Center for Management Science Research
ISSN: 3006-5291 3006-5283 Volume 4, Issue 6, 2026

PORTFOLIO OPTIMIZATION USING MODERN PORTFOLIO THEORY (MPT) AND MARKOWITZ EFFICIENT FRONTIER MODELING TECHNIQUES

Zahid Latif[*1], Syed Ghazanfer Inam[2], Fazle Adil[3], Talha[4]

*[*1]PhD Scholar, National College of Business Administration & Economics (NCBA&E), Rahim Yar Khan Campus, Pakistan*
[2]Mohammad Ali Jinnah University, Karachi, Pakistan
[3]Local Government & Rural Development Department Government of Khyber Pakhtunkhwa & MSC in International Business Department of Ulster Unitersity Business School at Ulster Unitersity London United Kingdom.
[3]University of Makran, Pakistan

[*1]zahidlatif.ms@gmail.com, [2]ghazanfer@jinnah.edu, [3]fazleadil@gmail.com [3]talha@uomp.edu.pk

DOI: https://doi.org/10.5281/zenodo.21187198

Keywords
Modern Portfolio Theory (MPT); Black-Litterman Model; Downside Risk; Robust Portfolio Optimization; Smart Predict-then-Optimize (SPO); Machine Learning

Article History
Received: 25 April 2026
Accepted: 04 June 2026
Published: 21 June 2026

Corresponding Author: *
Zahid Latif

Abstract

This review paper provides a comprehensive taxonomic synthesis of portfolio optimization frameworks, tracking the evolution of asset allocation from classical quantitative foundations to advanced data-driven models. It begins with Harry Markowitz's 1952 Modern Portfolio Theory (MPT), outlining its pioneering use of variance as a symmetrical risk metric and its vulnerability to estimation errors, which often turn optimizers into "estimation-error maximizers". To address these challenges, the paper analyzes Post-Modern Portfolio Theory (PMPT), which incorporates asymmetrical risk metrics like downside deviation and integrates high-frequency intraday data to capture downside realized semivariance and price jumps. The evaluation then shifts to parameter stabilization techniques, highlighting the Bayesian framework of the Black-Litterman model, which combines market equilibrium priors with subjective views to mitigate input sensitivity. Moving beyond classical boundaries, the paper reviews Robust Portfolio Optimization (RPO) and Mixed Conditional Value-at-Risk (MCVaR) frameworks designed to counter tail risk under non-parametric uncertainty sets. It also examines the computational challenges introduced by real-world market frictions and mixed-integer constraints, assessing how metaheuristic algorithms navigate these non-convex spaces. Finally, the paper highlights contemporary breakthroughs in artificial intelligence, comparing the traditional "predict-then-optimize" paradigm with decision-focused, end-to-end Smart Predict-then-Optimize (SPO) loss functions, alongside deep neural and graph networks. By structuring these methodologies into a comparative matrix, this review highlights the transition from rigid historical estimation toward adaptive, decision-focused, and robust risk management frameworks capable of navigating modern financial markets.

Introduction

Modern portfolio theory (MPT), established by Harry Markowitz in his seminal 1952 paper, introduced a mathematical formulation that transformed asset allocation from a qualitative, security-specific selection process into a quantitative, portfolio-level framework (Markowitz, 1952). Before MPT, investment practices focused primarily on evaluating individual financial instruments in isolation (Elton & Gruber, 1997). MPT established that an asset's risk and return should not be assessed independently, but rather by how its inclusion impacts the joint probability distribution of the entire portfolio. The core assumption of this paradigm is that investors are rational, risk-averse agents who make decisions based purely on expected returns and risk (Rutterford & Sotiropoulos, 2016).

Mathematically, MPT models the returns of a portfolio of N risky assets as a linear combination of individual asset returns. Let $w = [w_1, w_2, \ldots, w_N]^T$ represent the portfolio weight vector, subject to the budget constraint $w^T\mathbf{1} = 1$, where $\mathbf{1}$ is an $N \times 1$ column vector of ones (Elton & Gruber, 1997). If the expected returns of the constituent assets are denoted by the vector μ, the expected return of the portfolio, $E(R_p)$, is formulated as:

$$E(R_p) = \sum_{i=1}^{N} w_i\, E(R_i) = w^T\mu$$

The total risk of the portfolio is quantified by its variance, σ_p^2, which depends on both individual asset variances and the pairwise covariances between assets (Markowitz, 1991). The algebraic representation is:

$$\sigma_p^2 = \sum_{i=1}^{N} w_i^2\, \sigma_i^2 + \sum_{i=1}^{N}\sum_{j\neq i}^{N} w_i\, w_j \sigma_{ij} = w^T\Sigma w$$

where Σ represents the symmetric $N \times N$ covariance matrix, σ_i^2 represents the variance of asset i, and σ_{ij} represents the covariance between asset i and asset j. The reliance on covariance demonstrates the mathematical mechanism of diversification. If assets exhibit low or negative covariance ($\sigma_{ij} \leq 0$), the portfolio variance is reduced below the weighted average of the individual asset variances (Kolari & Pynnönen, 2023). Diversification eliminates non-systematic (idiosyncratic) risk, leaving the portfolio exposed primarily to non-diversifiable systematic market risk (Edlinger & Parent, 2014).

