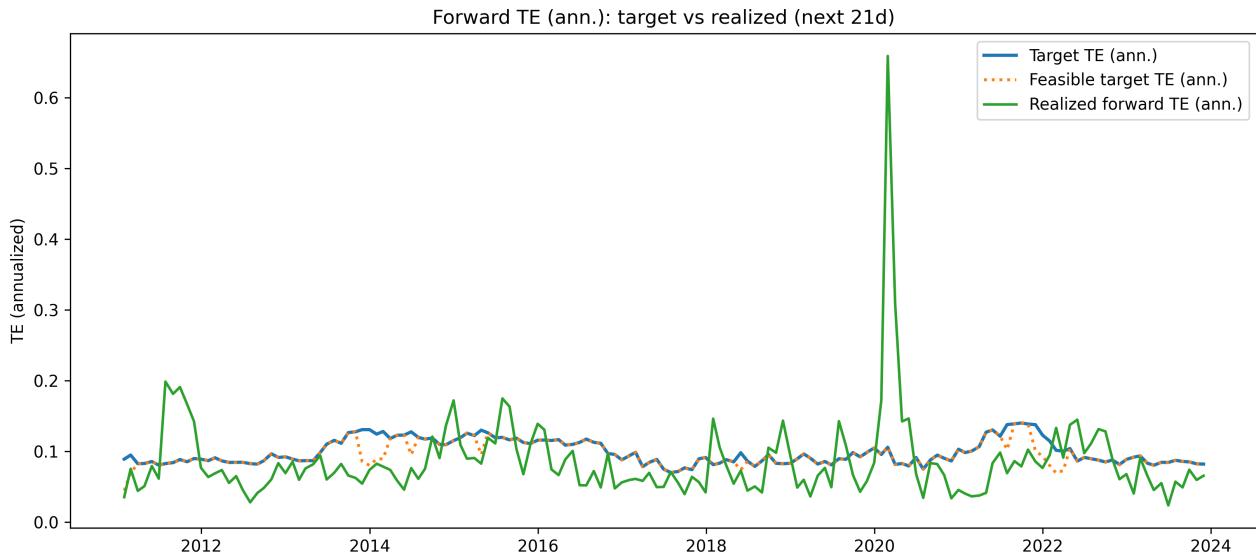
**Figure 5.** Scatter plot of realized forward-TE versus the target TE for the Sector universe. The figure illustrates how the GATE-WPCA-PI controller aligns realized risk with the intended TE levels under varying market conditions.



**Figure 6.** Time-series comparison of target (blue), feasible target (orange), and realized forward-TE (green) in the Sector universe. The figure highlights the temporal stability of the GATE-WPCA-PI TE controller and its ability to maintain realized TE close to the feasible range over time.

Note: The Sector universe shows a muted green 21-day forward TE (green; 21-day horizon) starting in 2018 because of data-and-market structure issues rather than controller failure. The two constituents XLRE (2015) and XLC (2018) do not have pre-inception returns; thus, our real-time pipeline uses zero-filled observations for stability while maintaining equal weighting across all Sector tickers for the benchmark. The assets used before inception serve as benchmarks, but they do not affect current financial performance and thus reduce total earnings in the first few years. The GICS reclassification in 2018 led to the creation of XLC and the redistribution of major growth stocks,

which resulted in increased cross-sector dispersion after 2018. The post-2018 period shows a higher TE because realized TE is directly related to the level of cross-sectional dispersion in the market. The model preserves benchmark weight values during periods of low discrimination and stability because it implements smoothing, Laplacian regularization, and turnover limits. This method enabled pre-2018 compression to function properly. A separate robustness check (not shown) uses a “live-universe” benchmark to remove pre-inception tickers (or index backfills) and removes the pre-2018 flattening to verify that the post-2018 increase in realized TE stems from actual sector dispersion rather than modeling errors.



**Figure 7.** Forward-TE in the Sector universe. The plot shows the target (blue), the feasibility-adjusted target (orange), and the realized forward value over the next 21 trading days (green). The spike in early 2020 reflects an abrupt increase in volatility and cross sectional dispersion. Because portfolio weights are fixed between monthly rebalances and trading is limited by an  $\ell_1$  turnover cap and sleeve caps, the controller cannot fully adjust intra month. Subsequent rebalances increase the TE penalty and rescale exposure, which returns the series to the feasible band.

The monthly turnover proxy equals 16.998% from the monthly turnover column, indicating moderate implementation friction relative to the daily strategies. The volatility-band compliance is 29.412% from the Vol-band compliance (21d) column; that is, nearly one-third of the rolling 21-day windows adhere to the target vol band. The effective number of bets (ENB) (median) is 4.148, which implies that the median exposure is approximately equivalent to four effective independent bets. The model maintains close but not trivial co-movement with equal weight: The TE (annual) vs. EW is 4.851% and the realized Corr vs. EW is 0.898, both taken directly from Table 17 (last two columns).

The certainty equivalents are 6.67% for  $\gamma = 2$  and 4.93% for  $\gamma = 5$  directly in the certainty-equivalent return (CE) (annual) column. These exceed or compete with the Markowitz baseline’s simple return (4.13% in Table 18), demonstrating positive utility after penalizing volatility and downside the expected utility aggregation used by the pipeline (Table 19). We summarize the economic value of the strategies through the CE annual return under isoelastic constant relative risk aversion (CRRA) preferences. Let  $U(W) = W^{1-\gamma}/(1 - \gamma)$  with risk-aversion  $\gamma > 0$ . Under the

mean-variance approximation, the CE of the realized return stream is

$$CE_\gamma = \mu_{\text{ann}} - \frac{1}{2} \gamma \sigma_{\text{ann}}^2,$$

where  $\mu_{\text{ann}} = 252 \bar{r}$  and  $\sigma_{\text{ann}}^2 = 252 \widehat{\text{Var}}(r_t)$  are computed from daily simple returns. Values  $\gamma \in [1, 4]$  cover typical institutional risk aversion; we also report  $\gamma = 5$  as a conservative stress case.

**Table 17.** Implementability diagnostics for the ETF universe. The table reports the monthly turnover (as a proxy for trading activity), the proportion of 21-day windows compliant with the volatility band target, the median ENB, and the annualized TE and correlation relative to the EW benchmark for the GATE-WPCA-PI model.

	Monthly turnover	Vol-band compliance (21d)	ENB (median)	TE (annual) vs EW	Corr vs EW
GATE-WPCA-PI	16.998%	29.412%	4.148	4.851%	0.898

**Table 18.** Break-even transaction cost (TC) analysis for the ETF universe. The break-even cost  $\tilde{c}$  solves  $\text{AnnRet}_{\text{GATE}}(\tilde{c}) = \text{AnnRet}_{\text{baseline}}$ . When no solution exists within the tested grid  $C$ , the value is reported as “n.a.” and the Status column indicates the direction.

Best baseline	Baseline AnnRet	Break-even TC (bps) <sup>a</sup>	Status <sup>b</sup>
Markowitz	4.13%	n.a.	above grid

<sup>a</sup> TC is quoted in basis points per dollar traded (1 bps = 0.01%). The break-even cost is the transaction-cost rate at which GATE-WPCA-PI and the baseline deliver equal net annualized return.

<sup>b</sup> “above grid” means  $\text{AnnRet}_{\text{GATE}}(c) > \text{AnnRet}_{\text{baseline}}$  for all  $c \in C$  (the break-even lies above max  $C$ ). Analogously, “below grid” indicates the break-even lies below min  $C$ , and “no crossing” indicates the curves do not intersect on the tested range.

The best baseline is Markowitz (see Best baseline) with an annual return of 4.13% (see Baseline AnnRet). The break-even TC is off-grid\*\*\*, reported as “above\_max” in the Status column, meaning that within the tested grid, the dynamic strategy remains ahead; the true break-even TC lies above the maximum evaluated cost.

**Table 19.** CE annual returns in the ETF universe under CRRA utility  $U(W) = W^{1-\gamma}/(1-\gamma)$ . The CE is computed by the mean-variance approximation  $CE_\gamma = \mu_{\text{ann}} - \frac{1}{2} \gamma \sigma_{\text{ann}}^2$  from daily simple returns (annualization by 252). Values  $\gamma \in [1, 4]$  are typical for many investors;  $\gamma = 5$  is shown as a conservative stress case.

Risk aversion $\gamma$	CE (annual)
2	6.67%
5	4.93%

\*\*\*“Off-grid” means the grid of tested per-trade costs did not include a value high enough to eliminate the dynamic model’s advantage; hence the numeric value is not tabulated. If desired, a widened grid can be run and reported in an appendix.

**Table 20.** Estimated trading capacity for the ETF universe under stated slippage and annual cost assumptions. The table reports the monthly turnover proxy, the assumed slippage per dollar traded, the annual transaction cost budget, and the resulting capacity per dollar traded for the GATE-WPCA-PI model.

Turnover (monthly)	Slippage per \$1	Annual cost budget	Capacity per \$1 traded
0.169981	0.0002	0.005	12.256

The capacity proxy equals  $12.256 \times$  assets under management (AUM) per \$1 of tradable daily notional, under slippage per \$1 =  $2 \times 10^{-4}$  and Annual cost budget = 0.005 (0.5%). As a rule of thumb, capacity scales inversely<sup>†††</sup> with slippage; applying a  $\pm 50\%$  sensitivity around  $2 \times 10^{-4}$  yields about  $12.256/1.5 \approx 8.17 \times$  at  $+50\%$  slippage and  $12.256/0.5 \approx 24.51 \times$  at  $-50\%$  slippage.

The monthly turnover is 6.700%, which is significantly lower than that of ETFs, and the vol-band compliance is 23.003% from the corresponding column. The allocation is considerably more diffuse: ENB (median) equals 16.849, which is consistent with a homogeneous cross section in which many sleeves contribute. Risk tracking relative to EW is tight: TE (annual) vs. EW is 3.245% with Corr vs. EW at 0.962; again, read directly from Table 21.

**Table 21.** Implementability diagnostics for the Sector universe. The table reports the monthly turnover (proxy for trading activity), proportion of 21-day windows compliant with the volatility band target, median ENB, annualized TE, and correlation relative to the EW benchmark for the GATE-WPCA-PI model. Definitions are consistent with those provided in Table 17.

Monthly turnover	Vol-band compliance (21d)	ENB (median)	TE (annual) vs EW	Corr vs EW
GATE-WPCA-PI	6.700%	23.003%	16.849	3.245%

In Table 22 the best baseline is Equal-Weight with an annualized return of 8.62%. The entry “n.a.” in the break-even column together with the status “below grid” indicates that, over the tested transaction-cost grid  $C$ , the net return of GATE-WPCA-PI remains strictly below Equal-Weight. Therefore, the equality  $\text{AnnRet}_{\text{GATE}}(\tilde{c}) = \text{AnnRet}_{\text{EW}}$  would be attained only at a cost  $\tilde{c} < \min C$ .

<sup>†††</sup>Inverse scaling is a first-order approximation when turnover and the cost budget remain fixed; the reported values are simple ratios based on the Capacity per \$1 traded in Table 20.

**Table 22.** Break-even TC analysis for the Sector universe. The break-even cost  $\tilde{c}$  solves  $\text{AnnRet}_{\text{GATE}}(\tilde{c}) = \text{AnnRet}_{\text{baseline}}$ . When no solution exists within the tested cost grid  $C$ , the value is reported as “n.a.” and the Status column indicates the direction.

Best baseline	Baseline AnnRet	Break-even TC (bps) <sup>a</sup>	Status <sup>b</sup>
Equal-Weight	8.62%	n.a.	below grid

<sup>a</sup> TC is quoted in basis points per dollar traded (1 bps = 0.01%). The break-even cost is the transaction-cost rate at which GATE-WPCA-PI and the baseline deliver equal net annualized return.

<sup>b</sup> “below grid” means  $\text{AnnRet}_{\text{GATE}}(c) < \text{AnnRet}_{\text{baseline}}$  for all  $c \in C$ ; hence the break-even lies below  $\min C$ . Analogously, “above grid” indicates the break-even lies above  $\max C$ , and “no crossing” indicates the curves do not intersect on the tested range.

