

Portfolio Optimization Platform for Investment

Abstract

This dissertation investigates the integration of classical financial theory and machine learning in the domain of portfolio optimisation. Classical models such as mean–variance optimisation, Capital Asset Pricing Model (CAPM), Monte Carlo simulations, and Value-at-Risk (VaR) have long provided the theoretical foundation for investment decision-making. However, their reliance on strong assumptions, particularly normality of returns and stable covariances, often results in fragile and concentrated allocations. Machine learning techniques, including Random Forests and Gradient Boosting, offer improved predictive accuracy but tend to produce unstable and high-turnover portfolios.

To address these challenges, this study develops and evaluates a hybrid framework that incorporates machine learning forecasts as probabilistic “views” within the Black–Litterman model. Using equity and ETF datasets, the research compares classical, ML-driven, and hybrid portfolios across multiple metrics including Sharpe ratio, Sortino ratio, Value-at-Risk, Conditional VaR, Jensen’s alpha, and turnover.

The results show that the hybrid portfolios are consistently superior in terms of risk-adjusted returns, stability, and lower downside exposure compared to the classical and ML-only portfolios. The strength of the hybrid structure is demonstrated through stress testing under bearish, bullish, and crisis conditions. The study also designs an effective dashboard that integrates efficient frontiers, allocations, and performance measures, making the platform accessible to both institutional and retail investors.

The primary contribution of this dissertation is the demonstration that introducing predictive intelligence into conventional optimisation models allows portfolio managers to achieve both theoretically sound and practically robust portfolios in the dynamic environment of modern finance.

Contents

1	Introduction	12
1.1	Problem Statement	12
1.2	Research Objectives	13
1.3	Research Questions	13
1.4	Methodological Overview	13
1.5	Significance of the Study	14
1.6	Structure of the Thesis	14
2	Literature Review	15
2.1	Introduction	15
2.2	Classical Portfolio Theories	15
2.2.1	Modern Portfolio Theory (MPT)	15
2.2.2	Capital Asset Pricing Model (CAPM)	15
2.2.3	Factor Models	15
2.3	Risk and Valuation Models	16
2.3.1	Black–Scholes Option Pricing	16
2.3.2	Monte Carlo Simulation	16
2.3.3	Value-at-Risk (VaR)	16
2.4	Black–Litterman Model	16
2.5	Machine Learning in Finance	16
2.5.1	Predictive Modelling	16
2.5.2	Deep Learning and Time Series	16
2.5.3	Hybrid Approaches	17
2.6	Existing Platforms and Dashboards	17
2.7	Research Gap	17
3	Methodology	18
3.1	Research Design	18
3.2	Data Collection	18
3.3	Data Preprocessing	18
3.3.1	Cleaning and Missing Values	18

3.3.2	Normalisation	18
3.3.3	Smoothing	18
3.4	Statistical and Exploratory Analysis	19
3.5	Model Development	19

3.5.1	Financial Models	19
3.5.2	Machine Learning Models	19
3.6	Optimisation and Integration	19
3.7	Performance Evaluation	20
3.8	Ethical Considerations	20
4	Implementation	21
4.1	Introduction	21
4.2	System Architecture	21
4.2.1	Layer 1 – Data Layer	21
4.2.2	Layer 2 – Analytical Layer	21
4.2.3	Layer 3 – Model Layer	21
4.2.4	Layer 4 – Optimization Layer	22
4.2.5	Layer 5 – Presentation Layer	22
4.3	Data Layer	22
4.3.1	Data Sources	22
4.3.2	Handling Missing Data	22
4.3.3	Data Transformation	22
4.3.4	Normalisation and Scaling	22
4.3.5	Outlier Treatment	23
4.3.6	Smoothing	23
4.4	Analytical Layer	23
4.4.1	Descriptive Statistics	23
4.4.2	Correlation Structure	23
4.4.3	Stationarity Tests	23
4.4.4	Volatility Clustering	23
4.4.5	Factor Sensitivity Analysis	23
4.5	Model Layer (Part 1 – Financial Models)	24
4.5.1	Capital Asset Pricing Model (CAPM)	24
4.5.2	Monte Carlo Simulation	24
4.5.3	Value-at-Risk (VaR)	24
4.5.4	Black–Litterman Model	24
4.6	Optimization Layer	24
4.6.1	Classical Mean–Variance Optimisation	25
4.6.2	Machine Learning–Driven Optimisation	25
4.6.3	Hybrid Optimisation (Black–Litterman with ML Views)	25
4.6.4	Stress Testing and Robust Optimisation	25
4.7	Presentation Layer	26
4.7.1	Dashboard Design Principles	26
4.7.2	Core Visualisations	26