Input Sensitivity, Estimation Error, and the Challenges of Classical Optimization

Despite its theoretical elegance and foundational role in modern finance, classical mean-variance optimization exhibits severe limitations that often degrade its performance in practical investment environments. The primary operational failure of modern portfolio theory (MPT) stems from its extreme sensitivity to input parameters—specifically, the expected returns vector μ and the covariance matrix Σ (Best & Grauer, 1991). Because expected returns are highly unstable and notoriously difficult to forecast, minor estimation errors in historical means are mathematically amplified during the optimization process (Ledoit & Wolf, 2025).

This sensitivity causes mean-variance optimizers to function as "estimation-error maximizers". The model tends to allocate disproportionately large weights to assets with high estimated returns and low estimated variances or negative correlations. Often, these extreme statistical anomalies are simply the result of sampling noise or historical measurement error, which leads to highly concentrated, unstable, and poorly diversified portfolios that perform poorly out-of-sample (Lai et al., 2011).

Furthermore, MPT relies on several strict assumptions that do not align with real-world market dynamics. The assumption of normally distributed asset returns is frequently violated in practice. Empirical financial asset returns routinely exhibit fat tails (excess kurtosis) and significant asymmetry (skewness), especially during systemic market stress or black swan events (Li, 2023).

By using variance as the sole measure of risk, MPT penalizes upside volatility (highly desirable positive returns) in the exact same manner as downside volatility (losses). This symmetrical treatment of risk is inconsistent with actual investor behavior, which displays loss aversion—specifically, a disproportionate sensitivity to downside losses compared to upside gains. Finally, classical MPT assumes frictionless markets with infinitely divisible assets, ignoring the operational realities of transaction costs, illiquidity premiums, taxes, and discrete lot constraints (Michaud, 1989).

An additional challenge in empirical asset pricing is the "factor zoo," where high-dimensional data streams introduce severe noise and collinearity into expected return forecasts. To mitigate this noise, quantitative managers apply advanced dimension reduction and regularization techniques (Fuhrer & Hock, 2023). For example, the least absolute shrinkage and selection operator (LASSO) uses an L_1 penalty term to enforce sparsity, creating parsimonious regression models by minimizing the number of active features and increasing prediction accuracy. In parallel, unsupervised clustering algorithms such as K-means partition the asset universe into k distinct, uncorrelated clusters, helping to ensure that the asset covariance matrix remains well-conditioned prior to optimization (Beardsley, 2008).

figure 1: The Classical Markowitz Framework: Capital Market Line (CML), Asset Allocation Spectrum, and the Symmetrical Efficient Frontier.