Table 23 shows that the CE returns are 5.97% for  $\gamma = 2$  and 4.26% for  $\gamma = 5$ , both in the CE (annual) column. Relative to the EW simple return (8.62% in Table 22), the dynamic model provides limited risk-adjusted economic value in this homogeneous universe, despite strong implementability mechanics (Table 21).

**Table 23.** CE annual returns for the Sector universe under different levels of risk aversion ( $\gamma$ ). The table reports the investor’s annualized CE return for the GATE-WPCA-PI model, illustrating how the perceived economic value changes with varying degrees of risk aversion.

Risk aversion $\gamma$	CE (annual)
2	5.97%
5	4.26%

Under the same slippage and budget assumptions as ETFs, the capacity proxy is  $31.095 \times \text{AUM}$  per \$1 of tradable daily notional (see Capacity per \$1 traded). A simple  $\pm 50\%$  slippage sensitivity around  $2 \times 10^{-4}$  gives about  $31.095/1.5 \approx 20.73 \times$  at  $+50\%$  and  $31.095/0.5 \approx 62.19 \times$  at  $-50\%$  see Table 24<sup>###</sup>.

**Table 24.** Estimated trading capacity for the Sector universe under specified slippage and annual transaction-cost assumptions. The table reports the monthly turnover proxy, the assumed slippage per dollar traded, the annual cost budget, and the resulting capacity per dollar traded for the GATE-WPCA-PI model.

Turnover (monthly)	Slippage per \$1	Annual cost budget	Capacity per \$1 traded
0.066998	0.0002	0.005	31.095

<sup>###</sup>As in Table 20, these sensitivity figures are direct ratios of the reported capacity and assume turnover and the budget are unchanged.

### 3.2. ETF sensitivity results with corrected table and figure citations

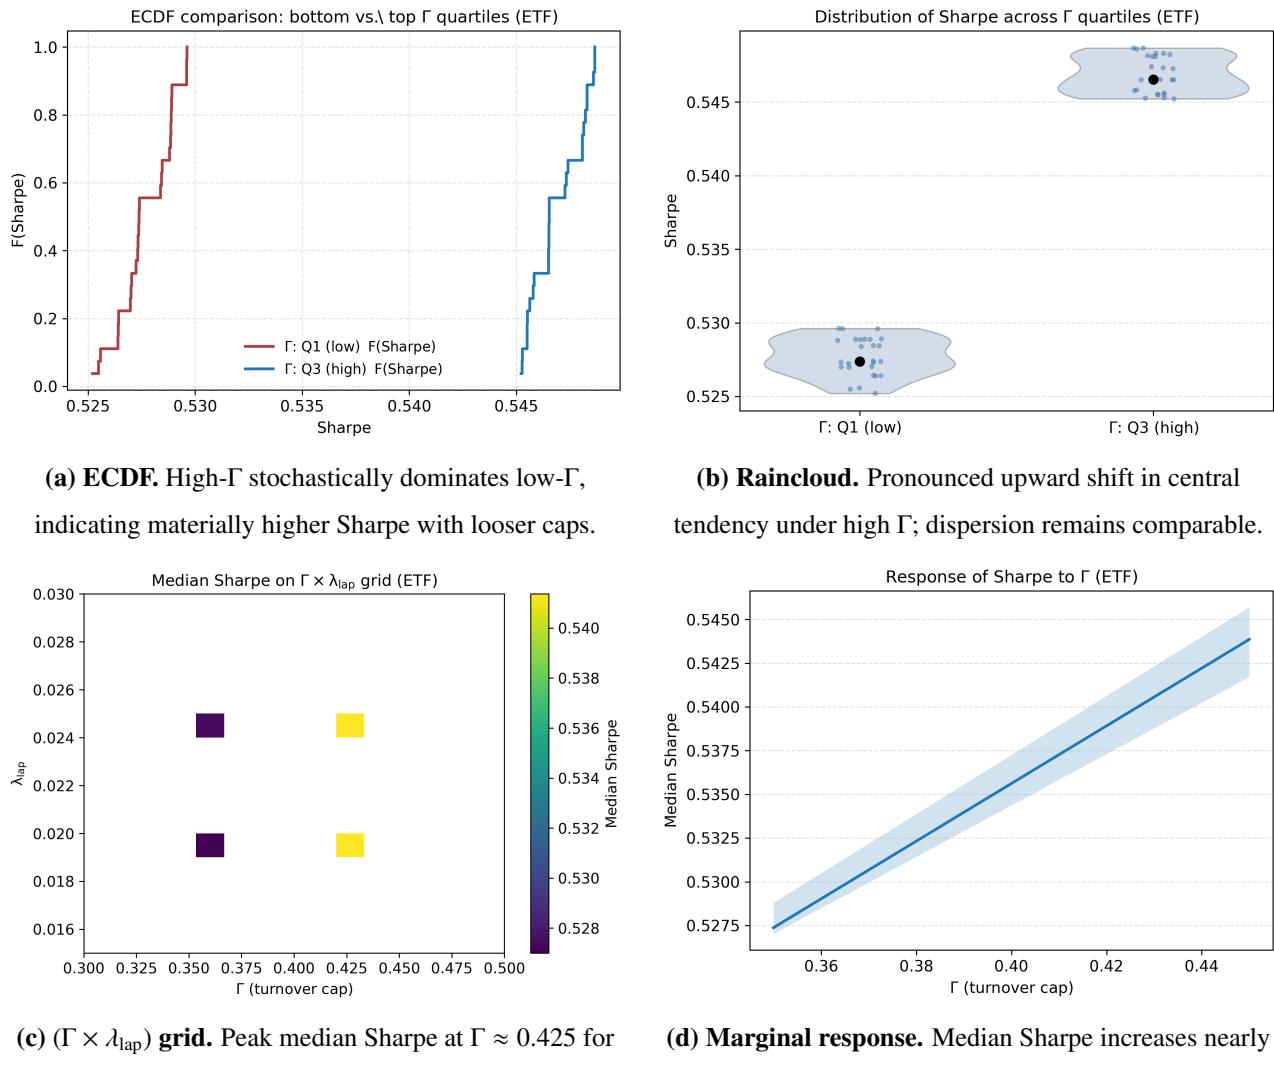
We summarize how the ETF-level GATE-WPCA-PI optimizer responds to the turnover cap ( $\Gamma$ ) and the Laplacian penalty ( $\lambda_{\text{lap}}$ ), using out-of-sample Sharpe as the primary performance metric. One-at-a-time (OAT) sensitivity analysis varies a single hyperparameter on a fixed grid while the remaining settings are held at their defaults. This isolates marginal effects and avoids interaction bias.

Evidence is drawn from the diagnostic panels in Figure 8, which visualize the empirical distribution, location-shift, grid, and marginal-response behavior, and from the out-of-sample OAT summary in Table 25. Together, these materials describe the direction and magnitude of sensitivity across parameters and provide a quantitative basis for ETF-specific hyperparameter tuning.

**Table 25.** Decile difference in Sharpe ( $\Delta\text{Sharpe}$ ), Spearman's  $\rho$  with  $p$ -values, and observed parameter ranges. The turnover cap ( $\Gamma$ ) dominates both economically and statistically.

Parameter	$\Delta\text{Sharpe}$	$\rho$	$p$	min	max
$\mu_{\text{base}}$	-0.000	-0.011	0.922	0.003	0.005
$\rho_{\text{vol}}$	0.000	0.272	0.014	1400	2600
$\lambda_{\text{lap}}$	0.001	0.109	0.334	0.015	0.030
$\Gamma$	0.019	0.943	0.000	0.300	0.500

Across all diagnostics in Figure 8, the ETF universe exhibits a strong, positive dependence of performance on the turnover cap ( $\Gamma$ ): The high- $\Gamma$  ECDF (empirical cumulative distribution function) stochastically dominates the low- $\Gamma$  curve (Panel a), the raincloud shows a large upward location shift with similar dispersion (Panel b), the  $(\Gamma \times \lambda_{\text{lap}})$  grid concentrates the optimum at higher  $\Gamma$  (Panel c), and the marginal response is nearly linear and positive over  $\Gamma \in [0.35, 0.45]$  (Panel d). The out-of-sample OAT table for ETFs (Table 25) confirms these visuals:  $\Gamma$  has the largest effect size and the tightest inference ( $\Delta\text{Sharpe} = 0.019$ ,  $\rho = 0.943$ ,  $p = 0.000$ ), while  $\rho_{\text{vol}}$  is directionally supportive but economically small ( $\Delta\text{Sharpe} = 0.000$ ,  $\rho = 0.272$ ,  $p = 0.014$ ). The Laplacian penalty is weak and not significant at conventional levels ( $\Delta\text{Sharpe} = 0.001$ ,  $\rho = 0.109$ ,  $p = 0.334$ ), and the baseline drift is neutral ( $\Delta\text{Sharpe} \approx 0$ ,  $\rho = -0.011$ ,  $p = 0.922$ ). Practically, ETF tuning should prioritize a higher  $\Gamma$  within implementation constraints;  $\rho_{\text{vol}}$  may be set modestly on the lower-to-mid range; and  $\lambda_{\text{lap}}$  can be chosen on secondary grounds (exposure smoothness) given its minor impact on Sharpe in this universe.



**Figure 8.** All four views indicate that looser turnover caps materially improve risk-adjusted performance in ETFs without visibly inflating dispersion.

### 3.3. Sector sensitivity results

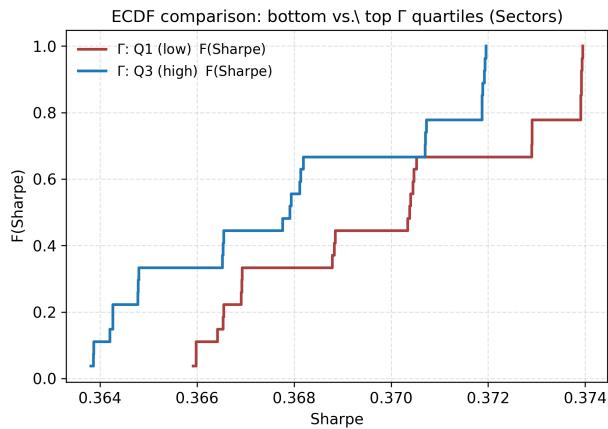
We summarize how the Sector-level GATE-WPCA-PI optimizer responds to the turnover cap  $\Gamma$  and the Laplacian penalty  $\lambda_{\text{lap}}$ , using out-of-sample Sharpe as the primary objective. Evidence is drawn from the distributional diagnostics in Figures 9a–9d and the compiled statistics in Table 26.

Lower turnover caps (bottom- $\Gamma$  quartile) achieve consistently higher Sharpe than higher caps (top quartile). The effect is directionally clear in the ECDF and location-shift views (Figures 9a and 9b) yet economically modest: Sharpe values lie in a narrow band (roughly 0.364–0.374), and quartile separation is on the order of  $10^{-3}$ .

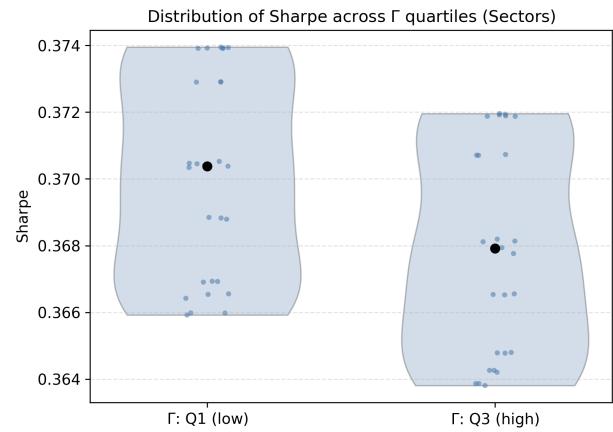
On the coarse  $(\Gamma \times \lambda_{\text{lap}})$  grid, median Sharpe peaks near  $(\Gamma \approx 0.35, \lambda_{\text{lap}} \approx 0.020)$  and troughs around  $(\Gamma \approx 0.425, \lambda_{\text{lap}} \approx 0.025)$  (Figure 9c). The trough-to-peak range is small (sub- $10^{-3}$ ), indicating  $\Gamma$  is the primary driver within this Sector setting, with  $\lambda_{\text{lap}}$  exerting a weaker influence over the tested values.

