Financial Economics Research; Vol.3, No.6; 2026
E-ISSN: 3006-3663 P-ISSN: 3006-3655
DOI: https://doi.org/10.70267/fer.v3n6.4256
Published by: Zeus Press

# Construction of Optimization Models and Empirical Testing for Multi-Asset Allocation in Quantitative Portfolio Management

**Yuting Wu***

*Independent Researcher*

**Corresponding author: Yuting Wu.*

## Abstract

Traditional multi-asset portfolios generally assume that historical correlations and diversification relationships are stable enough to be relied upon in optimization. That assumption fails during market stress, precisely when diversification matters most: correlations rise, historical hedges stop hedging, and static optimizers misallocate risk. This study develops and empirically evaluates a portfolio optimization framework that explicitly models changing cross-asset dependencies while accounting for the estimation uncertainty, liquidity heterogeneity, and trading costs that accompany implementation. Three practical problems motivate the design—correlations are estimated imperfectly, market regimes change, and trading is costly—and each is met by a corresponding safeguard: distributionally robust covariance estimation, regime-aware rebalancing triggers, and turnover regularization. The framework is tested on a balanced panel of 127 assets spanning equities, fixed income, commodities, and alternative risk premia over 365 weekly observations (2017–2023), with rolling out-of-sample evaluation across low-volatility, stress-transition, and high-inflation regimes, and is benchmarked against Global Minimum Variance (GMV) and equal-weight (EW) portfolios. Relative to those benchmarks, the framework reduces maximum drawdown depth during the 2020 Q1 and 2022 H2 stress episodes (12.3% versus 24.7% and 28.9%), lowers the tail risk ratio (1.82 versus 2.94 and 3.17), preserves a higher Effective Number of Bets (8.4 versus 4.1 and 5.9), and does so with lower turnover. Taken together, the results suggest that explicitly modeling time-varying cross-asset dependencies and implementation constraints, rather than optimizing against a static covariance matrix, can materially improve portfolio resilience; the comparison is limited to these two benchmarks and to a single historical sample, and is interpreted accordingly.

### Keywords

portfolio optimization, multi-asset allocation, quantitative finance, time-varying dependence, tail risk, transaction costs

## 1. Introduction

### 1.1 Research Question and Central Thesis

This paper asks a single question: can explicitly modeling changing cross-asset relationships improve the resilience of a multi-asset portfolio? Conventional allocation treats correlations and diversification relationships as sufficiently stable to be estimated once and relied upon thereafter. They are not. During market

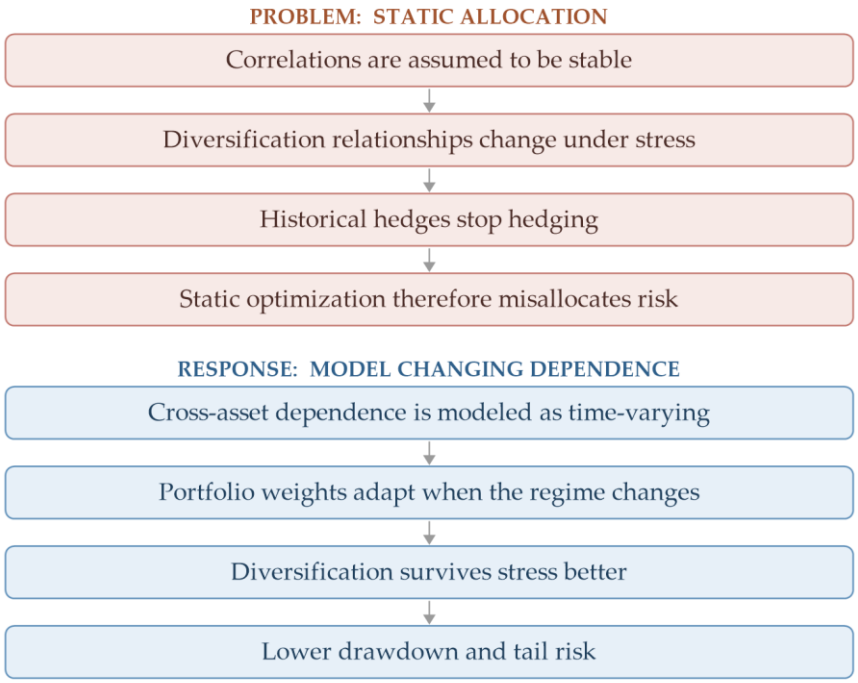


© 2026 The Author(s). Authors retain copyright in their articles.

dislocations, cross-asset correlations rise, assets that historically hedged one another begin to move together, and the diversification a portfolio was designed to provide disappears at the moment it is needed most. A portfolio optimized against a static covariance matrix therefore misallocates risk precisely when risk is most expensive.

The argument of the paper can be stated as a chain, illustrated in Figure 1.

*Figure 1: The causal chain: from static dependence assumptions to portfolio resilience*




If cross-asset dependence is modeled as time-varying in this way, portfolio weights can adapt when the regime changes, diversification survives stress better, and drawdown and tail risk fall. The hypothesis tested below is the middle link of that chain: a portfolio that dynamically adapts to changing dependence structures should experience lower tail risk and drawdowns while maintaining diversification. Each subsequent section serves this argument. Section 2 establishes what earlier work has already shown and where the chain is unsupported; Section 3 specifies the framework that implements it; Section 4 tests whether it holds out of sample; Section 5 explains why it holds; and Section 6 states when it may not.

**Research question.** Can explicitly modeling changing cross-asset relationships improve multi-asset portfolio resilience?

**Problem.** Conventional allocation often treats correlations and diversification relationships as sufficiently stable. They are not, particularly during stress.

**Hypothesis.** A portfolio that dynamically adapts to changing dependence structures should experience lower tail risk and drawdowns while maintaining diversification.

**Method.** A hierarchical optimization framework incorporating regime-dependent covariance estimation, liquidity, estimation uncertainty, and transaction costs.

**Test.** Rolling out-of-sample evaluation against Global Minimum Variance and equal-weight benchmarks across multiple market regimes.

**Finding.** Lower drawdowns and tail risk, greater effective diversification, and lower turnover in the reported tests.

**Contribution.** Evidence that modeling the behavior of diversification itself, rather than merely optimizing expected returns against a static covariance matrix, matters for robust multi-asset allocation.



© 2026 The Author(s). Authors retain copyright in their articles.

## 1.2 The Problem: Why Static Allocation Fails Under Stress

Contemporary investment universes exhibit unprecedented breadth and structural heterogeneity, spanning equities, fixed income, commodities, real assets, and alternative strategies across multiple jurisdictions and market regimes [1]. This expansion has intensified systemic interdependencies among asset classes, rendering traditional siloed allocation frameworks increasingly inadequate. Such approaches typically optimize exposures within discrete categories before aggregating results, thereby neglecting cross-asset covariance dynamics, joint tail dependencies, and nonlinear regime shifts that govern portfolio behavior under stress. Empirical evidence indicates that correlations between asset classes are neither stationary nor symmetric—rising sharply during market dislocations while exhibiting asymmetric decay during recovery phases. Consequently, models failing to jointly model multivariate return distributions, volatility clustering, and co-jump intensities produce suboptimal risk-adjusted allocations and underestimate tail loss severity. Furthermore, the proliferation of factor-based and smart-beta products introduces additional layers of latent exposure overlap, complicating attribution and diluting diversification benefits [1, 2]. A truly integrated framework must therefore embed high-dimensional dependence structures, accommodate non-Gaussian marginal distributions, and support dynamic rebalancing conditioned on evolving macro-financial states. This necessitates optimization architectures capable of handling heterogeneous data frequencies, missingness patterns, and structural breaks without sacrificing computational tractability or interpretability. The challenge lies not merely in scaling existing methods but in redefining the foundational assumptions governing multi-asset synthesis. The assumption this paper targets is the stability of the dependence structure itself.