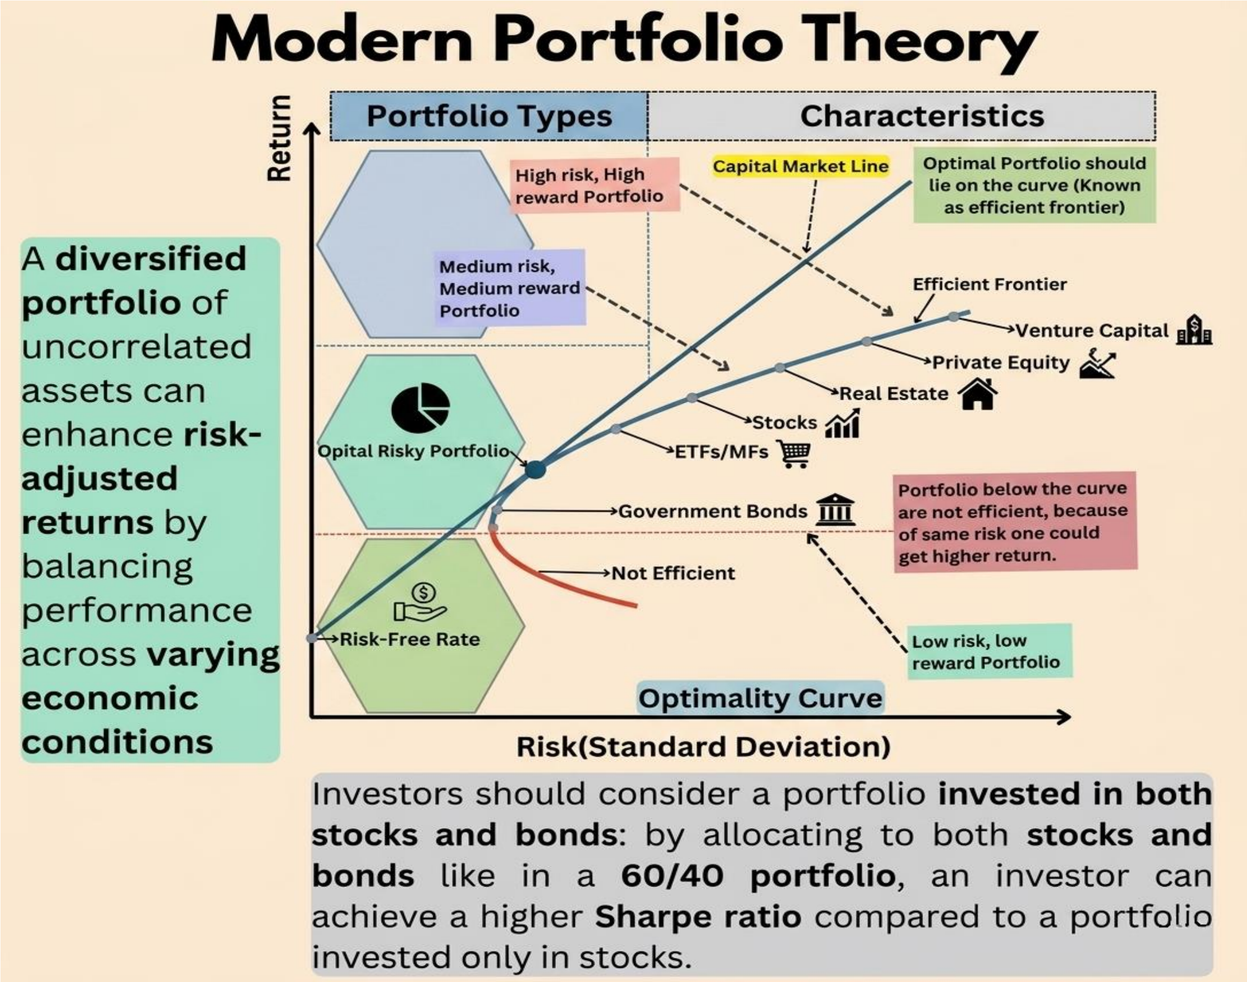

Post-Modern Portfolio Theory and High-Frequency Downside Risk

To address the limitations of treating risk symmetrically, post-modern portfolio theory (PMPT) replaces portfolio variance with asymmetric risk metrics that focus exclusively on downside risk. Under the PMPT framework, volatility above a designated threshold is not penalized, as it represents capital growth rather than risk (Rasiah, 2011). The core metric used is downside deviation (also known as target semi-deviation), which measures the dispersion of returns that fall below a user-defined minimum acceptable return (MAR), denoted by τ (Bayram & Aktaş, 2024). Mathematically, downside deviation (d_i) is calculated as the square root of the target semi-variance:

$$d_i = \sqrt{\frac{\sum_{t=1}^{T} \min(R_{it} - \tau, 0)^2}{T}}$$

where R_{it} represents the realized return of asset i at time t, and T is the total number of observations. By squaring only the negative deviations, this metric penalizes larger failures quadratically, aligning with behavioral observations of investor loss aversion (Kupiainen, 2024).

Recent econometric advances have integrated high-frequency intraday data to improve downside risk estimation. Realized variance (RV) over a daily interval can be decomposed into upside realized semivariance (RS^+) and downside realized semivariance (RS^-) (Barndorff-Nielsen et al., 2010). Under in-fill asymptotics, where the frequency of intraday observations approaches infinity, downside realized semivariance converges to the sum of continuous downside volatility and squared negative price jumps:

$$RS^- = \sum_{j=1}^{n}(Y_{t_j} - Y_{t_{j-1}})^2 \mathbb{1}_{\{Y_{t_j}-Y_{t_{j-1}} \le 0\}} \xrightarrow{p} \frac{1}{2}\int_0^1 \sigma_s^2 \, ds + \sum_{s\le 1}(\Delta Y_s)^2 \mathbb{1}_{\{\Delta Y_s \le 0\}}$$

where Y_t is the logarithmic asset price process, σ_s is continuous volatility, and ΔY_s represents discrete jump components. This high-frequency decomposition isolates the risk of sudden market declines from standard, continuous price fluctuations, providing a more precise risk metric than daily standard deviation (Barndorff-Nielsen et al., 2010).

To evaluate performance under asymmetric risk, the classical Sharpe ratio is replaced by the Sortino ratio (S), which measures downside-risk-adjusted excess return (Sortino & van der Meer, 1991):

$$S = \frac{\bar{R}_p - \tau}{d_p}$$

where $\bar{R}_p$ is the portfolio's average return and d_p is its downside deviation. Under conditions of high volatility skewness (where the ratio of variance above the mean to variance below the mean deviates from unity), the Sharpe and Sortino ratios can lead to vastly different conclusions regarding an asset's efficiency.

Alternative metrics, such as the multipliers method, further distinguish between different segments of the downside distribution, separating unrealized return deviations from actual capital losses to establish a more granular risk measure (Sortino & Satchell, 2001).

Post-modern portfolio theory (PMPT) also reformulates the capital asset pricing model (CAPM) into the downside capital asset pricing model (D-CAPM) by replacing classical beta with downside beta (β_i^D). This adjustment measures an asset's systematic sensitivity to market declines (Cheryomushkin, 2012):

$$\beta_i^D = \frac{\text{Cov}(R_i, R_m^-)}{\text{Var}(R_m^-)}$$

$$\text{Cov}(R_i, R_m^-) = \frac{1}{T}\sum_{t=1}^{T}[\min(R_{it} - E(R_i),0) \cdot \min(R_{mt} - E(R_m),0)]$$

$$\text{Var}(R_m^-) = \frac{1}{T}\sum_{t=1}^{T}\min(R_{mt} - E(R_m),0)^2$$

where R_m^- represents the negative deviations of the market portfolio returns. D-CAPM computes the expected return of asset i as:

$$E(R_i) = R_f + \beta_i^D\big(E(R_m) - R_f\big)$$

This model provides a more robust pricing mechanism in markets characterized by high volatility skewness, where classical CAPM systematically misprices assets with asymmetric return distributions (Estrada, 2002).