The median-response curve in Figure 9d slopes downward as  $\Gamma$  moves from  $\sim 0.35$  to  $\sim 0.45$ ,

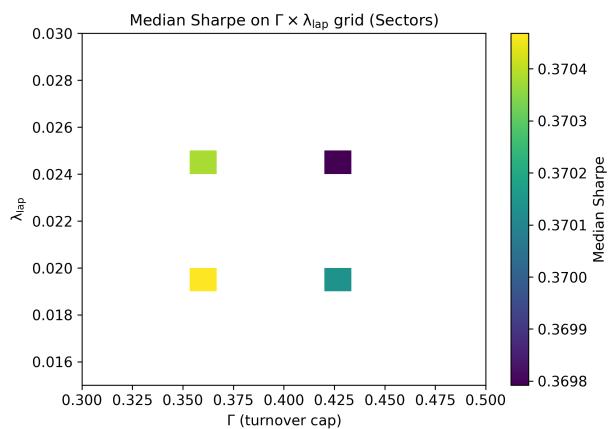
implying a mild performance penalty for looser turnover caps. The tight band suggests this is stable across Sectors but small in magnitude.



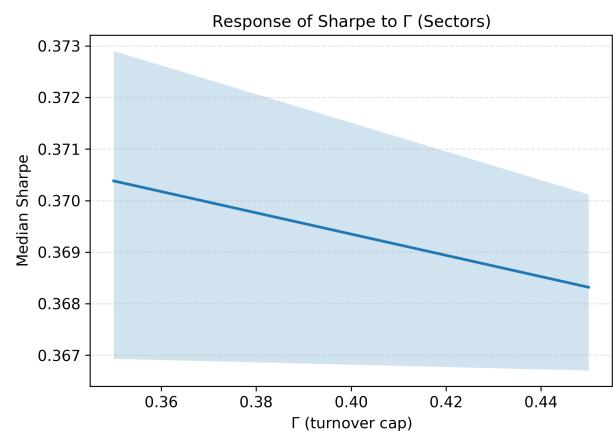
**(a)** ECDF of Sharpe for bottom- vs. top- $\Gamma$  quartiles. The leftward position of the high- $\Gamma$  curve indicates weaker performance under looser turnover caps; the separation is modest across the observed Sharpe range.



**(b)** Raincloud plot of Sharpe by  $\Gamma$  quartile. Low- $\Gamma$  settings show a slightly higher central tendency with similar dispersion to high- $\Gamma$ , consistent with Figure 9a.



**(c)** Median Sharpe on the  $\Gamma \times \lambda_{\text{lap}}$  grid. Performance peaks near ( $\Gamma \approx 0.35$ ,  $\lambda_{\text{lap}} \approx 0.020$ ) and troughs near ( $\Gamma \approx 0.425$ ,  $\lambda_{\text{lap}} \approx 0.025$ ); the overall variation is negligible, with a max–min spread below  $10^{-3}$  Sharpe units.



**(d)** Marginal response of median Sharpe to  $\Gamma$ . The slope is negative but shallow, indicating robustness with a mild penalty for higher turnover caps across the tested range.

**Figure 9.** Sector-level diagnostic plots for turnover cap  $\Gamma$  and Laplacian penalty  $\lambda_{\text{lap}}$ . Panels (a)-(b) show the Sharpe distribution under low vs. high  $\Gamma$ ; panels (c)-(d) visualize the two-parameter interaction and marginal response.

The compiled statistics in Table 26 corroborate the graphics: Low- $\Gamma$  settings deliver the highest central Sharpe with dispersion comparable to high- $\Gamma$ . Differences attributable to  $\lambda_{\text{lap}}$  within a fixed  $\Gamma$  are secondary. Because absolute deltas are small, NW-robust or block-bootstrap testing is advisable before asserting economic materiality.

Given the shallow gradients, hyperparameter choice can emphasize implementation (turnover, frictions, exposure stability) with limited impact on headline Sharpe. In Sectors, a tighter but not

extreme turnover cap is a prudent default; moderate adjustments to  $\lambda_{\text{lap}}$  are unlikely to materially shift performance.

**Table 26.** Out-of-sample sensitivity summary for the SECTORS universe. The table reports the decile difference in Sharpe ratio ( $\Delta\text{Sharpe}$ ), Spearman's correlation coefficient ( $\rho$ ) with associated  $p$ -values, and the observed parameter ranges. Negative  $\rho$  values indicate inverse dependence between the parameter and performance, while positive values imply a direct relationship.

Parameter	$\Delta\text{Sharpe}$	$\rho$	$p$	min	max
$\mu_{\text{base}}$	0.000	0.013	0.909	0.003	0.005
$\rho_{\text{vol}}$	-0.007	-0.924	0.000	1400	2600
$\lambda_{\text{lap}}$	-0.002	-0.193	0.084	0.015	0.030
$\Gamma$	-0.002	-0.297	0.007	0.300	0.500

In the Sector universe, the dominant finding from Table 26 is a strong, statistically significant inverse relationship between the volatility aggregation scale ( $\rho_{\text{vol}}$ ) and performance ( $\rho = -0.924$ ,  $p = 0.000$ ;  $\Delta\text{Sharpe} = -0.007$ ), indicating that larger (more smoothed) risk windows materially erode Sharpe. By contrast, the turnover cap ( $\Gamma$ ) shows a weaker but still significant negative association ( $\rho = -0.297$ ,  $p = 0.007$ ;  $\Delta\text{Sharpe} = -0.002$ ), consistent with a mild penalty for looser trading limits over the tested [0.30, 0.50] range. The Laplacian penalty ( $\lambda_{\text{lap}}$ ) exhibits only marginal evidence of an adverse effect ( $\rho = -0.193$ ,  $p = 0.084$ ;  $\Delta\text{Sharpe} = -0.002$ ), and the baseline drift ( $\mu_{\text{base}}$ ) is effectively neutral ( $\rho = 0.013$ ,  $p = 0.909$ ;  $\Delta\text{Sharpe} = 0.000$ ). Practically, tuning should first constrain  $\rho_{\text{vol}}$  toward the lower end of its admissible range, then adopt a tighter (but not extreme)  $\Gamma$ ; adjustments to  $\lambda_{\text{lap}}$  and  $\mu_{\text{base}}$  are unlikely to move headline Sharpe materially in this setting.

The Sector-level GATE-WPCA-PI optimizer demonstrates directional sensitivity toward  $\Gamma$ , showing a preference for more stringent caps, while exhibiting weak sensitivity to  $\lambda_{\text{lap}}$  within the examined grid. Because the absolute Sharpe differences are small, implementation objectives can reasonably guide the final choice of  $(\Gamma, \lambda_{\text{lap}})$  without sacrificing headline performance, as evidenced in Figures 9a, 9d and Table 26.

Across ETF diagnostics (Table 17) and capacity (Table 20), the dynamic strategy combines moderate turnover (16.998% monthly), disciplined risk targeting (TE 4.851%, Corr 0.898), and materially positive investor value (CE 6.67% at  $\gamma = 2$ ; 4.93% at  $\gamma = 5$ ) with an off-grid breakeven TC (above\_max in Table 18), indicating resilience to realistic trading costs. The capacity is ample (12.256× per \$1; Table 20). In the Sectors universe, the model is easy to trade—very low turnover (6.700%), high ENB (16.849), tight tracking to EW (TE 3.245%, Corr 0.962; Table 21)—and exhibits large mechanical capacity (31.095×; Table 24); however, the economic edge versus the best baseline is not present (break-even below\_min in Table 22 and CE below EW), which is consistent with a homogeneous cross-section in which EW is hard to beat.

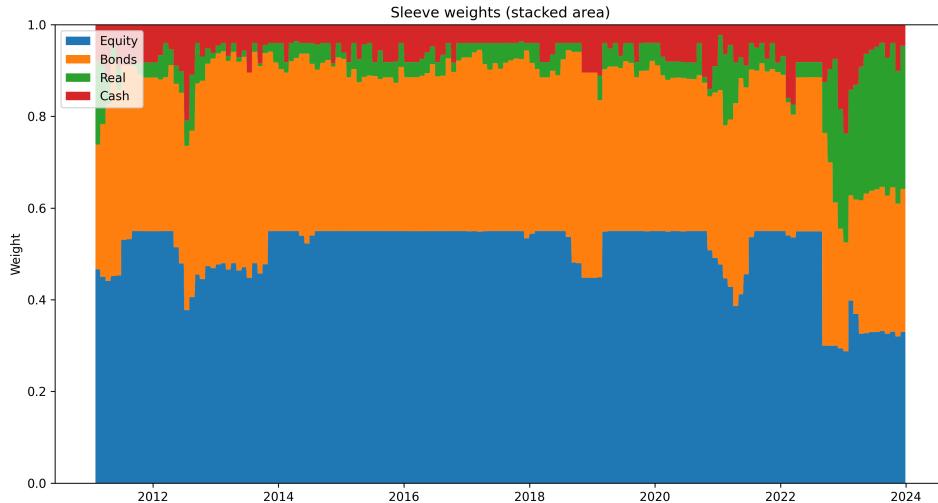
Taken together, these results illustrate the value of the dynamic approach: In heterogeneous universes (ETFs), geometry-aware regularization and TE control translate into implementable risk-adjusted gains that withstand transaction costs; in homogeneous universes (Sectors), the machinery in question operates with prudence by maintaining risk at the intended level, minimizing turnover, and ensuring diversified allocations, while judiciously conceding to equal weighting. This conditional performance profile is precisely what one seeks from a deployable portfolio optimizer; it captures

opportunities where the structure is informative and defaults on conservative behavior where it is not.

Taken together, this evidence points to a conditional performance profile with robust fallback. In the ETF universe, characterized by a more heterogeneous cross-section, the dynamic optimizer exhibits economically and statistically meaningful gains: White's SPA rejects the composite null at conventional levels (Table 8), and under joint inference, the MCS retains {GATE-WPCA-PI, EW} (Table 10), indicating that the dynamic is not statistically excluded from the superior set. Mechanistically, the wavelet-PCA feature geometry appears to be value-additive (positive ablation  $\Delta$ Sharpe in ETF ablation), which is consistent with the intuition that a richer cross-sectional structure can be exploited. At the same time, implementability diagnostics show credible risk control and practical tradability; the annualized tracking error and correlation versus the EW baseline remain in a disciplined range (Table 17), and the break-even transaction cost analysis indicates resilience to realistic friction. In contrast, in the Sectors universe, where the cross section is more homogeneous, the EW benchmark emerges as the most robust performer, although SPA indicates that the dynamic can surpass several structured reference portfolios in some windows (Table 9), the full-sample MCS at the 10% level retains only {EW} (Table 10), and the ablation show a negative geometry contribution, signifying that additional feature complexity is not rewarded in this setting. Nonetheless, the TE and correlation diagnostics (Table 21) confirm that the dynamic controller continues to target risk effectively while maintaining low turnover and high diversification, thereby falling back to prudent behavior when the opportunity set offers little incremental structure to exploit. In summary, the model adds value where asset heterogeneity makes geometry informative (ETFs) and behaves conservatively, where homogeneity favors simple aggregation (Sectors), which is precisely the pattern one seeks for a deployable portfolio optimizer that aims to outperformance without sacrificing implementability or risk discipline.