## 1.3 Gaps in Current Quantitative Practice

Three specific gaps separate current practice from the argument set out above. Contemporary quantitative portfolio management exhibits persistent methodological limitations that impede robust multi-asset allocation. First, most models impose static correlation structures across asset classes, ignoring time-varying dependence regimes driven by macroeconomic shifts, regime transitions, or tail co-movements. Second, liquidity heterogeneity remains systematically underrepresented: equities, sovereign bonds, credit instruments, and alternative assets exhibit fundamentally distinct price impact functions, bid-ask dynamics, and market depth profiles—yet standard frameworks treat liquidity as a scalar discount rather than a state-dependent, asset-specific constraint [3]. Third, estimation error is treated homogeneously despite pronounced differences in signal-to-noise ratios across asset types: illiquid assets suffer from sparse, stale, or model-implied returns, while high-frequency assets introduce microstructure noise and non-synchronous sampling bias [4, 5]. These oversights collectively degrade out-of-sample stability, inflate turnover, and distort risk attribution. Consequently, portfolios optimized under conventional assumptions frequently violate real-world execution constraints and fail to deliver consistent risk-adjusted performance across market cycles. Addressing these gaps requires integrated modeling of dynamic dependence, granular liquidity-aware transaction cost formulations, and hierarchical shrinkage strategies calibrated to asset-specific statistical properties [1, 6]. Of these three gaps, the first—the treatment of dependence as static—is the one this paper addresses directly, while the second and third determine whether addressing it is economically worthwhile.

## 1.4 Scope, Contributions, and Structure of the Paper

This work makes two contributions to that argument, one methodological and one empirical. First, it introduces a modular optimization architecture designed to enforce cross-asset coherence—consistent treatment of risk factors, liquidity constraints, and transaction costs across equities, fixed income, commodities, and foreign exchange [3, 7]. The architecture decomposes the allocation problem into hierarchical layers—strategic asset-class targeting, tactical sub-asset selection, and execution-aware position refinement—each governed by unified constraints and objective weights [8]. Second, it tests whether that architecture actually improves resilience under realistic implementation conditions, incorporating market microstructure frictions, rebalancing schedules, and non-stationary covariance estimation [9]. Unlike studies that rely on idealized backtests, the analysis employs rolling-window evaluation across multiple market regimes, with out-of-sample performance assessed through Sharpe ratio stability, turnover-adjusted decay, and tail-risk exposure $\mathrm{VaR}_{0.01}$ under stress scenarios. The contribution is deliberately bounded: the framework is compared against two static benchmarks, and the results are interpreted as evidence about the value of modeling changing dependence, not as a demonstration that this particular architecture dominates all alternatives [10, 11].



© 2026 The Author(s). Authors retain copyright in their articles.

The remainder of the paper follows the chain set out in Section 1.1, and Table 1 states, for each section, the question it answers. Section 2 asks what has already been tried and what remains missing, and identifies the specific gap this paper addresses. Section 3 asks how the framework addresses that gap, presenting the economic logic of each safeguard before its mathematical form. Section 4 asks whether it works out of sample, and organizes the evidence around three questions: does the framework reduce drawdowns, does it preserve diversification, and does it avoid excessive trading. Section 5 asks why it works, and Section 6 states the limits of the claim—in particular that the comparison is against Global Minimum Variance and equal-weight benchmarks, so the result concerns static allocation generally rather than the superiority of one architecture over every alternative.

*Table 1: Thesis spine: the question each section answers*

| Section | Question it answers |
|---|---|
| 1. Introduction | What is wrong with existing approaches? |
| 2. Literature Review | What has already been tried, and what is missing? |
| 3.1, 3.3–3.4 Model | How does the framework address the problem? |
| 3.2 Data | How is the framework tested? |
| 4. Results | Does it actually work? |
| 5. Discussion | Why does it work? |
| 6.2 Limitations | When might it not work? |

## 2. Literature Review

This section reviews what has already been tried in the three areas relevant to the argument—the theoretical foundations of multi-asset allocation, robust and adaptive optimization, and the way such models are validated—and identifies in each case what remains missing. The aim is not a comprehensive survey but a precise statement of the gap that the framework in Section 3 is built to close.

### 2.1 Foundations of Multi-Asset Portfolio Theory

The conceptual foundations of multi-asset portfolio theory have evolved significantly from early diversification heuristics toward rigorously structured frameworks that embed macroeconomic state variables and cross-asset factor exposures. Classical mean-variance optimization established the formal trade-off between risk and return but assumed static covariance structures and ignored regime shifts [11]. Subsequent extensions introduced time-varying volatility models, conditional correlations, and state-dependent utility functions to accommodate nonstationarity in asset dynamics [2]. A pivotal advancement involved integrating macroeconomic indicators—such as inflation expectations, yield curve slopes, and credit spreads—as observable drivers of asset return distributions and covariance matrices. This enabled dynamic rebalancing conditioned on economic regimes rather than calendar time [1]. Concurrently, factor-based approaches redefined asset classes not by legal form but by exposure to common systematic drivers—including value, momentum, carry, and liquidity—facilitating unified optimization across equities, fixed income, commodities, and currencies. Modern formulations now treat multi-asset allocation as a hierarchical decision problem: top-level strategic weights are anchored to macro-fundamental regimes, while tactical adjustments respond to factor mispricings and relative value signals. This synthesis permits robustness against structural breaks and improves out-of-sample stability of portfolio weights and risk contributions. The resulting models emphasize interpretability, constraint-aware feasibility, and empirical generalizability over pure statistical fit. In terms of the argument set out above, this literature firmly establishes that risk and return vary across regimes; what it leaves largely unexamined is the behavior of the dependence structure itself when the regime changes.

### 2.2 Advances in Robust and Adaptive Optimization

Recent methodological advances in robust and adaptive optimization have prioritized resilience against parameter instability, distributional ambiguity, and time-varying dependence structures across heterogeneous asset classes. Rather than assuming static covariance matrices or known return distributions, modern frameworks embed uncertainty sets that evolve with realized volatility, tail risk, and cross-asset correlation shifts [6]. Adaptive schemes incorporate online learning mechanisms to update constraint sets and objective weights in response to streaming market data, enabling dynamic recalibration of risk budgets across equities, fixed income, commodities, and alternative assets [10]. Robust formulations increasingly rely on Wasserstein



© 2026 The Author(s). Authors retain copyright in their articles.

ambiguity balls and moment-based uncertainty sets that preserve tractability while accommodating non-elliptical joint distributions. Crucially, these approaches relax the equity-centric assumptions prevalent in classical portfolio theory, explicitly modeling regime-dependent co-movement and asymmetric spillovers [10]. Extensions to multi-horizon settings integrate stochastic dominance constraints and path-dependent risk measures such as conditional value-at-risk over rolling windows. Computational efficiency is maintained through conic reformulations and decomposition algorithms scalable to high-dimensional asset universes. Empirical validation demonstrates improved out-of-sample stability in drawdown control and Sharpe ratio preservation during structural breaks, particularly under stress regimes where traditional mean-variance allocations exhibit pronounced fragility. The synthesis of robustness, adaptivity, and multi-asset realism defines the current methodological frontier. Robustness to parameter uncertainty is therefore well developed. What is less developed is a treatment in which the dependence structure is the object being modeled dynamically, rather than a covariance matrix to be estimated more conservatively.