Table 1: Comparative Analysis of Risk Frameworks: From Classical MPT to High-Frequency Downside Metrics

Risk Dimension	Classical MPT Framework	PMPT Downside Framework	High-Frequency Realised Framework
Primary Metric	Standard Deviation (Σ)	Downside Deviation (d_i)	Downside Realised Semivariance (RS^-)
Statistical Input	Historical Covariance Matrix (Σ)	Semi-covariance / Downside Beta (\beta_i^D)	Intraday High-Frequency Jumps (\Delta Y_s)
Symmetry Bias	Symmetrical (penalizes upside volatility)	Downside Asymmetrical (penalizes below MAR)	Jump Asymmetrical (penalizes negative tail jumps)
Performance Metric	Sharpe Ratio	Sortino Ratio	Downside Jump-Adjusted Ratios

The Black-Litterman Model: A Bayesian Framework for Parameter Stabilization

The Black-Litterman model, developed by Fischer Black and Robert Litterman in 1990, addresses the high input sensitivity and poor out-of-sample performance of classical mean-variance optimization. Instead of relying solely on historically derived means, the model uses a Bayesian framework to combine a neutral market equilibrium prior with subjective investor views (Estrada, 2002).

The Bayesian Formulation

In the Black–Litterman framework, the expected asset returns μ are modeled as a latent random variable rather than a fixed vector of historical averages (Black & Litterman, 1992). The model uses a prior distribution representing market equilibrium, which is updated by a likelihood distribution representing subjective views, yielding a posterior distribution of combined expected returns.

To establish a stable, neutral starting point, the model uses "reverse optimization" (Sharpe, 1974). Under the assumption that the market capitalization-weighted portfolio (w_{mkt}) is optimal, the implied excess equilibrium returns vector (Π) is backed out using the historical covariance matrix (Σ) and the market's risk aversion coefficient (δ):

$$\Pi = \delta \Sigma w_{\text{mkt}}$$
$$\delta = \frac{E(R_m) - R_f}{\sigma_m^2}$$

where δ represents the market price of risk, scaled by the expected excess market return over its variance. The prior expected return vector μ is assumed to be normally distributed around this implied equilibrium:

$$\mu \sim N(\Pi, \tau\Sigma)$$

where τ is a small scalar (typically between 0.01 and 0.05) representing the estimation uncertainty of the prior belief.

The investor specifies k views on the assets using three main parameters:

- **Pick matrix** (P)**:** A $k \times N$ matrix where each row represents a view. For an absolute view on asset i, $P_{k,i} = 1$ and all other elements are zero. For a relative view (e.g., asset i outperforming asset j by 2%), $P_{k,i} = 1$ and $P_{k,j} = -1$.
- **Views vector** (Q)**:** A $k \times 1$ vector containing the expected returns of the views.
- **Uncertainty matrix** (Ω)**:** A diagonal $k \times k$ covariance matrix representing the uncertainty of each view. A common convention is to scale this matrix by the asset covariance (He & Litterman, 2002):

$$\Omega = \mathrm{diag}(P(\tau\Sigma)P^T)$$

The views are expressed as:

$$P\mu = Q + \epsilon, \epsilon \sim N(0, \Omega)$$

Using Bayesian updating, the prior distribution and the likelihood of the views are combined to derive the posterior expected returns vector (μ_{BL}) and posterior covariance matrix (Σ_{BL}) (Theil, 1971; Satchell & Scowcroft, 2000):

$$\mu_{BL} = [(\tau\Sigma)^{-1} + P^T\Omega^{-1}P]^{-1}[(\tau\Sigma)^{-1}\Pi + P^T\Omega^{-1}Q]$$
$$\Sigma_{BL} = \Sigma + [(\tau\Sigma)^{-1} + P^T\Omega^{-1}P]^{-1}$$

To avoid numerical instability when inverting $(\tau\Sigma)$ if it is ill-conditioned, the master formulas can be rewritten in a computationally stable form using the Woodbury matrix identity (He & Litterman, 2002):

$$\mu_{BL} = \Pi + \tau\Sigma P^T(P\tau\Sigma P^T + \Omega)^{-1}(Q - P\Pi)$$
$$\Sigma_{BL} = \Sigma + \tau\Sigma - \tau\Sigma P^T(P\tau\Sigma P^T + \Omega)^{-1}P\tau\Sigma$$

Alternatively, Meucci's formulation avoids the use of the prior uncertainty scalar τ by introducing a distinct confidence parameter $c > 0$ to construct the view uncertainty matrix directly (Meucci, 2010):

$$\Omega = \frac{1}{c}P\Sigma P^T$$

To determine user-specified confidence levels systematically without requiring direct variance inputs, Idzorek's method computes the optimal weights under complete certainty ($\Omega = 0$) and then scales the implied portfolio tilt using a user-specified confidence percentage between 0% and 100%, solving back for the unique diagonal elements of Ω (Idzorek, 2007).