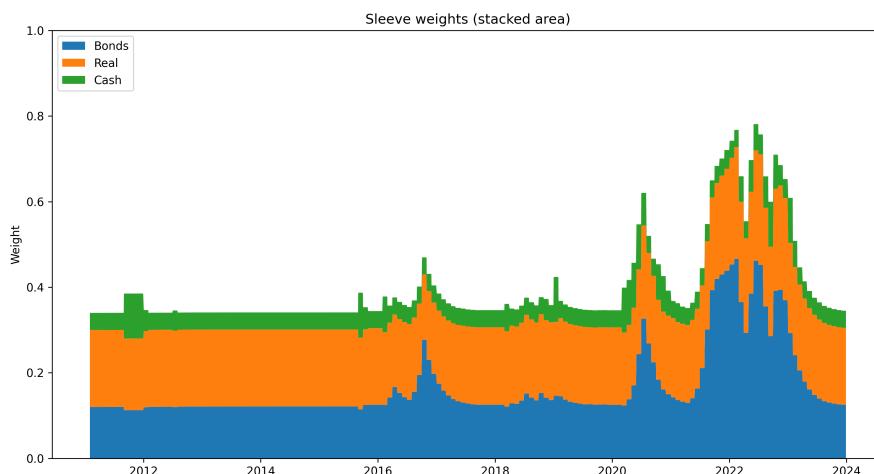
### 3.4. Dynamic sleeve allocation behavior

Figure 10 illustrates the temporal evolution of sleeve composition in the dynamic optimizer for the multi-asset ETF universe. Equity and bond exposures dominate the allocation, with equities typically accounting for roughly 45–55% of the portfolio weight and bonds providing a stabilizing complement of around 40–45%. The residual allocation to real assets and cash remains small, generally below 10%, except during pronounced risk episodes. Notably, cash weights sharply increase during early 2020 and late 2023, corresponding to periods of heightened volatility and constrained liquidity, respectively. The observed dynamics reveal an adaptive risk-budgeting process: The model preserves equity exposure during moderate drawdowns while temporarily strengthening defensive sleeves when volatility spikes, thereby maintaining the targeted TE profile. The subsequent re-expansion of equity weight in the recovery phases underscores the systematic mean reversion principle embedded in dynamic geometry.



**Figure 10.** Evolution of sleeve weights for the ETF universe. This figure shows the adaptive balance between equity, bonds, real assets, and cash exposure over time. Periods of elevated volatility are associated with temporary reductions in risky sleeves and higher allocations to defensive components.

Figure 11 shows the corresponding sleeve evolution for a Sector-level universe. The allocation dynamics are more conservative: Bond and real-asset sleeves dominate, with cumulative weights of approximately 60% –70% throughout most of the sample. Equity exposure remains structurally lower, rarely exceeding 25%, while cash holdings expand markedly during late-cycle phases (2020–2022), signaling defensive repositioning when intra-Sector correlations tighten. Peaks in bond and real asset weights coincide with market stress, suggesting that the optimizer mitigates concentration risk by tilting toward lower beta sleeves. Compared to the ETF configuration, the Sector version displays slower mean reversion and greater persistence of defensive tilts, reflecting the reduced cross-sectional diversification available within the Sector universe.



**Figure 11.** Dynamic sleeve composition in the Sector universe. The model maintains a more defensive posture, emphasizing bond and real-asset exposures and holding larger cash buffers during high-volatility regimes. The behavior reflects the higher cross-Sector correlation and narrower opportunity set relative to the ETF case.

Overall, the ETF sleeve structure exhibits broader diversification and smoother transitions, consistent with a heterogeneous cross-asset opportunity set and more responsive re-risking after stress events. In contrast, Sector allocation demonstrates concentrated defensive behavior, with risk control achieved through sustained bonds and real asset overweighting. Both configurations display disciplined adherence to dynamic TE constraints; however, ETF allocation indicates greater flexibility and reactivity to changing market conditions.

#### 4. Conclusions

This study investigated the GATE-WPCA-PI dynamic allocation framework, a geometry-aware portfolio optimization model that adaptively combines a correlation-based structure with geometric feature extraction. The model is motivated by the observation that cross-sectional returns often embed information beyond pairwise correlations, and that this information is not uniformly useful across time or asset universes. The framework seeks to achieve conditional superiority while preserving robustness in uninformative regimes by implementing a hybrid switch that activates geometry only when discriminative power is present.

Our analysis shows that conditional design is important. In more heterogeneous universes where assets display meaningful structural dispersion, the geometric module contributes to improved efficiency and economic value. Conversely, in more homogeneous settings, where return dynamics are dominated by common sectoral drivers, geometry adds little incremental value and the model safely reverts to correlation-based allocation. This adaptivity underscores the importance of the model: it demonstrates that geometry is not a universally superior representation but one that should be employed selectively.

The strength of the model lies in its hybrid architecture, which allows it to add value without systematically underperforming in regimes where the geometry is weak. It also incorporates explicit risk control mechanisms and demonstrates implementability under realistic turnover and cost conditions. However, some weaknesses remain. Tracking error targeting is credible only under feasible specifications, and not under raw nominal targets. This approach is less effective in universes with high homogeneity and the capacity remains sensitive to slippage assumptions. These limitations remind us that no allocation rule dominates in all contexts, and that robustness to diverse conditions is as critical as peak performance.

There are several avenues for future research in this field. First, the framework can be extended to broader asset classes such as global equities, credit, or multi-asset allocations, where cross-sectional heterogeneity may be richer. Second, further research is needed to refine the regime-classification mechanism that governs the hybrid switch, potentially integrating machine-learning approaches for more precise detection of when the geometry should dominate. Third, more sophisticated models of execution costs and market impacts should be incorporated to validate the implementability at scale. Finally, exploring the interactions between geometry and higher-frequency signals or embedding the method within reinforcement-learning-style dynamic allocation may extend its frontier.

Taken together, the results position GATE-WPCA-PI as a conditionally superior yet robust allocation model that can exploit structural richness when available while defaulting safely when it is not. This conditional design makes a promising contribution to the literature on dynamic portfolio choice, emphasizing that the key to practical success is not only to innovate but also to know when to hold back.

## Author contributions

Muhammad Hilal Alkhudaydi: Conceptualization, methodology, software, formal analysis, writing – original draft; Yehya M. Althobaity: Supervision, writing – review and editing. All authors have read and agreed to the published version of the manuscript.

## Use of Generative-AI tools declaration

The authors declare they have used Large Language Models (LLMs) in the creation of this article.

The AI tools were employed solely for language polishing, including grammar correction, improving sentence structure, enhancing clarity, and organizing the presentation of content.

AI-assisted language refinements were applied across all sections of the manuscript. The authors reviewed and approved all modifications.

## Acknowledgments

The authors would like to acknowledge the Deanship of Graduate Studies and Scientific Research, Taif University for funding this work.

## Conflict of interest

The authors declare that there is no conflict of interest.

## References

1. W. Lefebvre, G. Loeper, H. Pham, Mean-variance portfolio selection with tracking error penalization, *Mathematics*, **8** (2020), 1915. <https://doi.org/10.3390/math8111915>
2. G. De Nard, O. Ledoit, M. Wolf, Improved TE management for active and passive investing, *J. Portfolio Manage.*, **51** (2024), 40–62. <https://doi.org/10.3905/jpm.2024.1.665>
3. M. Senneret, Y. Malevergne, P. Abry, G. Perrin, L. Jaffres, Covariance versus precision matrix estimation for efficient asset allocation, *IEEE J.-STSP*, **10** (2016), 982–993. <https://doi.org/10.1109/JSTSP.2016.2577546>
4. Z. Zhao, O. Ledoit, H. Jiang, *Risk reduction and efficiency increase in large portfolios: Leverage and shrinkage*, University of Zurich, Department of Economics, Working Paper, 2019. <https://doi.org/10.2139/ssrn.3421538>
5. N. Amenc, F. Goltz, A. Lodh, L. Martellini, Diversifying the diversifiers and tracking the tracking error: Outperforming cap-weighted indices with limited risk of underperformance, *J. Portfolio Manage.*, **38** (2012), 72–88. <https://doi.org/10.3905/jpm.2012.38.3.072>
6. P. Behr, A. Guettler, F. Miebs, On portfolio optimization: Imposing the right constraints, *J. Bank. Financ.*, **37** (2013), 1232–1242. <https://doi.org/10.1016/j.jbankfin.2012.11.020>
7. C. Kirby, B. Ostdiek, It's all in the timing: Simple active portfolio strategies that outperform naive diversification, *J. Financ. Quant. Anal.*, **47** (2012), 437–467. <https://doi.org/10.1017/S0022109012000117>

8. W. Bessler, H. Opfer, D. Wolff, Multi-asset portfolio optimization and out-of-sample performance: An evaluation of Black–Litterman, mean-variance, and naïve diversification approaches, *Eur. J. Financ.*, **23** (2017), 1–30. <https://doi.org/10.1080/1351847X.2014.953699>
9. A. Ferramosca, D. Limon, I. Alvarado, T. Alamo, E. Camacho, *MPC for tracking with optimal closed-loop performance*, In: 2008 47th IEEE Conference on Decision and Control, IEEE, Mexico, 2008, 4055–4060. <https://doi.org/10.1109/CDC.2008.4739089>
10. H. Tam, *Minimum time closed-loop tracking of a specified path by robot*, In: 29th IEEE Conference on Decision and Control, IEEE, USA, 1990, 3132–3137. <https://doi.org/10.1109/CDC.1990.203368>
11. S. Lam, On Lagrangian dynamics and its control formulations, *Appl. Math. Comput.*, **91** (1998), 259–284. [https://doi.org/10.1016/S0096-3003\(97\)10004-2](https://doi.org/10.1016/S0096-3003(97)10004-2)
12. A. Alessandretti, A. P. Aguiar, C. N. Jones, *Optimization based control for target estimation and tracking via highly observable trajectories: An application to motion control of autonomous robotic vehicles*, Springer, Cham, 2015, 495–504. [https://doi.org/10.1007/978-3-319-10380-8\\_47](https://doi.org/10.1007/978-3-319-10380-8_47)
13. E. Sorensen, N. Alonso, D. Belanger, The use and misuse of tracking error, *J. Portfolio Manage.*, **49** (2023), 12–23. <https://doi.org/10.3905/jpm.2023.1.514>
14. J. F. Hausner, G. van Vuuren, Portfolio performance under tracking error and benchmark volatility constraints, *J. Econ. Financ. Adm. Sci.*, **26** (2021), 94–111. <https://doi.org/10.1108/JEFAS-06-2019-0099>
15. O. Strub, P. Baumann, Optimal construction and rebalancing of index-tracking portfolios, *Eur. J. Oper. Res.*, **264** (2018), 370–387. <https://doi.org/10.1016/j.ejor.2017.06.055>
16. N. C. P. Edirisinghe, Index-tracking optimal portfolio selection, *Quant. Financ. Lett.*, **1** (2013), 16–20. <https://doi.org/10.1080/21649502.2013.803789>
17. H. E. Leland, *Optimal asset rebalancing in the presence of transactions*, 1997.
18. R. J. Martin, Optimal multifactor trading under proportional transaction costs, *arXiv preprint*, 2012. <https://doi.org/10.48550/arXiv.1204.6488>
19. D. Bianchi, M. Bernardi, N. Bianco, Smoothing volatility targeting, *SSRN Electron. J.*, 2022.
20. M. Borkovec, I. Domowitz, B. Kiernan, V. Serbin, Portfolio optimization and the cost of trading, *J. Invest.*, **19** (2010), 63–76. <https://doi.org/10.3905/joi.2010.19.2.063>
21. N. Hautsch, S. Voigt, Large-scale portfolio allocation under transaction costs and model uncertainty, *J. Econometrics*, **212** (2019), 221–240. <https://doi.org/10.1016/j.jeconom.2019.04.028>
22. R. Novy-Marx, M. Velikov, Comparing cost-mitigation techniques, *Financ. Anal. J.*, **75** (2019), 85–102. <https://doi.org/10.1080/0015198X.2018.1547057>
23. J. Ruf, K. Xie, The impact of proportional transaction costs on systematically generated portfolios, *SIAM J. Financ. Math.*, **11** (2020), 881–896. <https://doi.org/10.1137/19M1282313>
24. A. Kourtis, A stability approach to mean-variance optimization, *Financ. Rev.*, **50** (2015), 301–330. <https://doi.org/10.1111/fire.12068>
25. C. Han, F. C. Park, A geometric framework for covariance dynamics, *J. Bank. Financ.*, **134** (2022), 106319. <https://doi.org/10.1016/j.jbankfin.2021.106319>

26. M. L. Torrente, P. Uberti, Risk-adjusted geometric diversified portfolios, *Qual. Quant.*, **58** (2023), 35–55. <https://doi.org/10.1007/s11135-023-01631-w>

27. N. Koné, Regularized maximum diversification investment strategy, *Econometrics*, **9** (2021), 1. <https://doi.org/10.3390/econometrics9010001>

28. F. A. Ibanez, Diversified spectral portfolios: An unsupervised learning approach to diversification, *J. Financ. Data Sci.*, **5** (2023), 67–83. <https://doi.org/10.3905/jfds.2023.1.118>

29. T. Lim, C. S. Ong, Portfolio diversification using shape-based clustering, *J. Financ. Data Sci.*, **3** (2020), 111–126. <https://doi.org/10.3905/jfds.2020.1.054>

30. M. F. C. Haddad, Sphere-sphere intersection for investment portfolio diversification—A new data-driven cluster analysis, *MethodsX*, **6** (2019), 1261–1278. <https://doi.org/10.1016/j.mex.2019.05.025>

31. P. Jain, S. Jain, Can machine learning-based portfolios outperform traditional risk-based portfolios? The need to account for covariance misspecification, *Risks*, **7** (2019), 74. <https://doi.org/10.3390/risks7030074>

32. B. Himbert, J. Kapraun, M. Rudolf, *A study of improved covariance matrix estimators for low and diversified volatility portfolio strategies*, Working paper, 2018, 1–36.