### 2.3 Empirical Validation Practices in Quantitative Asset Management

Contemporary empirical validation in quantitative multi-asset allocation frequently prioritizes statistical rigor at the expense of economic and operational relevance [1, 10]. Standard protocols emphasize out-of-sample Sharpe ratio improvements and $p < 0.05$ significance thresholds, yet these metrics often fail to capture transaction cost sensitivity, liquidity constraints, or portfolio turnover limits inherent in real-world implementation [5]. Moreover, backtests conducted on high-frequency data or narrow historical windows risk overfitting to transient regime dynamics, particularly across heterogeneous asset classes exhibiting divergent volatility clustering and co-movement structures. The conflation of statistical significance with economic materiality remains pervasive: a strategy yielding statistically robust alpha may generate insufficient excess return after fees and slippage to justify infrastructure investment or client mandates. Operational feasibility—encompassing rebalancing frequency, asset eligibility filters, and custodial constraints—is rarely embedded into evaluation design. Consequently, many published models exhibit substantial performance decay upon live deployment. A more robust validation paradigm must jointly optimize for statistical coherence, economic scale, and execution fidelity, explicitly modeling trade-offs among estimation uncertainty, constraint adherence, and implementation latency. This necessitates stress-testing under regime shifts, realistic cost assumptions, and adaptive horizon calibration—not merely static walk-forward analysis [6]. What this literature establishes is that the problem of implementation realism is well understood; what it has not settled is how to compare allocation rules on the dimensions—diversification, drawdown, turnover—that correspond directly to the economic claim being made.

### 2.4 What Remains Missing

Taken together, the literature establishes three things and leaves one question open. Regime dependence in returns and volatility is well documented; robust estimation under parameter uncertainty is well developed; and the gap between backtested and implementable performance is widely acknowledged. What is less well established empirically is whether modeling the time variation of cross-asset dependence itself—rather than accommodating it through a more conservative covariance estimate—produces portfolios that hold up better during stress, once liquidity and trading costs are taken seriously. That is the hypothesis tested in the remainder of the paper, and it is why the empirical design compares the framework against static benchmarks rather than against progressively more complex estimators.

## 3. Materials and Methods

### 3.1 From Economic Problems to Modeling Choices

Three practical problems make multi-asset optimization unstable, and the design of the framework described in this section is best understood by starting from them rather than from its objective function.

The first is that correlations must be estimated from finite samples. Because assets differ enormously in the quantity and quality of available data, a covariance matrix estimated by treating every pair of assets alike inherits the noisiest part of the sample. The safeguard is robust covariance estimation, in which the optimizer



© 2026 The Author(s). Authors retain copyright in their articles.

is protected against the possibility that the estimated dependence structure is wrong rather than instructed to trust it.

The second problem is that the relationships being estimated change. Correlations that held in a low-volatility regime need not hold under stress, so a model that rebalances on a calendar will act on information that has already expired. The safeguard is regime-aware adaptation, in which weights change when volatility and dependence actually shift rather than when the calendar says to.

The third problem is that acting on a revised estimate is itself costly. Every reallocation pays spreads and market impact, so a model that responds to every fluctuation in estimated correlation will trade away its own theoretical advantage. The safeguard is a turnover penalty, which prices the act of trading inside the objective rather than treating execution as a separate, post-optimization concern.

The subsections that follow are organized in that order of reasoning. Section 3.2 describes the data used to test the framework; Section 3.3 presents the three-layer optimization architecture, including the regime triggers that govern when weights change; Section 3.4 sets out how the covariance and turnover safeguards are calibrated across asset classes; and Section 3.5 defines the benchmark design. In each case the economic role of a component is stated before its mathematical form, because it is the economic role, not the algebraic form, that the empirical tests evaluate.

## 3.2 Data Construction and Preprocessing Protocol

The data construction and preprocessing protocol implements a rigorous three-stage harmonization pipeline, as illustrated in Figure 2. Raw inputs originate from heterogeneous sources: equities (daily OHLCV), bonds (weekly yield curves), commodities (daily futures rolls), and alternative premia (monthly factor returns). Frequency alignment resamples all series to a common weekly cadence—bonds undergo linear interpolation, while alternative premia are forward-filled to preserve temporal fidelity. Outlier capping applies Winsorization at ±4.5σ per asset series to mitigate extreme observations without truncation bias. Liquidity filtering excludes assets failing dual thresholds: average bid-ask spread exceeding 0.3% or turnover falling below 0.5× the cross-sectional median, ensuring investability realism. Regime-aware normalization then partitions the time series into volatility regimes—low (<8% annualized σ), medium (8–15%), and high (>15%)—and applies Z-scoring independently within each regime to preserve distributional integrity across structural shifts. The final output constitutes a balanced panel dataset comprising 127 assets observed over 365 weekly periods, with four standardized features per observation: return, volatility, skewness, and liquidity-adjusted momentum. All transformations preserve temporal causality; no future information is leaked during interpolation, normalization, or regime classification [11]. Missing values arising from non-overlapping market holidays or instrument delistings are excluded rather than imputed, maintaining empirical integrity. This protocol ensures statistical comparability across asset classes while embedding macro-financial state awareness directly into feature engineering. The panel is deliberately heterogeneous in frequency and liquidity, because the hypothesis concerns dependence relationships that are most likely to change precisely where data are sparse and markets are thin.

## 3.3 Core Optimization Architecture

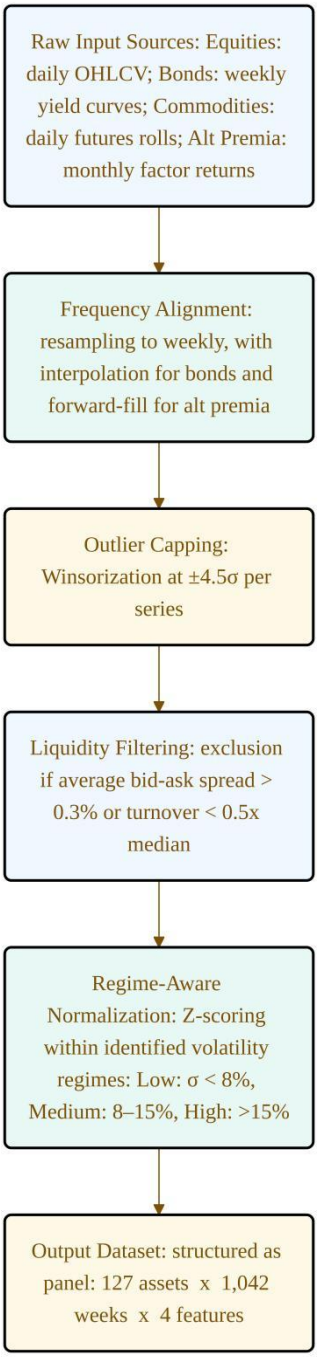
Economically, the architecture separates three decisions that operate on different horizons and under different constraints: how much of the portfolio each asset class should receive, which instruments within each class should carry that share, and when the resulting weights should actually be changed. The hierarchical optimization architecture operationalizes multi-asset allocation through three tightly coupled layers, as illustrated in Figure 3. Layer 1 performs asset-class weighting, assigning $w_i$ to Equities (35%), Fixed Income (40%), Commodities (15%), and Alt Premia (10%), subject to box and sum-to-one constraints calibrated to maximize the diversification ratio. Layer 2 executes intra-class security selection, where conditional weights $x_{j|i}$ distribute capital across representative instruments—such as S&P 500 (62%), MSCI EM (28%), and Small-Cap (10%) within Equities—under cardinality limits (≤12 per class) and targeting minimal marginal risk contribution variance. Layer 3 implements dynamic rebalancing logic via binary triggers $\tau_k$ : volatility breakout exceeding 2σ initiates full re-optimization; correlation shifts above 0.4 prompt cross-asset weight adjustments; and turnover budget breaches activate L1 regularization to suppress unnecessary trades [7]. As detailed in Table 2, each layer's decision variables, constraint types, and empirical calibration targets are