The mathematical equivalence between the Bayesian formulation and classical multivariate projection is established using classical normal projection equations (Chen et al., 2015):

$$E(\mu \mid y = Q) = \Pi + \Sigma P^T(P\Sigma P^T + \Omega)^{-1}(Q - P\Pi)$$

This representation demonstrates that the Black–Litterman expected return is a linear projection of the prior mean, adjusted by the standardized prediction errors of the investor's views (Black & Litterman, 1992).

Robust Optimization and Coherent Extreme Risk Modeling

Robust Portfolio Optimization (RPO) addresses parameter uncertainty by incorporating error margins directly into the mathematical formulation. Instead of relying on single-point estimates of returns (mu) and covariances (Σ), RPO assumes these parameters lie within bounded uncertainty sets centered around their estimated values. The model then optimizes against the worst-case scenario within those defined boundaries (Flández & Rubilar-Torrealba, 2025).

Mixed Conditional Value-at-Risk (MCVaR) and Kernel Extensions

To capture tail risks across multiple segments of a distribution, managers use mixed conditional value-at-risk (MCVaR). This framework applies a weighted combination of conditional value-at-risk (CVaR) measures calculated across different confidence thresholds (Kalinchenko et al., 2012):

$$\mathrm{MCVaR}_{\boldsymbol{\theta},\boldsymbol{\beta}}(w) = \sum_{j=1}^{m} \theta_j \ \mathrm{CVaR}_{\beta_j}(w)$$

where $\sum\theta_j = 1$, and each $\theta_j > 0$ determines the priority assigned to different levels of tail severity. This formulation allows managers to construct portfolios that are customized to specific risk preferences, focusing on different quantiles of the return distribution (Nakagawa et al., 2020).

To handle uncertainty without imposing strict parametric assumptions on the return distribution, the chance constraints associated with expected returns can be reformulated using kernel methods in a reproducing kernel Hilbert space (RKHS). By mapping the empirical returns to a high-dimensional space via a kernel function, the chance constraint can be restructured as a deterministic second-order cone constraint (Yadav & Mehra, 2025). Under empirical testing, robust MCVaR portfolios optimized with kernel-based chance constraints demonstrate a strong ability to limit drawdowns during bear market phases. In contrast, during bullish and neutral phases, their performance remains highly competitive with standard nominal models, showing the value of robust formulations during periods of systemic market stress (Long et al., 2024).

Modeling Real-World Frictions and Mixed-Integer Constraints

When real-world market frictions are incorporated into the portfolio construction process, the classical mean-variance framework is no longer solvable using standard quadratic programming algorithms. These realistic constraints introduce non-convexities, discrete decision boundaries, and discontinuous feasible regions, transforming the optimization into a Mixed-Integer Nonlinear Programming (MINLP) problem, which is NP-hard (Goel & Sharma, 2020).

Heuristic Optimization and Constraint Interaction

To navigate the non-convex search spaces created by these real-world constraints, quantitative managers use metaheuristic algorithms, such as Genetic Algorithms (GA), Tabu Search (TS), Simulated Annealing (SA), and Threshold Accepting (TA). These algorithms search the solution space iteratively, balancing exploration (discovering new regions) with exploitation (refining local solutions) (Meucci, 2008).

Table 2: Comparative Analysis of Metaheuristic Strategies in Portfolio Optimization

Metaheuristic Strategy	Algorithmic Mechanism	Key Advantage	Performance Limitation
Genetic Algorithm	Recombination, mutation, and selection of weight vectors.	High global convergence in discontinuous spaces.	Computationally intensive.
Tabu Search	Local neighborhood search with a memory-driven tabu list.	Avoids cycling and local minima traps.	Performance degrades as assets scale.
Simulated Annealing	Parameter-controlled cooling with stochastic acceptance.	Strong global convergence guarantees.	Slow execution speeds.
Threshold Accepting	Local search with deterministic threshold-bound steps.	Faster execution than Simulated Annealing.	Highly sensitive to threshold parameters.

Comparative empirical evaluations on large indices (up to 1,318 assets) reveal that Genetic Algorithms outperform Tabu Search by up to three orders of magnitude when optimizing portfolios with floor, ceiling, and cardinality constraints (Rollinger & Hoffman, 2013).

Crucially, empirical models show that floor constraints actually degrade overall portfolio performance. When portfolios incorporate non-linear transaction costs alongside cardinality limits, the optimal allocation often includes several assets with very low weights (Glawischnig & Seidl, 2011). These small positions serve a critical mathematical function: they diversify risk without triggering large transaction penalties. Floor constraints prevent these small allocations, forcing the optimizer to either exclude the asset entirely or purchase it in an economically inefficient quantity, which damages diversification and reduces risk-adjusted returns (Rollinger & Hoffman, 2013).

Additionally, transaction costs act as an operational friction over a specified holding period, H. In-sample portfolio optimization must incorporate H to calibrate the trade-off between the upfront cost of rebalancing and the expected risk-reduction benefits (Rom & Ferguson, 1993). In the presence of transaction costs, trading typically reduces both expected returns and realized risk. Ignoring these frictions leads to high turnover and portfolios that perform poorly out-of-sample compared to their in-sample projections (Lee & Eid Junior, 2018).