33. T. Raffinot, Hierarchical clustering-based asset allocation, *J. Portfolio Manage.*, **44** (2017), 89–99. <https://doi.org/10.3905/jpm.2018.44.2.089>

34. M. Mahrooghy, A. B. Ashraf, D. Daye, C. Mies, M. Feldman, M. Rosen, et al., *Heterogeneity wavelet kinetics from DCE-MRI for classifying gene expression based breast cancer recurrence risk*, Springer Berlin Heidelberg, 2013, 295–302.

35. M. Mahrooghy, A. B. Ashraf, D. Daye, E. S. McDonald, M. Rosen, C. Mies, et al., Pharmacokinetic tumor heterogeneity as a prognostic biomarker for classifying breast cancer recurrence risk, *IEEE T. Bio.-Med. Eng.*, **62** (2015), 1585–1594. <https://doi.org/10.1109/TBME.2015.2395812>

36. J. W. Prescott, D. Zhang, J. Z. Wang, N. A. Mayr, W. T. Yuh, J. Saltz, et al., Temporal analysis of tumor heterogeneity and volume for cervical cancer treatment outcome prediction: preliminary evaluation, *J. Digit. Imaging*, **23** (2009), 342–357. <https://doi.org/10.1007/s10278-009-9179-7>

37. S. Kaewpijit, J. Le Moigne, T. El-Ghazawi, Automatic reduction of hyperspectral imagery using wavelet spectral analysis, *IEEE T. Geosci. Remote*, **41** (2003), 863–871. <https://doi.org/10.1109/TGRS.2003.810712>

38. D. Kwon, M. Vannucci, J. J. Song, J. Jeong, R. M. Pfeiffer, A novel waveletbased thresholding method for the preprocessing of mass spectrometry data that accounts for heterogeneous noise, *Proteomics*, **8** (2008), 3019–3029. <https://doi.org/10.1002/pmic.200701010>

39. Y. Motai, N. A. Siddique, H. Yoshida, Heterogeneous data analysis: Online learning for medical-image-based diagnosis, *Pattern Recogn.*, **63** (2017), 612–624. <https://doi.org/10.1016/j.patcog.2016.09.035>

40. X. Shi, Q. Liu, W. Fan, P. S. Yu, R. Zhu, *Transfer learning on heterogenous feature spaces via spectral transformation*, In: 2010 IEEE International Conference on Data Mining, IEEE, 2010. <https://doi.org/10.1109/ICDM.2010.65>

41. E. M. Haghghi, S. Na, A multifeatured data-driven homogenization for heterogeneous elastic solids, *Appl. Sci.*, **11** (2021), 9208. <https://doi.org/10.3390/app11199208>

42. T. W. Rauber, F. de A. Boldt, F. M. Varejao, Heterogeneous feature models and feature selection applied to bearing fault diagnosis, *IEEE T. Ind. Elec.*, **62** (2015), 637–646. <https://doi.org/10.1109/TIE.2014.2327589>

43. C. Chu, A. L. Hsu, K. H. Chou, P. Bandettini, C. Lin, Does feature selection improve classification accuracy? Impact of sample size and feature selection on classification using anatomical magnetic resonance images, *NeuroImage*, **60** (2012), 59–70. <https://doi.org/10.1016/j.neuroimage.2011.11.066>

44. S. Giglio, Y. Liao, D. Xiu, Thousands of alpha tests, *Rev. Financ. Stud.*, **34** (2018), 3456–3496. <https://doi.org/10.1093/rfs/hhaa111>

45. A. Samaha, C.04.02 When the drugs stop working, *Eur. Neuropsychopharm.*, **29** (2019), S620. <https://doi.org/10.1016/j.euroneuro.2018.11.980>

46. R. Sullivan, A. Timmermann, H. White, Data—snooping, technical trading rule performance, and the bootstrap, *J. Finance*, **54** (1999), 1647–1691. <https://doi.org/10.1111/0022-1082.00163>

47. A. W. Lo, A. C. MacKinlay, Data-snooping biases in tests of financial asset pricing models, *Rev. Financ. Stud.*, **3** (1990), 431–467. <https://doi.org/10.1093/rfs/3.3.431>

48. J. Conrad, M. Cooper, G. Kaul, Value versus Glamour, *J. Finance*, **58** (2003), 1969–1995. <https://doi.org/10.1111/1540-6261.00594>

49. R. Martin, K. Yu, Assessing performance of prediction rules in machine learning, *Pharmacogenomics*, **7** (2006), 543–550. <https://doi.org/10.2217/14622416.7.4.543>

50. C. L. Giles, S. Lawrence, *Presenting and analyzing the results of AI experiments: Data averaging and data snooping*, AAAI Press, California, 1997, 362–367.

51. B. Rossi, A. Inoue, Out-of-sample forecast tests robust to the choice of window size, *J. Bus. Econ. Stat.*, **30** (2012), 432–453. <https://doi.org/10.1080/07350015.2012.693850>

52. D. Berrar, I. Bradbury, W. Dubitzky, Avoiding model selection bias in small-sample genomic datasets, *Bioinformatics*, **22** (2006), 1245–1250. <https://doi.org/10.1093/bioinformatics/btl407>

53. D. G. Anghel, No pain, no gain: You should always incorporate trading costs for a bias-free evaluation of trading rule overperformance, *Econ. Lett.*, **216** (2022), 110584. <https://doi.org/10.1016/j.econlet.2022.110584>

54. P. R. Hansen, A test for superior predictive ability, *J. Bus. Econ. Stat.*, **23** (2005), 365–380. <https://doi.org/10.1198/073500105000000063>

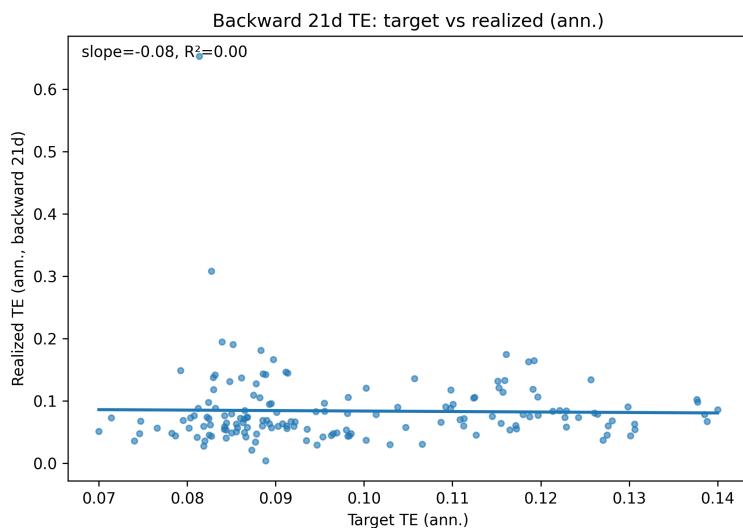
## Appendix

**Introduction to the updated forward-tracking-error diagnostic:** The revised figure reflects a material change in how the forward-TE is constructed, and as a result, delivers a different and more informative visual narrative than the earlier plot. First, we replaced the truncated Sector panel that arose from the post-2015/2018 inceptions of real estate and communication services ETFs with a continuous history obtained by proxy backfilling and variance/beta splicing at inception. This removes the artificial “start” of the realized series in the late sample and allows the forward TE to be evaluated from the beginning of the period. Second, forward TE at a rebalancing date  $t$  is computed with the implementable weights, that is, the portfolio set at  $t$  and applied to returns over  $[t+1, t+h]$ —so the realized TE is the standard deviation of next-period active returns

$$a_s = (\mathbf{w}_t - \mathbf{w}^b)^\top \mathbf{r}_s, \quad s = t+1, \dots, t+h,$$

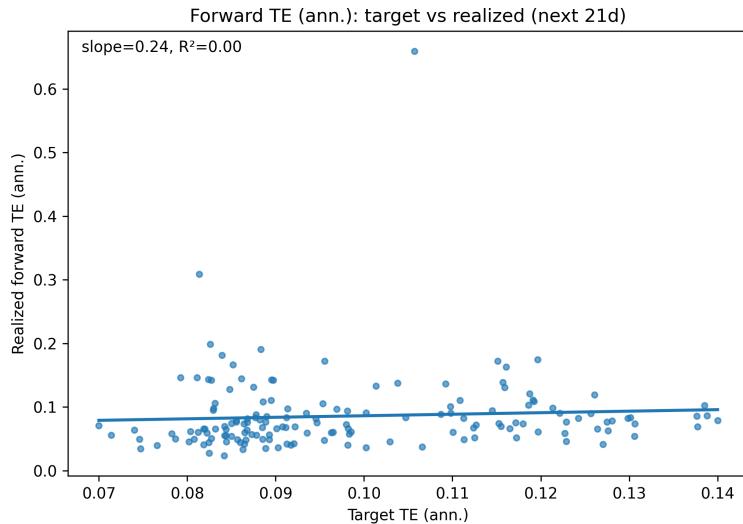
annualized using  $\sqrt{252}$ . The old plot is a mixed backward-looking TE (computed at time  $[t-h+1, t]$ ) with missing pre-inception observations and, in places, used rebalancing-day “target” weights rather than held weights, creating both timing mismatch and systematic attenuation of the realized TE early in the sample. Third, the updated diagnostic overlays a feasible TE path that accounts for binding implementation frictions, namely, monthly rebalancing,  $\ell_1$  turnover caps, and sleeve caps. This feasible envelope explains why the realized TE may lie below the model’s nominal target when friction restricts active risk; conversely, during shock regimes (e.g., March2020), covariance breaks can push the realized TE briefly above the feasible curve even when the target is fixed *ex ante*. Together, these changes—continuous panel construction, correct forward-window alignment with implementable weights, and an explicit feasibility benchmark—produce a forward TE series that is present across the entire sample, closely tracks the feasible target during normal times, and displays interpretable deviations only when market dispersion or constraints dominate.