4.7.3	Interactivity Features	26
4.8	Technical Challenges and Solutions	26
4.8.1	Data Quality Issues	26
4.8.2	Model Instability	27

4.8.3	Overfitting Risk.....	27
4.8.4	Computational Demands.....	27
4.8.5	User Experience.....	27
4.9	Summary.....	27
5	Results and Analysis	28
5.1	Introduction.....	28
5.2	Distribution Analysis.....	28
5.2.1	Return Distributions.....	28
5.2.2	Volatility Clustering.....	29
5.2.3	Correlations and Diversification Potential.....	29
5.2.4	Stationarity Tests.....	29
5.3	Comparative Analysis of Models.....	29
5.3.1	Classical Framework Results.....	29
5.3.2	Machine Learning Framework Results.....	29
5.3.3	Hybrid Framework Results.....	30
5.4	Performance Metrics.....	30
5.4.1	Sharpe Ratio and Alpha.....	30
5.4.2	Sortino Ratio and Treynor Ratio.....	30
5.4.3	Value-at-Risk and CVaR.....	30
5.4.4	Cumulative Returns and Stability.....	31
5.5	Visualization and Dashboard Outputs.....	31
5.5.1	Efficient Frontier Visualization.....	31
5.5.2	Allocation Dynamics.....	31
5.5.3	Performance Tables.....	31
5.5.4	Scenario Testing.....	31
5.5.5	User Interaction Features.....	32
5.6	Discussion of Findings.....	32
5.6.1	Comparative Strengths and Weaknesses.....	32
5.6.2	Theoretical Contributions.....	32
5.6.3	Practical Contributions.....	32
5.6.4	Alignment with Research Questions.....	33
5.6.5	Limitations.....	33
5.6.6	Future Research Directions.....	33
5.6.7	Practical Recommendations.....	34
6	Discussion	35
6.1	Interpretation of Results.....	35
6.1.1	Estimation Error and the Bias–Variance Trade-off.....	35

6.1.2	Regularisation through Black–Litterman.....	35
6.1.3	Covariance Estimation and Portfolio Stability.....	36
6.1.4	Risk-adjusted Performance vs. Turnover and Capacity.....	36
6.1.5	Explainability and Economic Intuition.....	36
6.2	Strengths and Weaknesses.....	36
6.2.1	Strengths.....	36
6.2.2	Weaknesses.....	37
6.3	Theoretical and Practical Contributions.....	38
6.3.1	Theoretical Contributions.....	38
6.3.2	Practical Contributions.....	38
6.4	Limitations and Future Research Directions.....	38
6.4.1	Empirical and Methodological Limitations.....	38
6.4.2	Research Directions.....	39
6.5	Practical Recommendations for Deployment.....	40
6.6	Concluding Remarks.....	40
7	Conclusion and Future Work	42
7.1	Summary of Findings.....	42
7.2	Contributions to Knowledge.....	43
7.2.1	Theoretical Contributions.....	43
7.2.2	Practical Contributions.....	43
7.2.3	Methodological Contributions.....	43
7.3	Final Remarks.....	44
	A Code Appendix: Figures	47
	B Complete Code Implementation	60
B.1	Complete Code Link.....	60