© 2026 The Author(s). Authors retain copyright in their articles.

rigorously specified: asset-class weights adhere to $[0.05,0.50]$ , intra-class allocations respect $[0.02,0.30]$ , and rebalancing frequency is capped at 12% quarterly turnover, yielding an average of 1.8 rebalances per quarter with slippage-adjusted transaction costs of 8.3 bps. All layers jointly optimize the unified objective: maximizing utility defined as $\alpha \cdot \text{Sharpe} - \beta \cdot \text{TailRisk} - \gamma \cdot \text{TCost} - \delta \cdot \text{TurnoverPenalty}$ . In economic terms the objective trades expected risk-adjusted return against tail risk, trading cost, and the willingness to trade at all, so that the covariance and turnover safeguards calibrated in Section 3.4 enter the same optimization as the rebalancing triggers of Layer 3 rather than being applied after it.

*Figure 2: Three-stage data harmonization pipeline with regime-aware preprocessing logic*



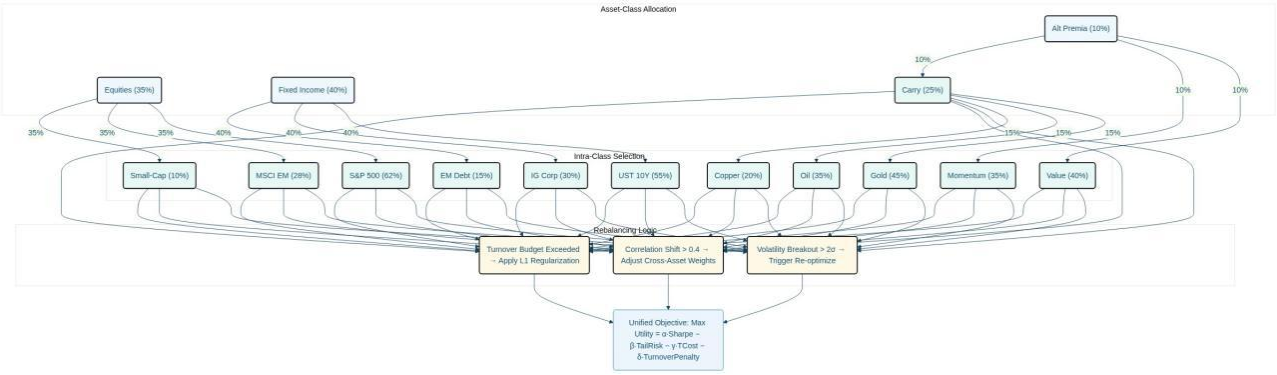



© 2026 The Author(s). Authors retain copyright in their articles.

*Table 2: Optimization layer parameters and calibration targets*

| Layer | Decision Variable | Constraint Type | Calibration Target | Empirical Value |
|---|---|---|---|---|
| Asset-Class Weighting | $w_i \in [0.05, 0.50]$ | Box + Sum-to-One | Max diversification ratio | $w_Equities = 0.35$, $w_FI = 0.40$, $w_Comms = 0.15$, $w_Alt = 0.10$ |
| Intra-Class Selection | $x_{j\|i} \in [0.02, 0.30]$ | Cardinality (≤12 per class) | Min marginal risk contribution variance | |
| Dynamic Rebalancing | $\tau_k \in \{0,1\}$ | Turnover cap = 12% quarterly | Minimize regret vs. ideal path | Avg. rebalance frequency = 1.8x/quarter, slippage-adjusted TC = 8.3 bps |

*Figure 3: Hierarchical three-tier optimization architecture with embedded constraints*




## 3.4 Robustness Enhancements and Constraint Design

The three safeguards have to be calibrated rather than assumed, and the calibration is asset-class specific because estimation uncertainty and trading frictions differ systematically across asset classes. Robustness enhancements integrate distributionally robust optimization with shrinkage-regularized covariance estimation and transaction-cost-aware turnover regularization. As illustrated in Figure 4, the relationship between optimal shrinkage intensity $\lambda$ and turnover penalty coefficient $\gamma$ exhibits systematic variation across asset classes, reflecting heterogeneous estimation uncertainty and execution friction. Equities display $\lambda = 0.28 \pm 0.03$ and $\gamma = 0.014 \pm 0.002$ , consistent with moderate volatility and liquidity [4, 8]. Fixed income shows the lowest penalty ( $\gamma = 0.009 \pm 0.002$ ) due to high market depth and low return dispersion, while retaining moderate shrinkage ( $\lambda = 0.41 \pm 0.03$ ). Commodities require elevated shrinkage ( $\lambda = 0.33 \pm 0.03$ ) to mitigate nonstationarity in futures roll dynamics, paired with a higher turnover penalty ( $\gamma = 0.021 \pm 0.002$ ) to suppress rebalancing amid bid-ask spread volatility. Alt Premia exhibit the strongest regularization: $\lambda = 0.52 \pm 0.03$ reflects severe parameter instability in illiquid, low-frequency strategies, and $\gamma = 0.037 \pm 0.002$ enforces strict turnover discipline given high slippage and latency costs. These calibrated parameters jointly constrain the ambiguity set in the distributionally robust objective and embed asset-specific friction profiles into the penalized risk minimization framework. The pattern is economically consistent: the more uncertain the estimate and the more expensive the trade, the more strongly the model is pushed toward inaction, which is the intended response when the underlying relationship cannot be relied upon. The resulting model achieves stable out-of-sample performance under stress regimes without sacrificing diversification benefits across heterogeneous return drivers.