Machine Learning, Deep Architectures, and Decision-Focused Optimization

The integration of artificial intelligence and machine learning has reshaped the landscape of portfolio construction, moving beyond the historical limits of classical MPT. These modern approaches focus on improving return predictions, capturing complex asset relationships, and aligning predictive modeling directly with downstream investment decisions (Markowitz, 1991).

The Predict-then-Optimize Paradigm vs. End-to-End SPO

Traditionally, portfolio optimization operates as a sequential, two-stage process:

1. **Prediction phase:** Statistical or machine learning models (e.g., random forests, support vector regression) are trained to predict asset returns or covariances. This phase typically minimizes a standard prediction error metric, such as mean squared error (MSE) (Rasheed et al., 2022).
2. **Optimization phase:** The predicted values are fed into an optimizer (e.g., a mean-variance engine) to calculate the final portfolio weights.

A key limitation of this two-stage approach is that it decouples prediction from optimization. Because the predictive model's loss function (such as MSE) is indifferent to how the predictions are used, small forecasting errors in directions that heavily impact the optimization can lead to highly suboptimal portfolio allocations (Markowitz, 1952).

To address this, the smart predict-then-optimize (SPO) paradigm integrates prediction and optimization into a single, end-to-end framework. SPO updates the weights of the predictive model by backpropagating a decision-focused loss function that measures the downstream portfolio performance (such as risk-adjusted return or drawdown) directly (Rutterford & Sotiropoulos, 2016) :

$$\text{Loss}_{\text{SPO}}(\hat{\mu},\mu) = f(w^*(\hat{\mu}),\mu) - f(w^*(\mu),\mu)$$

where $w^*(\hat{\mu})$ represents the optimal portfolio weights computed using the predicted returns, and $w^*(\mu)$ represents the theoretical optimal weights if the true returns were known. This approach ensures the machine learning model is trained to optimize the final decision quality rather than isolated statistical accuracy (Edlinger & Parent, 2014).

Advanced Neural Architectures and Graph Networks

To capture temporal dynamics and complex relationships across financial assets, managers use advanced deep learning architectures:

- **Temporal Modeling:** Long Short-Term Memory (LSTM) networks model sequential dependencies in historical financial time series, capturing long-term trends and volatility patterns (Estrada, 2002).
- **Relational Modeling:** Graph Attention Networks (GATs) model inter-asset dependencies dynamically. By treating assets as nodes in a graph with dynamically updated edges, GATs capture spillover effects, sector connections, and supply chain relationships, avoiding the limitations of static correlation estimates (Sortino & Satchell, 2001).

- **Alternative Data Integration:** Natural language processing models extract sentiment from financial news headlines. This sentiment is combined with temporal and relational features in a unified, end-to-end framework to learn optimal portfolio weights directly.

In parallel, reinforcement learning (RL) models are used for direct optimization, bypass the need for separate return and covariance estimations by training agents to maximize risk-adjusted returns directly through continuous market interaction (Idzorek, 2007).

Table 3: Comparative Evaluation of Portfolio Optimization Frameworks: From Classical MPT to Decision-Focused Machine Learning

Optimization Framework	Risk Metric	Parameter Input Source	Mathematical Optimization Class	Robustness to Estimation Error	Real-World Friction Handling
Classical MPT	Symmetrical Variance (\Σ _p^2)	Historical returns and covariance matrix	Quadratic Programming (QP)	Low (Estimation-error maximizer)	Poor (Assumes frictionless markets)
Post-Modern (PMPT)	Downside Deviation (d_i)	Target semi-variance and downside beta	Non-convex Non-linear Programming	Moderate (Focuses on downside tail)	Moderate (Can incorporate target returns)
Black-Litterman	Bayesian Posterior Covariance (\Σ _{BL})	Market equilibrium prior and subjective investor views	Quadratic Programming (utilizing posterior inputs)	High (Shrinks estimates toward market equilibrium)	Moderate (Allows systematic tilt inputs)
Robust Optimization	Worst-case Risk (within uncertainty set)	Parametric bounded uncertainty sets	Second-Order Cone Programming (SOCP)	Very High (Optimizes against worst-case scenario)	Moderate (Accommodates parameter bounds)
Machine Learning End-to-End (SPO)	Decision-focused SPO Loss	Dynamic neural representations (temporal/graphical)	Mixed-Integer Deep Learning Optimization	High (Directly aligns predictions with decisions)	High (Integrates transaction costs into loss function)

Conclusion

The evolution of portfolio optimization reflects a continuous effort to bridge the gap between elegant mathematical theory and the volatile, non-convex realities of financial markets. While Markowitz's mean-variance paradigm laid the groundwork for quantitative finance, its real-world utility remains constrained by its symmetrical definition of risk and high sensitivity to parameter estimation errors. Over several decades, the field has adapted by introducing asymmetrical risk dimensions through Post-Modern Portfolio Theory, stabilizing parameters via Bayesian Black-Litterman frameworks, and engineering robust, worst-case optimization strategies to buffer against tail risks and market frictions. The cutting-edge frontier of asset allocation now belongs to machine learning and end-to-end decision-focused architectures. By transitioning from the disjointed "predict-then-optimize" routine to integrated Smart Predict-then-Optimize (SPO) frameworks, quantitative finance can now align predictive statistical accuracy directly with ultimate portfolio performance. Coupled with deep temporal and graph attention networks that process alternative data and complex, dynamic inter-asset relationships, modern portfolio construction has shifted from a retrospective historical science to an adaptive, forward-looking engineering discipline. Future research and practical applications must continue to focus on this intersection, blending robust optimization techniques with deep learning to build portfolios that are both computationally scalable and resilient to systemic shocks.