List of Figures

- A.1 Portfolio weights per stock in sorted form. (This figure illustrates the distribution of portfolio weights assigned to individual stocks, arranged in descending order. It highlights which stocks hold the largest shares within the portfolio and provides a clear view of diversification across assets.) 47
- A.2 Portfolio distribution weight represented in a pie chart. (This pie chart shows the relative weight of each stock in the portfolio, providing a quick view of asset allocation and diversification.) 48
- A.3 Portfolio cumulative returns. (Shows the cumulative returns of the portfolio over time, highlighting the overall growth trend and performance.) 48
- A.4 21-day rolling volatility with dates. (Illustrates the 21-day rolling volatility of the portfolio, showing fluctuations in risk over time along with corresponding dates.) 49
- A.5 Correlation matrix of log returns. (This correlation matrix of log returns displays the interdependencies between asset classes, with values ranging from -1 (perfect negative correlation) to 1 (perfect positive correlation).) 49
- A.6 Cumulative portfolio value over time – Linear Regression. (A linear regression line fits historical data points, indicating overall portfolio growth direction and aiding in assessing risk-return patterns.) 50
- A.7 Cumulative portfolio value over time – Gradient Boosting. (Cumulative portfolio value over time with Gradient Boosting leverages sequential machine learning models to capture complex patterns in portfolio growth, handling non-linear relationships effectively.) 50
- A.8 Cumulative portfolio value over time – Random Forest. (Cumulative portfolio value over time with Random Forest uses multiple decision trees for robust modeling of portfolio growth. Captures complex interactions for insightful predictions on investment trends.) 51
- A.9 Weight comparison – Random Forest (ML vs BL). (Random Forest weight comparison (ML vs BL) contrasts machine learning-derived feature importance with traditional backtesting (BL) approaches for portfolio insights.) 51
- Optimized Sharpe: ML vs Black–Litterman. (Optimized Sharpe ratio comparison

between Machine Learning (ML) and Black–Litterman highlights how ML techniques and the Black–Litterman model differently enhance portfolio optimization, with Black–Litterman blending market equilibrium and investor views for potentially more robust risk-adjusted returns.)

Efficient frontier & Monte Carlo – Linear Regression (ML). (The efficient frontier paired with Monte Carlo simulations using Linear Regression (ML) generates probabilistic portfolio frontiers to visualize risk-return tradeoffs, aiding investors in optimizing asset allocations based on predicted performance distributions.)

52

A.10 Efficient frontier & Monte Carlo – Linear Regression (BL). (Efficient frontier & Monte Carlo with Linear Regression (BL) illustrates traditional portfolio optimization boundaries using baseline (BL) methods, showcasing risk-return profiles via simulated asset allocation scenarios.)

53

A.11 Efficient frontier & Monte Carlo – Random Forest (ML). (Efficient frontier & Monte Carlo with Random Forest (ML) demonstrates advanced portfolio optimization using machine learning’s non-linear capabilities, potentially uncovering complex risk-return patterns beyond traditional models.)

53

A.12 Efficient frontier & Monte Carlo – Random Forest (BL). (Efficient frontier & Monte Carlo with Random Forest (BL) contrasts baseline (BL) applications of Random Forest in portfolio optimization, highlighting traditional interpretations of risk-return tradeoffs via simulated asset allocations.)

54

A.13 Efficient frontier & Monte Carlo – Gradient Boosting (ML). (Efficient frontier & Monte Carlo with Gradient Boosting (ML) leverages powerful sequential learning for nuanced portfolio optimization, capturing intricate market patterns to shape risk-return frontiers for potentially enhanced investment strategies.)

54

A.14 Efficient frontier & Monte Carlo – Gradient Boosting (BL). (Efficient frontier & Monte Carlo with Gradient Boosting (BL) applies baseline Gradient Boosting techniques for traditional portfolio analysis, focusing on standard interpretations of simulated risk-return outcomes for asset allocation insights.)