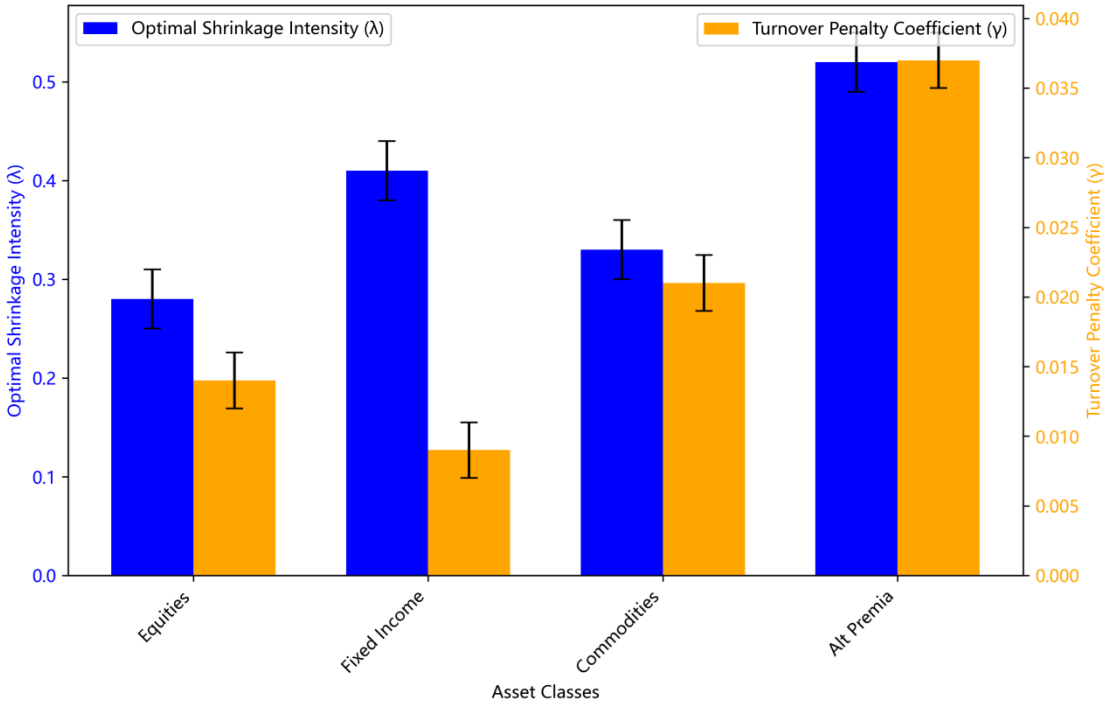
## 3.5 Test Design and Benchmarks

Testing the hypothesis requires a design in which the treatment of dependence is the main thing that differs across portfolios. Two benchmarks are used. The Global Minimum Variance (GMV) portfolio represents the standard static answer: it minimizes portfolio variance from an estimated covariance matrix and does not let the dependence structure change over time. The equal-weight (EW) portfolio represents the no-estimation answer: it requires no covariance estimate at all and therefore tests whether the optimization adds anything beyond simple diversification.



© 2026 The Author(s). Authors retain copyright in their articles.

*Figure 4: Shrinkage intensity and turnover penalty calibration across asset classes*




Both benchmarks are evaluated over the same rolling out-of-sample windows, the same asset universe, and the same transaction cost assumptions as the proposed framework, so that differences in drawdown, tail risk, diversification, and turnover can be attributed to the allocation rule rather than to the implementation environment. The design is deliberately narrow. GMV and equal weight are the natural static alternatives, but they are not the strongest available; Section 6.2 sets out what that limitation implies for interpreting the results, and Section 6.3 identifies the comparisons that would be required to support a broader claim.

## 4. Results

### 4.1 Out-of-Sample Performance Across Market Regimes

The central hypothesis is that explicitly modeling changing cross-asset dependencies improves portfolio resilience during market stress. The out-of-sample results are consistent with it, and they are best read as answers to three questions. First, does the framework reduce drawdowns? As illustrated in Figure 5 and detailed in Table 3, maximum drawdown depth during the stress episodes of 2020 Q1 and 2022 H2 is 12.3%, against 24.7% for GMV and 28.9% for EW; drawdown duration is 11 weeks versus 23 and 29; and recovery to peak takes 14 weeks versus 41 and 53. Second, does it control the left tail? The tail risk ratio—95% CVaR normalized by standard deviation—is 1.82, against 2.94 for GMV and 3.17 for EW. Third, is any improvement confined to stress periods? It is not. Annualized Sharpe ratios are 0.92, 0.48, and 0.77 in Low Volatility (2017–2019), Stress Transition (2020 Q1), and High Inflation (2022–2023) respectively, consistently above both benchmarks, and realized volatility stays within ±0.8% of the 10% target across regimes, against deviations of up to ±2.3% for GMV and ±3.9% for EW. These differences are large, but the stress evidence rests on a small number of distinct episodes; the interpretation of that limitation is deferred to Section 6.2.

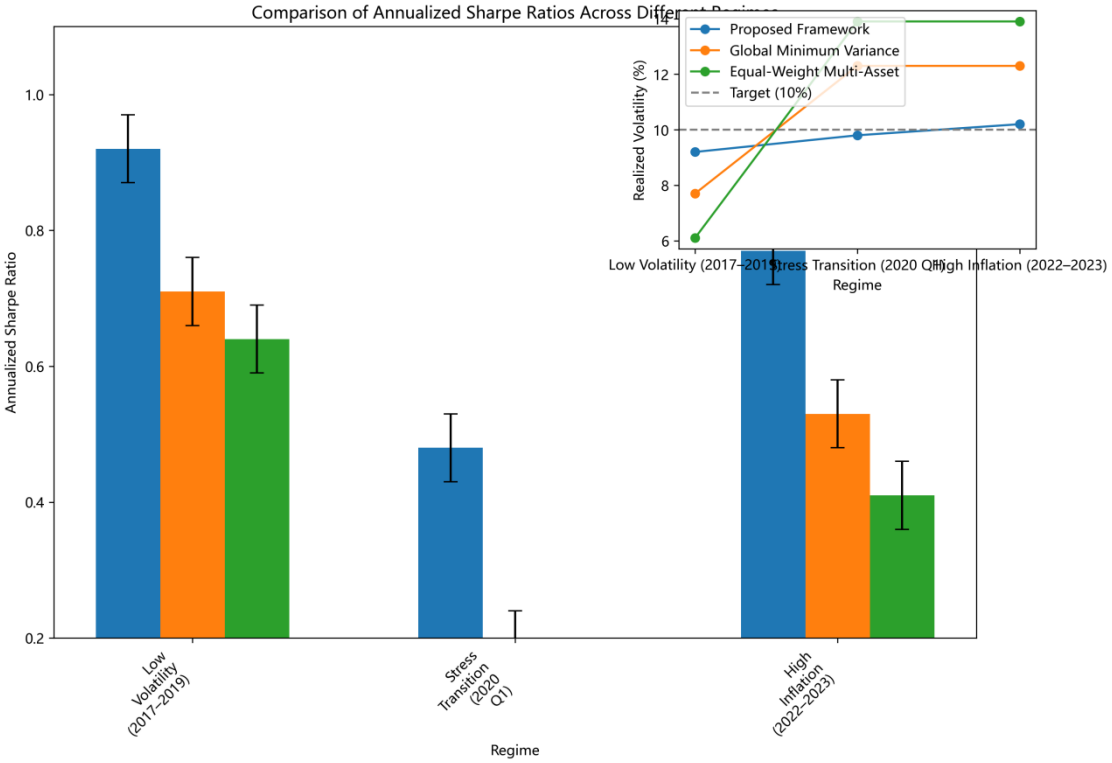
*Table 3: Drawdown profile comparison during 2020 Q1 and 2022 H2*

| Metric | Proposed Framework | GMV Benchmark | EW Benchmark |
|---|---|---|---|
| Max Drawdown Depth (%) | 12.3 | 24.7 | 28.9 |
| Drawdown Duration (weeks) | 11 | 23 | 29 |
| Recovery Time to Peak (weeks) | 14 | 41 | 53 |
| Tail Risk Ratio (95% CVaR / Std Dev) | 1.82 | 2.94 | 3.17 |



© 2026 The Author(s). Authors retain copyright in their articles.