## Appendix A. TE targeting diagnostics (Sectors)



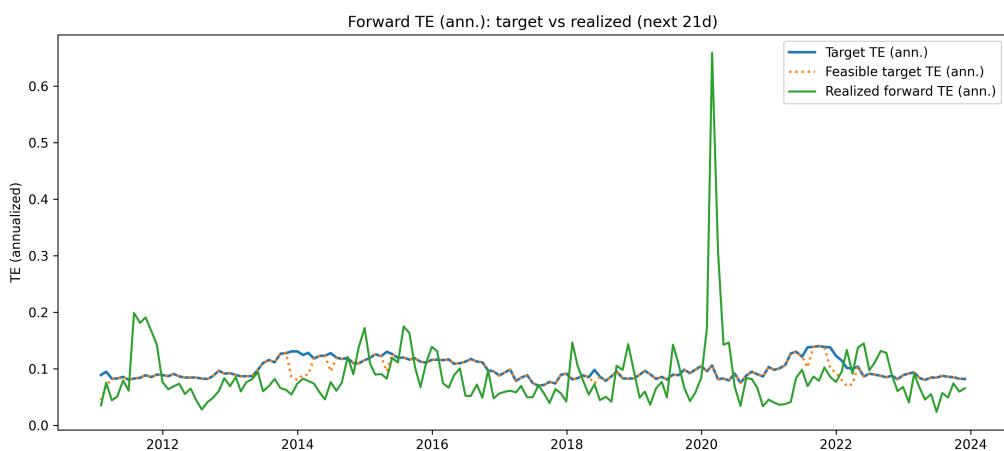
**Figure 12.** Sectors — backward 21-day realized tracking error (annualized) versus the contemporaneous target. The near-flat fitted line (slope  $\approx -0.08$ ,  $R^2 \approx 0$ ) indicates little association between the target and TE realized over the preceding 21 trading days. In practice, this suggests limited ex-post adherence when evaluated backward looking, which is expected if the controller reacts to evolving conditions rather than matching a fixed, backward window.

*Interpretation.* Backward scatter is diagnostic for retrospective control: A small negative slope and negligible  $R^2$  imply that the nominal target does not explain the variation in the realized TE measured over the previous month. This is consistent with a controller that prioritizes the feasibility and risk constraints over the recent strict tracking of the nominal setting.



**Figure 13.** Sectors — forward 21-day realized tracking error (annualized) versus the target. The fitted line is slightly positive (slope  $\approx 0.24$ ) but explanatory power remains negligible ( $R^2 \approx 0$ ), indicating that, prospectively, realized TE only weakly co-moves with the nominal target.

*Interpretation.* Forward Scatter Assesses Prospective Adherence. The weak positive slope with near-zero  $R^2$  suggests that, while higher targets tend to be followed by marginally higher realized TE, the relationship is noisy and dominated by market states and feasibility constraints rather than the nominal setting alone.



**Figure 14.** Sectors — time series of the nominal target (blue), the feasible target (orange, reflecting implementability constraints), and the realized forward-TE (green), all annualized. Realized TE tracks the feasible bound more closely than the nominal target and tends to undershoot during calm periods, with visible excursions in stress episodes (e.g., the spike around 2020).

*Interpretation.* The series underscores the control logic: Feasibility narrows the effective target, and the realized TE oscillates around the feasible envelope rather than the nominal line. Transient spikes during market stress reflect binding constraints and risk reduction responses. Outside these intervals,

the realized TE remains below the nominal setting, which is consistent with conservative execution.

## Appendix B. Sleeve allocation dynamics (Sectors)

**Sleeve allocation dynamics and interpretation.** Figure 15 presents the time evolution of Bonds, Real, and Cash sleeves for the Sector universe between 2011 and 2024. Two consistent patterns are observed. First, the Bonds and Real sleeves move in opposite directions; when one expands, the other contracts. This counterbalancing behavior reflects a risk-budgeted controller reallocating risk between sleeves while maintaining total active risk within a feasible TE band rather than targeting fixed capital weights. Second, the cash operates as a shock absorber. It remains modest during stable market conditions but increases sharply during stress periods, such as the 2011–2012 turbulence, late 2018 selloff, 2020 pandemic shock, and 2022 interest rate and volatility regime, before reverting toward baseline levels as the conditions stabilize. The gradual adjustments and bounded variation of sleeves are consistent with turnover discipline and implementability constraints.

The allocation behavior in Figure 15 aligns with the TE targeting diagnostics, the realized TE tend to track the feasible rather than nominal target, and the sleeve mechanism acts as a valve that opens or closes to maintain feasibility. In the risk-off phases, the controller increases bonds, reduces real exposure, and temporarily raises cash, lowering portfolio beta and correlation risk while remaining within the TE band. During recovery, cash is redeployed, real exposure increases as volatility subsides, and the available capacity expands. The persistent counterbalancing between Bonds and Real, combined with the limited but strategic use of Cash, confirms that the GATE-WPCA-PI allocator behaves as a geometry-aware, constraint-respecting system, defensive when risk tightens and more expansive when feasible risk increases, resulting in stable, implementable sleeve weights over time.



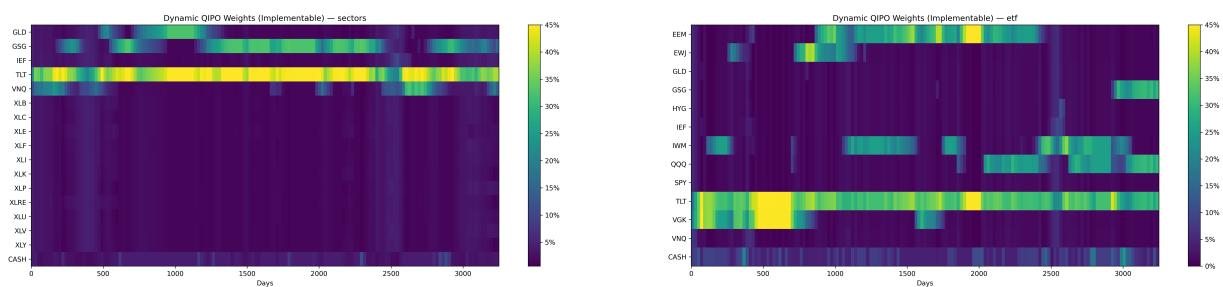
**Figure 15.** Sector universe: Stacked-area allocation showing the evolution of sleeve weights for Bonds (blue), Real assets (orange), and Cash (green) from 2011 to 2024. The figure illustrates the dynamic risk allocation of the GATE-WPCA-PI model under TE and turnover constraints.

**Weights behavior and portfolio allocation dynamics.** Figure 16 compares the time series of the implementable portfolio weights generated by the GATE-WPCA-PI model in the Sector (left, Figure 16a) and the ETF (right, Figure 16b) universes. The results highlight distinct allocation dynamics that reflect both the degree of cross-sectional heterogeneity and the geometry-aware optimization process of the model.

In the Sector universe, the model exhibits pronounced persistence in allocations to duration- and inflation-linked exposures, notably TLT, alongside the modest use of GLD and GSG. The individual equity-Sector sleeves (XLB–XLY) display low-amplitude adjustments, indicating that most of the portfolio's active risk is channelled through rate and real-asset sleeves, rather than through intra-Sector rotation. This allocation structure reflects the relatively homogeneous nature of sectoral returns, where the incremental benefit of geometry is limited, and equal weight remains a strong baseline. This pattern also demonstrates the model's preference for stability and low turnover, which is consistent with the implementability diagnostics reported earlier.

By contrast, the ETF universe displays a markedly more dynamic and diversified allocation structure. The model rebalances capital across global equities (SPY, QQQ, IWM, EEM), credit (HYG), and rates (TLT, IEF) in distinct phases, corresponding to shifts in market leadership between growth-oriented and duration-sensitive assets. The alternating intensity of equity and bond exposures is consistent with the geometry-aware controller's capacity to identify changing regimes of covariance dominance and to reallocate the risk budget across multiple independent factors. Cash holdings, while generally modest, exhibit brief increases during systemic stress episodes (e.g., the 2020 pandemic shock), consistent with the behavior observed in the TE control analysis.

Overall, Figure 16 illustrates that GATE-WPCA-PI adapts its portfolio composition conditionally to the structure of the investable universe. Heterogeneous markets such as ETFs achieve broad diversification and cross-factor rotation. In more homogeneous settings, such as Sectors, exposure is concentrated within duration and inflation hedges, while preserving implementability through smooth, band-limited weight adjustments. These findings confirm that the model's geometry-aware design enhances portfolio responsiveness, while maintaining economic plausibility and turnover control, which are key attributes for robust portfolio optimization.

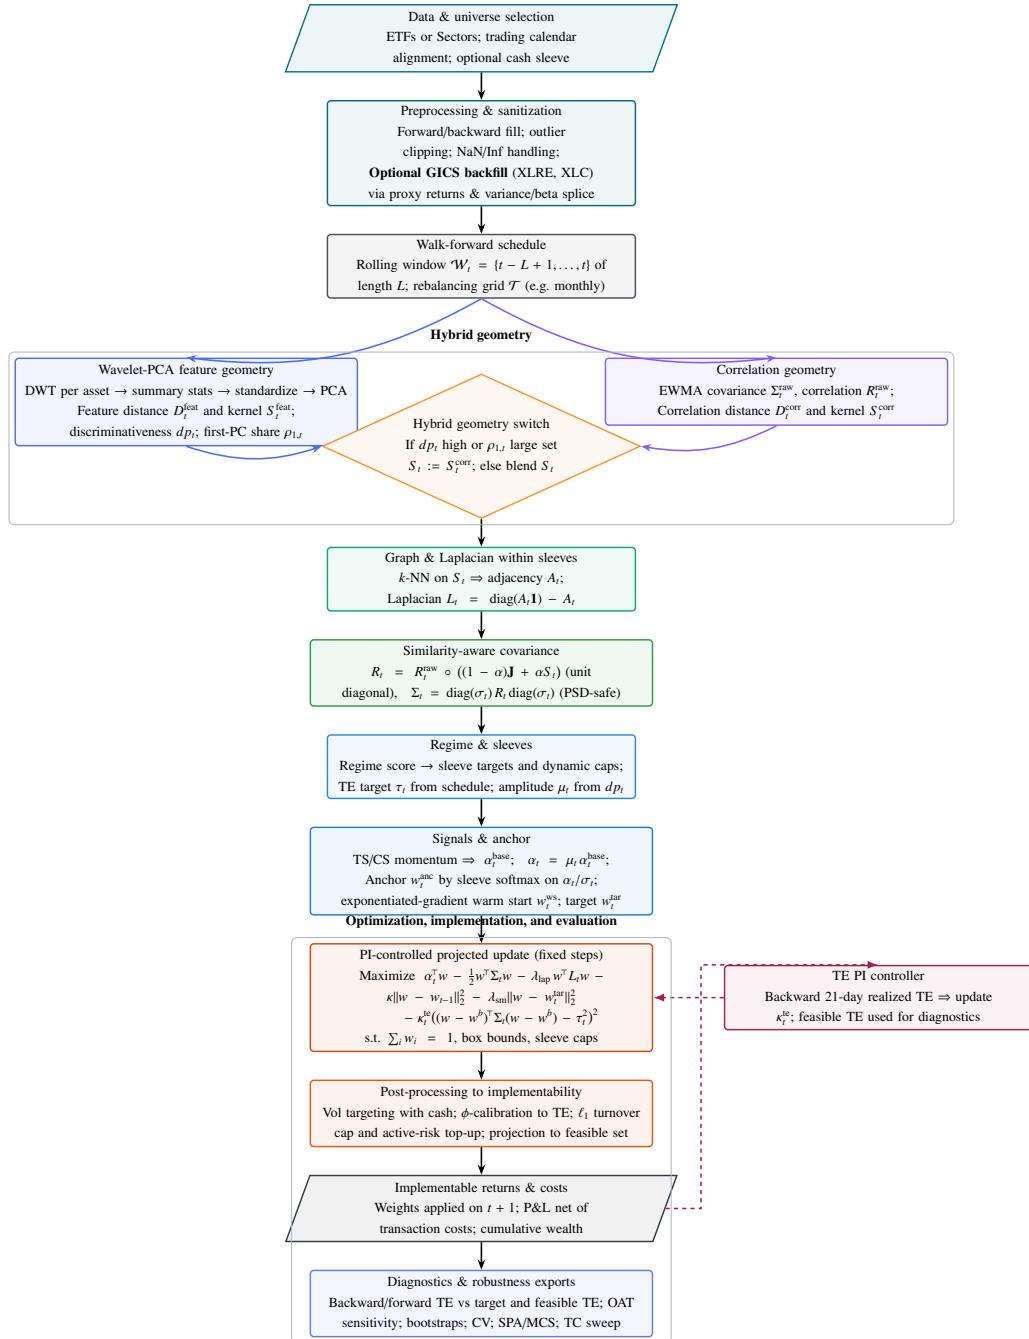


(a) Sectors: implementable weights heatmap.