55

A.15 Final cumulative return comparison of all models used. (Final cumulative return comparison of all models synthesizes performance outcomes across machine learning (ML) and baseline (BL) approaches, highlighting relative efficacy in capturing portfolio growth trends and informing investment decisions.)

55

A.16 Cumulative return comparison across models. (Cumulative return comparison across models visually contrasts the terminal wealth outcomes of diverse portfolio strategies, spotlighting performance disparities between machine learning-driven and traditional investment approaches.)

56

A.17 Prices overview series with their values. (Prices overview series with their values displays historical price trajectories of analyzed assets, providing contextual market data underlying the portfolio optimization and modeling efforts showcased in comparative performance assessments.)	56
A.18 Return correlation with ticker. (Return correlation with ticker illustrates inter-asset return relationships via correlation matrices, aiding identification of diversification opportunities and portfolio construction insights tied to specific securities.)	57
A.19 21-day rolling volatility (Annualized) with date and values. (21-day rolling volatility (Annualized) with date and values tracks short-term risk dynamics of assets over time, spotlighting fluctuations in annualized volatility to inform risk management and asset allocation decisions.)	57
A.20 Financial model Markowitz weights (Non-Negative) values with their assets and weights. (Financial model Markowitz weights (Non-Negative) values with their assets and weights showcases optimal asset allocations from Markowitz mean-variance optimization, highlighting portfolio composition under traditional Modern Portfolio Theory constraints.)	58
A.21 Financial model Black–Litterman model (Non-Negative) values with their assets and weights. (Financial model Black–Litterman model (Non-Negative) values with their assets depicts portfolio weights blending equilibrium returns and investor views via Black–Litterman, reflecting tailored asset allocations incorporating subjective insights.)	58
A.22 ML model daily composite views with their assets and weights. (ML model daily composite views with their assets and weights reveals machine learning-derived daily asset views for portfolio steering, indicating data-driven perspectives on relative attractiveness influencing investment positioning.)	59
A.23 Out of sample cumulative return of models used with their values and dates. (Out of sample cumulative return of models used with their values and dates contrasts realized performance trajectories of tested portfolio strategies beyond training periods, assessing comparative investment efficacy in unseen market conditions.)	59

Chapter 1

Introduction

The primary problem of financial decisions has always been the efficiency of portfolios. [Markowitz \(1952\)](#) was the first to advance modern portfolio theory (MPT) to give a formal description of the risk–return trade-off as a consequence of mean–variance optimisation. [Sharpe \(1964\)](#) further expanded this work into the Capital Asset Pricing Model (CAPM) which quantified the expected return using systematic risk. Subsequently, [? \(1973\)](#) introduced the option pricing model, which generalized classical finance theory to derivative securities and time series. These foundations illustrate that while classical models provide clarity, they depend on assumptions (normality of returns, stable covariances) that may not hold in practice.

These limitations have motivated the search for more flexible, data-driven solutions. In particular, machine learning’s ability to find non-linear relationships in complex data has become significant ([Gu et al., 2020](#)). By combining classical financial theory and machine learning, this study aims to formulate a portfolio optimisation model that targets potent, flexible, and interpretable investment strategies. The platform will integrate predictive modeling into established frameworks like the Black–Litterman model, enabling investors to benefit from advanced forecasts while retaining clear decision insights.

1.1 Problem Statement

Classical portfolio optimisation methods have several drawbacks. First, financial markets are noisy, and it is difficult to extract predictive information from short-term returns ([Lo, 2004](#)). Second, optimisation is highly sensitive to estimation errors in covariances, resulting in unstable allocations ([Ledoit and Wolf, 2004](#)). Third, investor equilibrium assumptions can be subjective and may lead to suboptimal forecasts ([Meucci, 2010](#)). Finally, many existing tools are neither interactive nor transparent, limiting practical adoption.

Therefore, the research question is: How can we build a portfolio optimisation model that leverages machine learning's predictive capabilities while retaining the interpretability of classical finance models?