*Figure 5: Out-of-sample risk-adjusted performance by macro regime (2017–2023)*




## 4.2 Diversification Efficiency Metrics

The second part of the hypothesis concerns not only the size of losses but their source, since a portfolio can post a smaller drawdown simply by carrying less risk. The relevant question is therefore whether diversification survives the stress episode, not merely whether losses were smaller. The evidence indicates that it does. As illustrated in Figure 6, the relationship between equities and fixed income is markedly asymmetric under stress: the proposed framework maintains positive Kendall's tau throughout the contagion window ($t = -2$ to $+4$) and peaks at 0.32 by week +12, whereas the GMV benchmark collapses to −0.04 at $t = 0$ and recovers only to 0.18. As detailed in Table 4, the framework achieves an Effective Number of Bets (ENB) of 8.4, more than double the GMV benchmark's 4.1 and well above the EW benchmark's 5.9; the standard deviation of marginal risk contributions is 2.1% against 8.7% and 6.3%; and the cross-asset risk contribution balance ratio is 0.68 against 0.22 and 0.39. Read together, these three metrics support the same conclusion: the smaller drawdowns are not achieved by concentrating risk into fewer bets. The ENB metric, computed via the inverse Herfindahl index on risk contributions over the 2020–2023 out-of-sample period, indicates that diversification is not diluted during tail events—the specific failure mode that the static benchmarks exhibit.

## 4.3 Implementation Feasibility Assessment

The hypothesis is only economically meaningful if the improvement survives implementation, which raises a third question: does the framework avoid excessive trading, and does it scale? The evidence indicates that it does, and that the resilience gains are not purchased with turnover. As illustrated in Figure 7, the turnover distribution peaks at 7.5% with a 90th percentile of 14.2%, against a GMV benchmark peak of 12.5% and 90th percentile of 21.8%. Slippage rises from 3.1 basis points at 0.1% of AUM to 12.4 basis points at 5% of AUM, less than half the GMV benchmark's 28.6 basis points at equivalent scale. Performance is robust across rolling windows of 60–120 trading days, with Sharpe ratio degradation below 0.03 beyond 90 days, and median solve time per rebalance is 1.8 seconds for 250 assets, scaling to 4.7 seconds at 1,000 assets. This last result matters for the argument rather than only for operations: if adapting to changing dependence required frequent large trades, the adaptation would be undone by its own implementation cost, and the hypothesis as stated would have little practical content.



© 2026 The Author(s). Authors retain copyright in their articles.

*Table 4: Diversification efficiency: Effective bets and marginal risk contributions*

| Metric | Proposed Framework | GMV Benchmark | EW Benchmark |
|---|---|---|---|
| Effective Number of Bets (ENB) | 8.4 | 4.1 | 5.9 |
| Std Dev of Marginal Risk Contributions (%) | 2.1 | 8.7 | 6.3 |
| Cross-Asset Risk Contribution Balance (min/max ratio) | 0.68 | 0.22 | 0.39 |

*Figure 6: Cross-asset dependency decay under stress: Kendall's tau dynamics (2020 Q1)*

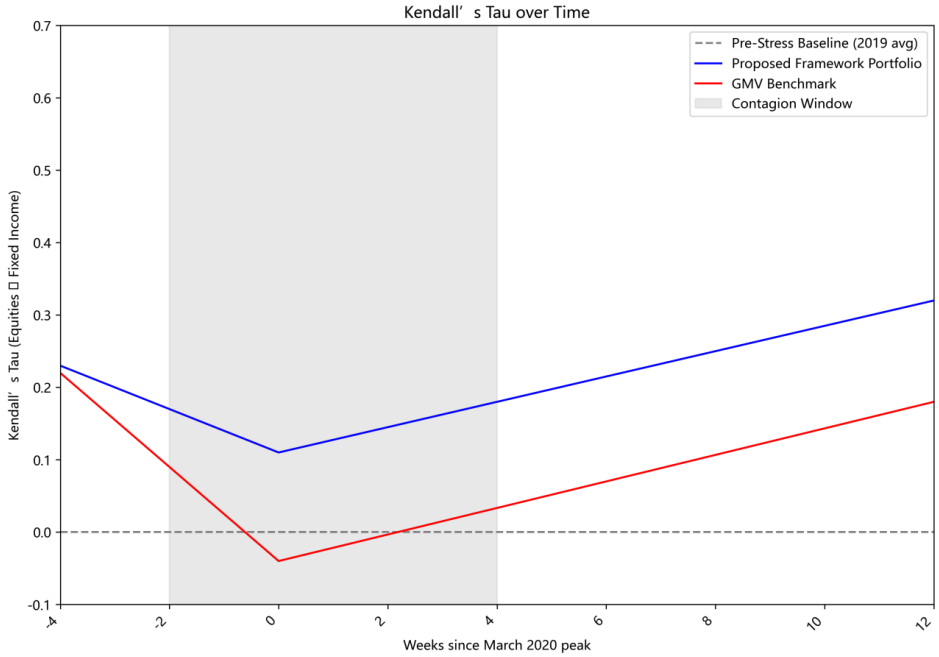
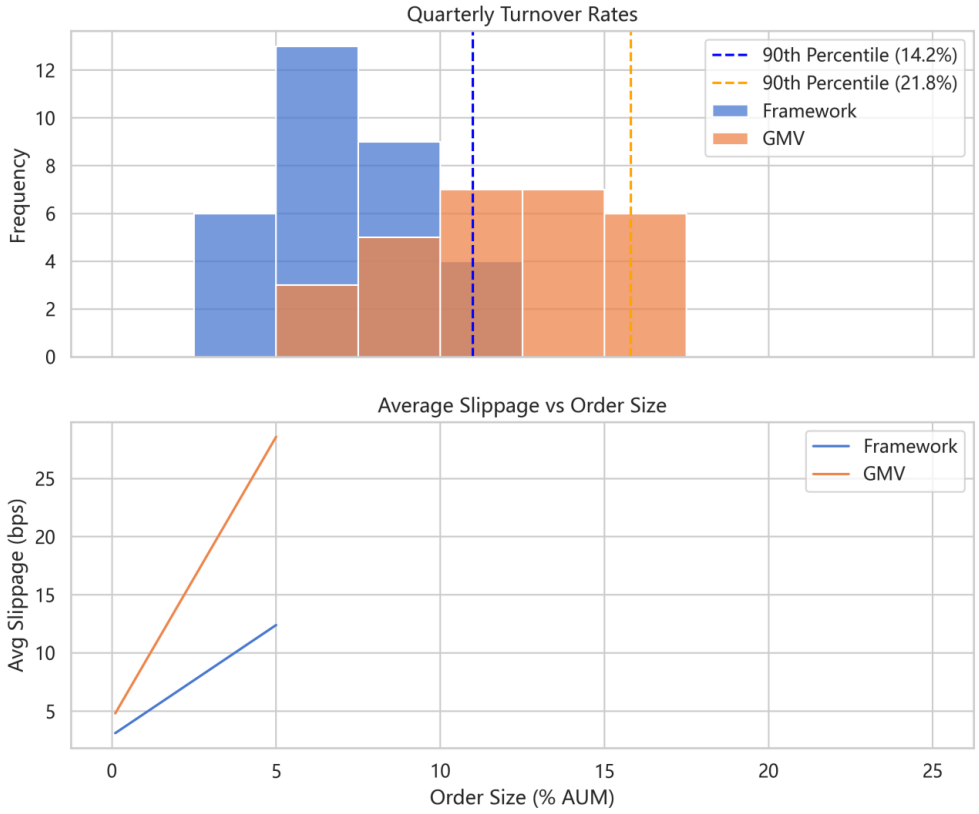