REFERENCES

Abdelhakmi, A., & Lim, A. E. B. (2024). A multi-period Black-Litterman model. *SSRN Electronic Journal*. https://doi.org/10.2139/ssrn.4811035

Barndorff-Nielsen, O. E., Kinnebrock, S., & Shephard, N. (2010). Measuring disadvantageous risk: Realized semivariances. *In Oxford Handbook of Economic Forecasting*, 391-423.

Bayram, E., & Aktaş, R. (2024). Portfolio construction with postmodern portfolio theory framework. *Journal of Business and Economic Studies*, *12*(1), 25-38.

Beardsley, X. W. (2008). Mean-Variance-Skewness-Kurtosis Portfolio Optimization with Return and Liquidity. *Journal of Applied Finance & Banking*, *3*(1), 1-22.

Best, M. J., & Grauer, R. R. (1991). On the sensitivity of mean-variance efficient portfolios to changes in asset means. *The Review of Financial Studies*, *4*(2), 315-342.

Black, F., & Litterman, R. (1991). Asset allocation: Combining investor views with market equilibrium. The Journal of Fixed Income, 1(2), 7–18. https://doi.org/10.3905/jfi.1991.408013

Chen, L., Da, Z., & Schaumburg, E. (2015). Implementing Black-Litterman using an equivalent formula and equity analyst target prices. *The Journal of Investing*, *24*(1), 34–47. https://doi.org/10.3905/joi.2015.24.1.034

Cheryomushkin, S. (2012). Internal Inconsistency of Downside CAPM Models. *Journal of Corporate Finance Research*, *5*(4), 90–111. https://doi.org/10.17323/j.jcfr.2073-0438.5.4.2011.90-111 Cited by: 4

Duhis, M., Mehmood, A., Ahmed, A., & Lawa, A. (2025). Sustainable Development in the Face of Climate Change Integration Climate Mitigation and Adaptation into Urban Planning. *The Critical Review of Social Sciences Studies*, 3(1), 2377-2390.

Edlinger, C., & Parent, A. (2014). THE BEGINNINGS OF A 'COMMON-SENSE' APPROACH TO PORTFOLIO THEORY BY NINETEENTH-CENTURY FRENCH FINANCIAL ANALYSTS PAUL LEROY-BEAULIEU AND ALFRED NEYMARCK. *Journal of the History of Economic Thought*, *36*(1), 23–44. https://doi.org/10.1017/s1053837214000017

Elton, E. J., & Gruber, M. J. (1997). Modern portfolio theory, 1950 to date. *Journal of Banking & Finance*, *21*(11), 1743–1759. https://doi.org/10.1016/s0378-4266(97)00048-4

Estrada, J. (2002). Systematic risk in emerging markets: the D-CAPM. *Journal of Applied Finance*, *12*(2), 36–42.

Flández, S., & Rubilar-Torrealba, R. (2025). Black–Litterman portfolio optimization with dynamic CAPM via ABC-MCMC. *Mathematics*, *13*(20), 3265. https://doi.org/10.3900/sub-placeholder-or-crossref

Fuhrer, A., & Hock, T. (2023). Uncertainty in the Black–Litterman model: Empirical estimation of the equilibrium. *Journal of Empirical Finance*, *72*, 251-275. https://doi.org/10.1016/j.jempfin.2023.03.009

Glawischnig, M., & Seidl, I. (2011). Portfolio optimization with serially correlated, skewed and fat tailed index returns. *Central European Journal of Operations Research*, *21*(1), 153-176. https://doi.org/10.1007/s10100-011-0219-2

Goel, A., & Sharma, A. (2020). Mixed value-at-risk and its numerical investigation. *Physica A: Statistical Mechanics and its Applications*, *541*, 123524. https://doi.org/10.1016/j.physa.2019.123524

He, G., & Litterman, R. (2002). The intuition behind Black-Litterman model portfolios. Goldman Sachs Asset Management Investment Management Research Paper.

Idzorek, T. M. (2007). A step-by-step guide to the Black-Litterman model: Incorporating user-specified confidence levels. In S. Satchell (Ed.), Forecasting Expected Returns in the Financial Markets (pp. 17–38). Butterworth-Heinemann. https://doi.org/10.1016/B978-075068321-0/50003-1

Kalinchenko, K., Uryasev, S., & Rockafellar, R. T. (2012). Calibrating risk preferences with the generalized capital asset pricing model based on mixed conditional value-at-risk deviation. *The Journal of Risk*, *15*(1), 45-70. https://doi.org/10.21314/jor.2012.249

Kolari, J. W., & Pynnönen, S. (2023). Portfolio Theory and Practice. *Investment Valuation and Asset Pricing*, 1–15. https://doi.org/10.1007/978-3-031-16784-3_1

Kupiainen, P. (2024). *A quantitative comparison of modern and post-modern portfolio theory optimal portfolios* (Master's thesis, Aalto University).