(b) ETFs: implementable weights heatmap.

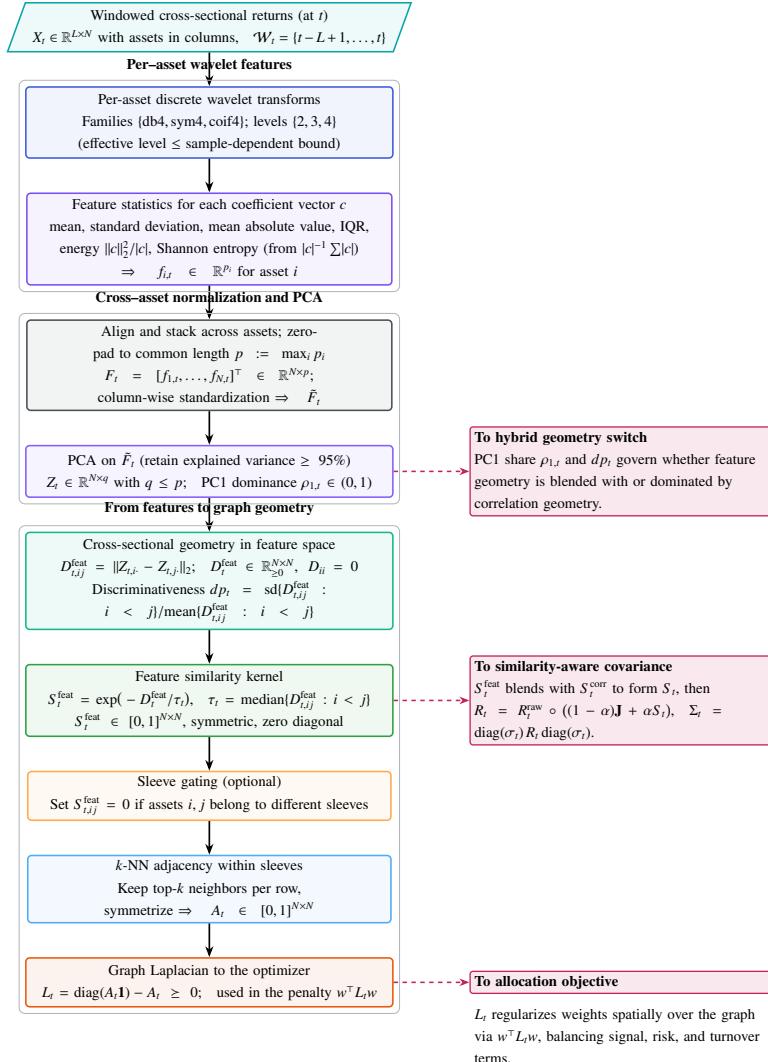
**Figure 16.** Implementable portfolio weights for the GATE-WPCA-PI model. Each panel illustrates the time evolution of the weights as a percentage of the total capital (color intensity) with rows corresponding to the instruments. The Sector universe includes GLD, GSG, IEF, TLT, VNQ, and GICS Sector sleeves, while the ETF universe spans global equity, fixed income, credit, and commodity exposures.

## Appendix C. Flow chart of the GATE–WPCA–PI dynamic allocation



**Figure 17.** Flow chart of the GATE–WPCA–PI dynamic allocation. From sanitized data (with optional GICS backfill) the method builds two geometries (correlation and wavelet-PCA), applies a hybrid switch to obtain  $S_t$ , constructs a  $k$ -NN Laplacian  $L_t$  and a similarity-aware covariance  $\Sigma_t$ , forms signals and an anchor, and performs a TE-controlled projected update under investment constraints. Post-processing enforces volatility targeting, TE calibration, and turnover limits, producing implementable returns and a full robustness suite.

## Appendix D. Feature-to-Graph Flow: Wavelet-PCA to Similarity Kernel and Laplacian (GATE-WPCA-PI)



**Figure 18.** Color flow-map from a windowed returns panel to a graph geometry. For each rebalancing date  $t$ , per-asset wavelet summaries are assembled into a standardized feature matrix  $F_t \in \mathbb{R}^{N \times p}$ , where  $p$  is the total number of aligned wavelet features per asset after zero-padding to a common length. PCA produces an embedding  $Z_t \in \mathbb{R}^{N \times q}$  with  $q \leq p$  principle components retained to explain at least 95% of the variance. Euclidean distances in this feature space give  $D_t^{\text{feat}}$ ; a median-bandwidth radial kernel yields  $S_t^{\text{feat}}$ . Optional sleeve gating and a within-sleeve  $k$ -NN rule produce a sparse adjacency  $A_t$  and its Laplacian  $L_t$ , which regularizes portfolio weights. The PC1 share  $\rho_{1,t}$  and discriminativeness  $d_{p_t}$  feed the hybrid geometry switch, while  $S_t^{\text{feat}}$  blends with correlation geometry in the similarity-aware covariance  $\Sigma_t$ . Here  $|c|$  denotes the length (cardinality) of a coefficient vector  $c$ .

## Notation

**Table 27.** Core sets, indices, data, and weights. Scalars are in Roman italics, vectors bold lowercase, and matrices bold uppercase. Units are stated when applicable.

Symbol	Type	Domain/Dim.	Meaning	Units
$t$	index	$t \in \mathcal{T}$	Rebalancing date index	n/a
$\mathcal{T}$	set	$\mathcal{T} \subset \mathbb{N}$	Set of rebalancing dates	n/a
$\mathcal{W}_t$	set	$\{t - L + 1, \dots, t\}$	Trailing window at $t$ of length $L$ trading days	days
$L$	scalar	$\mathbb{N}$	Window length used for estimation	days
$N$	scalar	$\mathbb{N}$	Number of assets in the universe	n/a
$i$	index	$\{1, \dots, N\}$	Asset index	n/a
$g$	index	$g \in \mathcal{G}$	Sleeve/group index	n/a
$\mathcal{G}$	set	partition of $\{1, \dots, N\}$	Sleeve (group) partition of assets	n/a
$\mathbf{r}_\tau$	vector	$\mathbb{R}^N$	Arithmetic return vector on day $\tau$	daily
$r_{\tau,i}$	scalar	$\mathbb{R}$	Return of asset $i$ on day $\tau$	daily
$\mathbf{X}_t$	matrix	$\mathbb{R}^{L \times N}$	Windowed return panel $[\mathbf{r}_\tau]_{\tau \in \mathcal{W}_t}$	daily
$\mathbf{w}_t$	vector	$\mathbb{R}^N$	Portfolio weights decided at rebalancing $t$	fraction
$w_{i,t}$	scalar	$[0, 1]$	Weight of asset $i$ at $t$	fraction
$\mathbf{w}^b$	vector	$\mathbb{R}^N$	Benchmark weights (equal-weight risky sleeve in experiments)	fraction
$l, u$	scalars	$0 \leq l < u \leq 1$	Element-wise lower/upper weight bounds	fraction
$c_g$	scalar	$[0, 1]$	Sleeve/group cap for $g \in \mathcal{G}$	fraction
$\Omega$	set	$\Omega \subset \mathbb{R}^N$	Feasible set $\{\mathbf{w} : \mathbf{w}^\top \mathbf{1} = 1, l \leq \mathbf{w} \leq u, \sum_{i \in g} w_i \leq c_g\}$	n/a
$\mathbf{1}$	vector	$\mathbb{R}^N$	All-ones vector	n/a

**Table 28.** Covariance, correlations, distances, similarity, graph geometry.

Symbol	Type	Domain/Dim.	Meaning	Units
$\Sigma_t^{\text{raw}}$	matrix	$\mathbb{R}^{N \times N}$	Ledoit–Wolf shrinkage covariance from $\mathbf{X}_t$	daily var
$\sigma_t$	vector	$\mathbb{R}_{\geq 0}^N$	Marginal volatilities $\sqrt{\text{diag}(\Sigma_t^{\text{raw}})}$	daily stdev
$\mathbf{D}_t$	matrix	$\mathbb{R}^{N \times N}$	Diagonal variance matrix $\text{Diag}(\text{diag}(\Sigma_t^{\text{raw}}))$	daily var
$\mathbf{R}_t^{\text{raw}}$	matrix	$\mathbb{R}^{N \times N}$	Raw correlation $\mathbf{D}_t^{-1/2} \Sigma_t^{\text{raw}} \mathbf{D}_t^{-1/2}$	dimensionless
$\mathbf{D}_t^{\text{corr}}$	matrix	$\mathbb{R}_{\geq 0}^{N \times N}$	Correlation distance, $d_{ij} = \sqrt{2(1 - R_{i,j}^{\text{raw}})}$	dimensionless
$\tau_t^{\text{corr}}$	scalar	$> 0$	Kernel bandwidth (median of off-diagonal entries of $\mathbf{D}_t^{\text{corr}}$ )	dimensionless
$\mathbf{S}_t^{\text{corr}}$	matrix	$\mathbb{R}_{\geq 0}^{N \times N}$	Correlation kernel $\exp(-\mathbf{D}_t^{\text{corr}}/\tau_t^{\text{corr}})$ (zeroed across sleeves)	dimensionless
$f_{i,t}$	vector	$\mathbb{R}^{p_t}$	Wavelet statistic features for asset $i$ over $\mathcal{W}_t$	dimensionless
$\mathbf{F}_t$	matrix	$\mathbb{R}^{N \times p_t}$	Stacked feature matrix (standardized)	dimensionless
$\mathbf{Z}_t$	matrix	$\mathbb{R}^{N \times q_t}$	PCA-reduced feature matrix (retain fixed variance fraction)	dimensionless
$\rho_{1,t}$	scalar	$(0, 1)$	PC1 dominance (share of explained variance)	dimensionless
$\mathbf{D}_t^{\text{feat}}$	matrix	$\mathbb{R}_{\geq 0}^{N \times N}$	Euclidean feature distance in $\mathbf{Z}_t$	dimensionless
$\tau_t^{\text{feat}}$	scalar	$> 0$	Feature kernel bandwidth (median off-diagonal of $\mathbf{D}_t^{\text{feat}}$ )	dimensionless
$\mathbf{S}_t^{\text{feat}}$	matrix	$\mathbb{R}_{\geq 0}^{N \times N}$	Feature kernel $\exp(-\mathbf{D}_t^{\text{feat}}/\tau_t^{\text{feat}})$	dimensionless

(continued on next page)

(continued)