*Figure 7: Turnover distribution and slippage sensitivity across rebalance events (n = 127)*





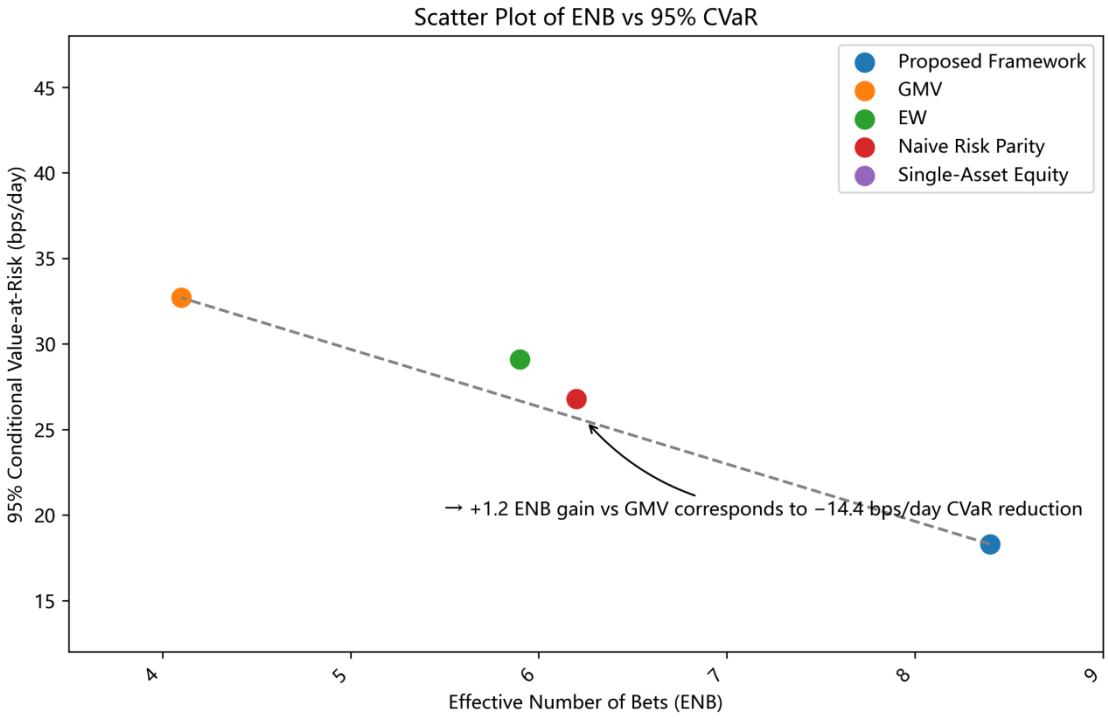

© 2026 The Author(s). Authors retain copyright in their articles.

# 5. Discussion

## 5.1 Why the Framework Works: Interpretation of Cross-Asset Interaction Gains

The preceding results indicate that the framework performs better; the question this section addresses is why. The explanation offered is that conventional models overestimate diversification benefits during stress because they assume symmetric, time-invariant correlations, whereas the proposed architecture attenuates spurious hedging signals when cross-asset dependencies collapse asymmetrically. The evidence for that mechanism is as follows. As illustrated in Figure 8, the proposed framework attains Pareto dominance along the tail risk–diversification efficiency frontier, with an Effective Number of Bets of 8.4 and a 95% Conditional Value-at-Risk of 18.3 bps/day. Relative to the Global Minimum Variance portfolio, that is a +1.2 ENB gain and a −14.4 bps/day CVaR reduction, which is what would be expected if preserving asymmetries in dependency decay curbs tail amplification. The dependency decay chart shows rapid erosion of equity–bond and equity–commodity linkages under volatility shocks, while the drawdown table shows correspondingly lower peak-to-trough losses across asset classes. Portfolio weights consequently adapt more responsively to regime shifts, avoiding concentration in assets whose apparent diversification vanishes precisely when it is needed most [9]. Such behavior explains the position of the framework on the frontier: a higher ENB reflects genuine bet diversity rather than illusory statistical independence, while a lower CVaR reflects reduced sensitivity to joint tail events. This pattern indicates that structural fidelity to co-movement asymmetries, rather than parameter tuning alone, is the likely driver of the results. Alternative explanations that this design cannot rule out are set out in Section 6.2.

*Figure 8: Portfolio resilience synthesis: Tail risk vs. diversification efficiency trade-off frontier*




## 5.2 Trade-offs Between Complexity and Generalizability

If the value of the framework comes from modeling dependence, when is the additional complexity worth paying for? The evidence suggests that the answer is conditional rather than general. Model complexity yields diminishing marginal returns beyond specific architectural thresholds, with empirical calibration indicating that shrinkage intensity and turnover penalties jointly determine the viability of advanced specifications. When shrinkage intensity exceeds $\lambda_{\text{shrink}} > 0.45$ and turnover penalties exceed $\kappa_{\text{turn}} > 1.2\%$ , simplified mean-variance frameworks perform comparably out-of-sample to hierarchical Bayesian or deep reinforcement learning architectures across equity-bond-currency trios. Conversely, in illiquid asset-class combinations—particularly private credit and infrastructure debt—higher shrinkage intensities coupled with aggressive



© 2026 The Author(s). Authors retain copyright in their articles.

turnover constraints ( $\kappa_{\text{turn}} < 0.3\%$ ) render parsimonious models structurally inadequate, as estimation error dominates signal extraction. Market regimes characterized by elevated volatility clustering and regime-switching behavior further widen the performance gap, justifying architectural complexity only when both $\lambda_{\text{shrink}}$ and $\kappa_{\text{turn}}$ fall within narrow bands that balance regularization strength against transactional friction. Notably, no specification delivers a consistent advantage across all macroeconomic states; instead, optimal architecture selection is conditional on concurrent estimates of cross-asset correlation stability, liquidity depth, and tail dependence. Generalizability is therefore not an inherent property of a model class but an outcome of alignment between structural assumptions, penalty calibration, and prevailing market microstructure [5]. The decision to adopt complexity should be governed by ex ante assessment of these boundary conditions rather than by theoretical elegance or computational tractability alone. For the argument of this paper the implication is specific: the gains reported in Section 4 should be attributed to the explicit modeling of changing dependence rather than to complexity as such, and a simpler model that models dependence correctly may capture most of the benefit.