Lai, T. L., Xing, H., & Chen, Z. (2011). Mean–variance portfolio optimization when means and covariances are unknown. *The Annals of Applied Statistics*, *5*(2A), 798–823. https://doi.org/10.1214/10-aoas422

Ledoit, O., & Wolf, M. (2025). Markowitz portfolios under transaction costs. *The Quarterly Review of Economics and Finance*, *100*, 101962. https://doi.org/10.1016/j.qref.2025.101962

Lee, S. C., & Eid Junior, W. (2018). Portfolio construction and risk management: Theory versus practice. *RAUSP Management Journal*, *53*(3), 345-365. https://doi.org/10.1108/rausp-04-2018-009

Li, J. (2023). Optimal Portfolio Choice with Fat Tails. *National Bureau of Economic Research (NBER) Working Paper*.

Long, Y., Ok, T., Scroccaro, P. Z., & Esfahani, P. M. (2024). Scalable kernel inverse optimization. *arXiv preprint arXiv:2410.23952*. https://doi.org/10.48550/arxiv.2410.23952

Markowitz, H. (1952). Portfolio selection. *The Journal of Finance*, 7(1), 77–91. https://doi.org/10.1111/j.1540-6261.1952.tb01525.x

Mehmood, A. (2023). Data-driven decision-making in local and global markets: The role of market research in pricing, branding, and distribution strategies. *Pakistan Social Sciences Review*, 7(3), 1225-1235.

Mehmood, A. (2025). Transformational leadership and gender diversity: A mixed-methods study on the effectiveness of women in US Fortune 500 companies. *The American Journal of Management and Economics Innovations*, 7(3), 30-38.

Mehmood, A., & Lawa, A. (2025). Ethical leadership as a catalyst for organizational resilience in times of crisis: a mixed-methods approach (American context). *Review of Applied Management and Social Sciences*, *8*(1), 407-420.

Mehmood, A., Siddiqui, H. A., & Rashid, M. (2025). Mentorship as a catalyst for leadership development: Strategies, outcomes, and implications for emerging leaders in organizational contexts. *Social Science Review Archives*, *3*(1), 2516-2525.

Meucci, A. (2008). The Black–Litterman approach: Common misconceptions and effective techniques. Risk Magazine, 21(9), 106–111.

Meucci, A. (2010). Fully flexible views: Theory and practice. Risk Magazine, 23(10), 78–83.

Michaud, R. O. (1989). The Markowitz optimization enigma: Is 'optimized' optimal?. *Financial Analysts Journal*, *45*(1), 31-42.

Nakagawa, R., Noma, S., & Abe, M. (2020). RM-CVaR: Regularized multiple \beta-CVaR portfolio. *arXiv preprint arXiv:2006.0629*. https://doi.org/10.24963/ijcai.2020/629

Rasheed, H., Ahmad, Habib., Javid, A. Y., & Khawaja, I. (2022). Is real estate a real diversifier in Pakistan? An ARDL approach. *Cogent Business & Management*, 9(1). https://doi.org/10.1080/23311975.2022.2082013

Rasiah, D. (2011). Post-modern portfolio theory supports diversification in an investment portfolio to measure investment's performance. *Journal of Financial and Econometric Research*, *1*(1), 68-78.

Rollinger, T. N., & Hoffman, S. T. (2013). Sortino: A 'sharper' ratio. *CME Group Education*, 1-4.

Rom, B. M., & Ferguson, K. W. (1993). Post-modern portfolio theory comes of age. *The Journal of Investing*, *2*(3), 11-17.

Rutterford, J., & Sotiropoulos, D. P. (2016). Financial diversification before modern portfolio theory: UK financial advice documents in the late nineteenth and the beginning of the twentieth century. *The European Journal of the History of Economic Thought*, *23*(6), 919–945. https://doi.org/10.1080/09672567.2016.1203968

Satchell, S., & Scowcroft, A. (2000). A demystification of the Black–Litterman model with applications to asset allocation. Journal of Asset Management, 1(2), 138–150. https://doi.org/10.1057/palgrave.jam.2240011

Sharpe, W. F. (1974). Implied risk premiums. Research Paper No. 192, Stanford University Graduate School of Business.

Sortino, F. A., & Satchell, S. (Eds.). (2001). *Managing downside risk in financial markets*. Butterworth-Heinemann.

Sortino, F. A., & van der Meer, R. (1991). Downside-risk measurement in the practice of asset management. *Financial Analysts Journal*, *47*(5), 27–31.

Theil, H. (1971). Principles of Econometrics. John Wiley & Sons.

Yadav, R., & Mehra, A. (2025). Robust MCVaR portfolio optimization with ellipsoidal support and reproducing kernel Hilbert space-based uncertainty. *arXiv preprint arXiv:2509.00447*. https://arxiv.org/abs/2509.00447