Symbol	Type	Domain/Dim.	Meaning	Units
$\psi(\mathbf{D})$	function	$\mathbb{R}^{N \times N} \rightarrow \mathbb{R}_{\geq 0}$	Discriminativeness: $\text{sd}\{D_{ij} : i < j\}/\text{mean}\{D_{ij} : i < j\}$	dimensionless
$d_t^{\text{norm}}$	scalar	$[0, 1]$	Min–max normalized $\psi(\mathbf{D}_t^{\text{feat}})$ across $t$ (or calibration window)	dimensionless
$\omega_t^{\text{feat}}, \omega_t^{\text{corr}}$	scalars	$[0, 1]$	Blend weights: $\omega_t^{\text{feat}} = \psi(\mathbf{D}_t^{\text{feat}})/(\psi(\mathbf{D}_t^{\text{feat}}) + \psi(\mathbf{D}_t^{\text{corr}}))$	dimensionless
$q_\ell$	scalar	$(0, 1)$	Discriminativeness quantile threshold for hybrid switch	dimensionless
$\theta_h$	scalar	$(0, 1)$	PC1 dominance threshold for hybrid switch	dimensionless
$S_t$	matrix	$\mathbb{R}_{\geq 0}^{N \times N}$	Final similarity kernel (hybrid of $S_t^{\text{corr}}$ and $S_t^{\text{feat}}$ )	dimensionless
$k$	scalar	$\mathbb{N}$	$k$ -NN degree (per sleeve) for graph construction	n/a
$\mathbf{A}_t$	matrix	$\mathbb{R}_{\geq 0}^{N \times N}$	Symmetric $k$ -NN adjacency (within sleeves)	dimensionless
$\mathbf{D}_t^{\text{deg}}$	matrix	$\mathbb{R}^{N \times N}$	Degree matrix $\text{Diag}(\mathbf{A}_t \mathbf{1})$	dimensionless
$\mathbf{L}_t$	matrix	$\mathbb{R}_{\geq 0}^{N \times N}$	Graph Laplacian $\mathbf{D}_t^{\text{deg}} - \mathbf{A}_t$ (PSD)	dimensionless
$\alpha$	scalar	$[0, 1]$	Similarity weight in correlation rescaling	dimensionless
$\mathbf{J}$	matrix	$\mathbb{R}^{N \times N}$	All-ones matrix $\mathbf{1}\mathbf{1}^T$	dimensionless
$\circ$	operator	n/a	Element-wise (Hadamard) product between matrices	n/a
$\mathbf{R}_t$	matrix	$\mathbb{R}^{N \times N}$	Similarity-aware correlation $\mathbf{R}_t^{\text{raw}} \circ ((1 - \alpha)\mathbf{J} + \alpha S_t)$	dimensionless
$\Sigma_t$	matrix	$\mathbb{R}_{\geq 0}^{N \times N}$	Adjusted covariance $\text{Diag}(\boldsymbol{\sigma}_t) \mathbf{R}_t \text{Diag}(\boldsymbol{\sigma}_t)$	daily var

**Table 29.** Alpha signals, anchors, warm start, objectives, penalties.

Symbol	Type	Domain/Dim.	Meaning	Units
$h_j$	scalar	$\mathbb{N}$	Time-series momentum horizon $j$ (days)	days
$H$	scalar	$\mathbb{N}$	Cross-sectional momentum horizon (days)	days
$\beta_j$	scalar	$\mathbb{R}$	Weight for TS horizon $h_j$	dimensionless
$\beta_{\text{TS}}, \beta_{\text{CS}}$	scalars	$\mathbb{R}$	Combination weights for TS and CS signals	dimensionless
$\text{Var}_{\text{EWMA}}(\cdot)$	operator	n/a	EWMA variance operator (per asset)	daily var
$s_t^{\text{TS}}$	vector	$\mathbb{R}^N$	Time-series momentum signal (vol-scaled)	dimensionless
$s_t^{\text{CS}}$	vector	$\mathbb{R}^N$	Cross-sectional momentum (demeaned, vol-scaled)	dimensionless
$\alpha_t^{\text{base}}$	vector	$\mathbb{R}^N$	Base alpha $\beta_{\text{TS}} s_t^{\text{TS}} + \beta_{\text{CS}} s_t^{\text{CS}}$	dimensionless
$\mu_{\min}, \mu_{\max}$	scalars	$\mathbb{R}$	Bounds for alpha amplitude schedule	dimensionless
$\mu_t$	scalar	$\mathbb{R}$	Alpha amplitude $\mu_{\min} + (\mu_{\max} - \mu_{\min}) d_t^{\text{norm}}$	dimensionless
$\alpha_t$	vector	$\mathbb{R}^N$	Final alpha vector $\mu_t \alpha_t^{\text{base}}$	dimensionless
$u_{g,t}$	scalar	$[0, 1]$	Sleeve target for group $g$ (sums to 1 over $g$ )	fraction
$\tau$	scalar	$> 0$	Softmax temperature used inside sleeves (anchor construction)	dimensionless
$k_g$	scalar	$\mathbb{N}$	Number of top assets per sleeve used in anchor	n/a
$\mathbf{w}_t^{\text{anc}}$	vector	$\mathbb{R}^N$	Sleeve-aware anchor portfolio	fraction
$\eta$	scalar	$> 0$	Step size for exponentiated-gradient warm start	dimensionless
$\tilde{\mathbf{w}}_t$	vector	$\mathbb{R}^N$	Warm-start weights $\propto \mathbf{w}_{t-1} \circ \exp(\eta(\alpha_t - \Sigma_t \mathbf{w}_{t-1}))$	fraction
$\mathbf{w}_t^{\text{tar}}$	vector	$\mathbb{R}^N$	Target for refinement: $\frac{1}{2} \tilde{\mathbf{w}}_t + \frac{1}{2} \mathbf{w}_t^{\text{anc}}$	fraction
$J_t(\mathbf{w})$	function	$\mathbb{R}^N \rightarrow \mathbb{R}$	Penalized objective (mean-variance, Laplacian, inertia, target shrink, vol band, TE penalty)	dimensionless
$\lambda_{\text{lap}}$	scalar	$\geq 0$	Laplacian penalty weight	dimensionless
$\kappa$	scalar	$\geq 0$	Inertia (smoothing) penalty toward $\mathbf{w}_{t-1}$	dimensionless

(continued on next page)

(continued)

Symbol	Type	Domain/Dim.	Meaning	Units
$\lambda_{\text{sm}}$	scalar	$\geq 0$	Target-shrink penalty toward $\mathbf{w}_t^{\text{tar}}$	dimensionless
$\rho_{\text{vol}}$	scalar	$\geq 0$	Volatility-band penalty weight	dimensionless
$\sigma_{\text{min}}, \sigma_{\text{max}}$	scalars	$> 0$	Lower/upper bounds of target vol band	daily stdev
$\sigma^*$	scalar	$> 0$	Desired operating volatility inside the band	daily stdev
$\sigma(\mathbf{w}; \Sigma_t)$	function	$\mathbb{R}^N \rightarrow \mathbb{R}_{\geq 0}$	Portfolio volatility $\sqrt{\mathbf{w}^T \Sigma_t \mathbf{w}}$	daily stdev
$(x)_+$	function	$\mathbb{R} \rightarrow \mathbb{R}_{\geq 0}$	Positive part $\max\{x, 0\}$	n/a
$\mathbf{d}$	vector	$\mathbb{R}^N$	Active displacement $\mathbf{w} - \mathbf{w}^b$	fraction
$\text{TE}_t^{\text{ann}}$	scalar	$> 0$	Annual TE target	annual stdev
$\Delta_{\text{TE}}$	scalar	$\geq 0$	Amplitude of annual TE schedule	annual stdev
$\tau_t$	scalar	$> 0$	Daily TE target $\text{TE}_t^{\text{ann}} / \sqrt{252}$	daily stdev
$\widehat{\tau}_t^{\text{real}}$	scalar	$> 0$	Backward realized daily TE (window length $h$ )	daily stdev
$h$	scalar	$\mathbb{N}$	Window length for realized TE in PI controller	days
$\kappa_t^{\text{te}}$	scalar	$\geq 0$	TE penalty weight (adaptive via PI)	dimensionless
$K_p, K_i$	scalars	$\geq 0$	Proportional and integral gains in PI update	dimensionless
$I_t$	scalar	$\mathbb{R}$	PI integral state at $t$ (clipped)	dimensionless
$\underline{I}, \bar{I}$	scalars	$\mathbb{R}$	Lower/upper bounds for integral state	dimensionless
$\underline{\kappa}, \bar{\kappa}$	scalars	$\geq 0$	Bounds for $\kappa_t^{\text{te}}$	dimensionless
$s$	scalar	$> 0$	Risky-sleeve scaling for volatility targeting	dimensionless
$\phi$	scalar	$> 0$	TE scaling along active ray (with projection)	dimensionless
$\tilde{\tau}_t$	scalar	$> 0$	Feasible TE target used when $\tau_t$ is unattainable	daily stdev
$\Gamma$	scalar	$\geq 0$	One-shot $\ell_1$ turnover cap at rebalancing	fraction
$\varepsilon$	scalar	$\geq 0$	Remaining turnover budget used for active-risk top-up	fraction
$\Pi_{\Omega}(\cdot)$	operator	$\mathbb{R}^N \rightarrow \Omega$	Euclidean projection onto $\Omega$ (budget, box, sleeve caps)	n/a

**Table 30.** Execution, costs, and returns.

Symbol	Type	Domain/Dim.	Meaning	Units
$c$	scalar	$\geq 0$	Proportional transaction cost per unit $\ell_1$ turnover	fraction
$R_{t+1}$	scalar	$\mathbb{R}$	Implementable portfolio return on $t+1$	daily
$\mathbf{1}\{t \in \mathcal{T}\}$	indicator	$\{0, 1\}$	Indicator of rebalancing day for charging costs	n/a
$M$	scalar	$\mathbb{N}$	Number of projected gradient refinement steps	n/a
$\gamma_m$	scalar	$> 0$	Step size at refinement step $m$	dimensionless

**Units and annualization.** Daily volatility and TE are converted to annualized units by multiplication by  $\sqrt{252}$ . We annualize the mean daily returns by multiplying them by 252. Kernel bandwidths  $\tau_t^{\text{corr}}$  and  $\tau_t^{\text{feat}}$  are dimensionless median distances. All blend weights, discriminativeness measures, and similarities were dimensionless.

**Notes on symbols used in two contexts.** The symbol  $\tau$  appears in two distinct roles: As a softmax temperature (scalar, dimensionless) in the anchor construction, and as a subscripted daily TE target  $\tau_t$  (scalar, daily stdev). These are unrelated, and the context and subscripting distinguish them. The symbol  $h$  denotes the window length used only for the realized TE inside the PI controller, and  $L$  is

the main estimation window length for the features, covariance, and geometry.

**Table 31.** Ticker abbreviations and underlying instruments used in the empirical evaluation. All tickers correspond to liquid, USD-denominated ETFs listed and traded on major U.S. exchanges.

Ticker	ETF (long name)	Intended exposure
<i>Panel A: Multi-asset universe (equities and credit)</i>		
SPY	SPDR S&P 500 ETF Trust	U.S. large-cap equities (S&P 500)
QQQ	Invesco QQQ Trust, Series 1	U.S. large-cap growth (Nasdaq-100)
IWM	iShares Russell 2000 ETF	U.S. small-cap equities (Russell 2000)
VGK	Vanguard FTSE Europe ETF	Developed Europe equities
EWJ	iShares MSCI Japan ETF	Japan equities
EEM	iShares MSCI Emerging Markets ETF	Emerging markets equities
HYG	iShares iBoxx \$ High Yield Corporate Bond ETF	U.S. high-yield corporate credit
<i>Panel B: Sector-sleeve universe (U.S. equity sectors)</i>		
XLB	Materials Select Sector SPDR Fund	U.S. materials sector equities
XLE	Energy Select Sector SPDR Fund	U.S. energy sector equities
XLF	Financial Select Sector SPDR Fund	U.S. financials sector equities
XLI	Industrial Select Sector SPDR Fund	U.S. industrials sector equities
XLK	Technology Select Sector SPDR Fund	U.S. information technology sector equities
XLP	Consumer Staples Select Sector SPDR Fund	U.S. consumer staples sector equities
XLU	Utilities Select Sector SPDR Fund	U.S. utilities sector equities
XLV	Health Care Select Sector SPDR Fund	U.S. health care sector equities
XLY	Consumer Discretionary Select Sector SPDR Fund	U.S. consumer discretionary sector equities
XLRE	Real Estate Select Sector SPDR Fund	U.S. real estate sector equities
XLC	Communication Services Select Sector SPDR Fund	U.S. communication services sector equities
<i>Panel C: Instruments used in both universes (duration and real assets)</i>		
TLT	iShares 20+ Year Treasury Bond ETF	Long-duration U.S. Treasuries
IEF	iShares 7–10 Year Treasury Bond ETF	Intermediate-duration U.S. Treasuries
GLD	SPDR Gold Shares	Gold (bullion-linked)
GSG	iShares S&P GSCI Commodity-Indexed Trust	Broad commodities (S&P GSCI)
VNQ	Vanguard Real Estate ETF	U.S. real estate/REIT equities

*Notes:* Panel C instruments enter both investable universes: in the multi-asset set as standalone duration/real-asset allocations and in the sector-sleeve set as hedging overlays.