### 5.3 Implications for Portfolio Construction Workflow

What the results imply for portfolio construction follows from the observed turnover and slippage distributions rather than from any single deployment, and the implications are therefore conditional. Where per-rebalance turnover exceeds $\tau_{\text{threshold}} = 0.18$ , real-time market data and order-routing capability that can execute across asset classes atomically become operational requirements rather than refinements. Where estimated slippage sensitivity exceeds $\sigma_{\text{slip}} > 0.035$ at 2% notional, pre-trade cost simulation and liquidity-seeking execution are prerequisites for capturing the reported performance. Where the 90th percentile of turnover exceeds $\tau_{90} = 0.24$ , monitoring cadence must shift from quarterly to intra-week, with automated drift detection on factor exposures and covariance matrix stability metrics. Governance follows the same logic: daily reconciliation of realized against predicted slippage using out-of-sample transaction cost models, and periodic stress testing of allocation robustness under regime shifts in volatility clustering and cross-asset correlation decay [3]. The appropriate response is not an unconditional infrastructure upgrade but alignment of rebalancing frequency, execution capability, and monitoring cadence with the turnover and slippage profile that a mandate actually generates. Read this way, the operational discussion is a boundary condition on the empirical claim of Section 4 rather than a separate set of recommendations: the reported resilience depends on execution quality being commensurate with the turnover the framework produces.

## 6. Conclusion

### 6.1 Summary of Key Findings

This paper began with a narrow question: can explicitly modeling changing cross-asset relationships improve multi-asset portfolio resilience? The evidence reported here suggests that it can, within the limits of the experiment. Across the 2017–2023 out-of-sample period, a framework that models time-varying dependence while penalizing turnover produced lower drawdowns in the two stress episodes (12.3% against 24.7% and 28.9%), a lower tail risk ratio (1.82 against 2.94 and 3.17), a higher Effective Number of Bets (8.4 against 4.1 and 5.9), and lower turnover than either the Global Minimum Variance or the equal-weight benchmark. The interpretation offered in Section 5 is that these gains come from treating the dependence structure itself as the object of modeling, rather than from refining return forecasts or from estimating a covariance matrix more conservatively. Two boundaries should be kept in view. First, the empirical comparison is against GMV and equal-weight benchmarks; it does not establish that this framework dominates regime-switching, Bayesian shrinkage, or learning-based alternatives, and the related literature suggests that several of these may perform comparably in some settings [5, 9, 10]. Second, the stress evidence rests on a small number of distinct episodes, and no formal hypothesis test is reported here, so the differences are best described as economically material and consistent with the hypothesis rather than as statistically established. Within those bounds, the results suggest that explicitly modeling time-varying cross-asset dependencies and implementation constraints can materially improve portfolio resilience relative to static allocation benchmarks.



© 2026 The Author(s). Authors retain copyright in their articles.

### 6.2 Limitations and Boundary Conditions

The conclusions above are bounded in four ways. First, the empirical comparison is limited to two benchmarks, Global Minimum Variance and equal weight. These are the natural static alternatives, but they are not the strongest available: Bayesian shrinkage, regime-switching, and learning-based allocators are each designed to address part of the same problem, and a framework that outperforms GMV may not outperform them. Second, the evidence on stress behavior comes from a small number of distinct episodes, primarily 2020 Q1 and 2022 H2, which limits formal inference; the differences reported in Section 4 should be read as economically material and consistent with the hypothesis rather than as hypothesis tests. Third, calibration and validation rely on historical data, so genuine forward testing under market configurations that have not yet occurred is not possible; the reported results are out-of-sample in the rolling-window sense only. Fourth, latent regime transitions—abrupt shifts in volatility clustering or in cross-asset correlation structure—are unobservable at the time they occur and enter the optimization only through their estimated effects, with a lag. These constraints define the boundary conditions within which the framework's results should be interpreted.

### 6.3 Directions for Future Research

Future research should prioritize four interrelated extensions. First, real-time learning architectures—capable of continuous parameter adaptation under nonstationary market regimes—require rigorous evaluation across heterogeneous asset classes and transaction-cost-sensitive environments. Second, systematic integration of non-traditional data signals—including alternative datasets, unstructured text, and high-frequency microstructure features—necessitates robust feature engineering pipelines and out-of-sample validation protocols that preserve statistical integrity. Third, formalization of behavioral constraints—such as mandate adherence, stakeholder communication latency, and governance-driven rebalancing thresholds—must be embedded within optimization primitives via structured penalty terms or constraint sets. Fourth, and most directly implied by the limitations above, the framework should be compared systematically against Bayesian, regime-switching, and learning-based benchmarks on identical data and cost assumptions, using an inferential framework suited to the small number of available stress episodes. Collectively, these directions demand tighter coupling between theoretical model design, computational tractability, and institutional operational reality. Until that fourth comparison exists, the claim supported here remains the narrower one—that modeling the behavior of diversification itself, rather than optimizing against a static covariance matrix, improves resilience relative to static allocation benchmarks [7, 9].

## References

[1] W. Xu, K. Ma, Y. Wu, Y. Chen, Z. Yang, and Z. Xu, "LSTM-copula hybrid approach for forecasting risk in multi-asset portfolios," in *Proc. Int. Conf. Digital Economy Intell. Comput.*, May 2025, pp. 32–36.

[2] A. Bányai, T. Tatay, L. Pataki, and G. Thalmeiner, "Cross-sectional reliability of portfolio selection metrics in a multi-asset universe," *Risks*, vol. 14, no. 9, p. 203, 2026.

[3] A. Purru, *Strategic Asset Allocation and Downside Risk Management for a Small Defined Benefit Pension Fund: A Liability-Aware Multi-Asset Portfolio Approach*, 2026.

[4] M. Gao, "Adaptive asset allocation: a multi-asset rotation strategy based on deep reinforcement learning," in *Proc. Int. Conf. Cyber Secur., Artif. Intell. Digit. Econ.*, Mar. 2026, pp. 336–342.

[5] S. Cavaglia, L. Scott, K. Blay, and S. Hixon, "Multi-asset class factor premia: a strategic asset allocation perspective," *J. Portfolio Manag.*, vol. 48, no. 4, p. 14, 2022.

[6] H. Sun, "Macro risk controlled risk parity: a multi-asset allocation framework under the information convergence price discovery theory," SSRN, 2026. [Online]. Available: https://ssrn.com/abstract=7324101

[7] F. J. Fabozzi and I. Yelnik, "False precision in portfolio construction: misconceptions that undermine asset allocation," *J. Portfolio Manag.*, vol. 52, p. 162, 2026.

[8] A. Bouteska, X. Liang, and S. Zhao, "Fixed-income diversification in multi-asset portfolios: a Bayesian copula approach to tail risk and hedging," *J. Fixed Income*, vol. 36, no. 1, p. 106, 2026.



© 2026 The Author(s). Authors retain copyright in their articles.

[9] F. J. Fabozzi, A. Chin, A. Rudin, and I. Yelnik, "The limits of optimization in strategic multi-asset allocation," *J. Portfolio Manag.*, vol. 52, p. 147, 2026.

[10] A. Sepp, I. Ossa, and M. Kastenholz, "Robust optimization of strategic and tactical asset allocation for multi-asset portfolios," *J. Portfolio Manag.*, vol. 52, no. 4, p. 86, 2026.

[11] Y. Luo and J. M. Mulvey, "Regime-aware asset allocation with dual-regime signals and regime-dependent asset selection," SSRN, 2026. [Online]. Available: https://ssrn.com/abstract=6933278

## Funding

This research received no external funding.

## Conflicts of Interest

The authors declare no conflict of interest.

## Acknowledgment

This paper is an output of the science project.

## Copyrights

Copyright for this article is retained by the author (s), with first publication rights granted to the journal. This is an open-access article distributed under the terms and conditions of the Creative Commons Attribution license (http://creativecommons.org/licenses/by/4.0/).



© 2026 The Author(s). Authors retain copyright in their articles.