1.2 Research Objectives

This thesis focuses on creating and testing a hybrid portfolio optimisation platform. The specific objectives are:

- To implement a structured process for portfolio construction, incorporating both statistical and predictive approaches.
- To evaluate classical models (CAPM, Monte Carlo simulation, VaR, and the Black–Litterman model) within the platform.
- To apply machine learning algorithms (Linear Regression, Random Forest, Gradient Boosting) to forecast expected returns.
- To compare outcomes of classical, machine learning, and hybrid approaches based on performance indicators such as return, volatility, and Sharpe ratio.
- To build an interactive platform for investors and researchers to visualize and analyze portfolio results.

1.3 Research Questions

The study addresses the following questions:

- To what extent do machine learning forecasts improve portfolio performance compared to traditional financial models?
- How does integrating predictive views into the Black–Litterman framework affect allocation stability and out-of-sample performance?
- Which modeling approaches provide the best balance between predictive accuracy, risk management, and interpretability?

- What features are essential for an interactive platform to serve both academic and practical investment needs?

1.4 Methodological Overview

A multi-stage methodology is adopted. Financial market data is collected for a diversified set of securities, then preprocessed by handling missing values, normalizing, and smoothing. Exploratory statistical analysis identifies distributional characteristics, correlations, and volatility clusters.

Classical models (CAPM, Monte Carlo simulation, VaR, and Black–Litterman) are used to compute portfolio allocations under theoretical assumptions. In parallel, machine learning algorithms are trained on lagged financial features to forecast asset returns. These forecasts are incorporated into optimisation both directly and as “views” in the Black–Litterman model. Performance evaluation uses both financial metrics (Sharpe ratio, alpha, VaR) and predictive metrics (mean squared error, R^2). Results are analyzed to determine which framework best meets multiple objectives.

1.5 Significance of the Study

Academically, this research contributes by empirically demonstrating how machine learning forecasts can be integrated into classical portfolio optimisation models. Practically, it delivers a functional, interactive platform that provides investors with enhanced decision-making tools. By balancing predictive accuracy with interpretability, the study addresses a central tension in quantitative finance and offers insights relevant to scholars and practitioners.

1.6 Structure of the Thesis

The remainder of the thesis is organized as follows:

- Chapter 2: Literature Review — Surveys classical portfolio theories, financial risk and valuation models, and applications of machine learning in finance.
- Chapter 3: Methodology — Describes data collection, preprocessing, model design, and evaluation methods.
- Chapter 4: Implementation — Explains the architecture of the platform and its

components.

- Chapter 5: Results and Analysis — Presents and compares the empirical results of different modelling approaches.
- Chapter 6: Discussion — Analyzes the strengths, limitations, and implications of the findings.
- Chapter 7: Conclusion and Future Work — Summarizes contributions and outlines directions for further research

Chapter 2

Literature Review

2.1 Introduction

Portfolio optimisation has been extensively studied, beginning with mean–variance theory and expanding to advanced risk and predictive models. This chapter reviews classical portfolio theories, financial risk models, and modern data-driven approaches to investment.

2.2 Classical Portfolio Theories

2.2.1 Modern Portfolio Theory (MPT)

[Markowitz \(1952\)](#) introduced mean–variance optimisation, arguing that investors should focus on portfolios rather than individual securities. By balancing expected return and variance, investors can achieve the efficient frontier of optimal portfolios. This theory is foundational but relies on simplifying assumptions (e.g., normality of returns and stable correlations).

2.2.2 Capital Asset Pricing Model (CAPM)

[Sharpe \(1964\)](#) extended MPT by developing CAPM, which relates an asset's expected return to its systematic risk (beta). CAPM posits that investors are rewarded only for bearing systematic risk, as idiosyncratic risk can be diversified. While widely used, empirical tests have found anomalies (size and value effects) that challenge CAPM's predictions (?).

2.2.3 Factor Models

Multi-factor models, such as the Fama–French three-factor and five-factor models, incor-

porate factors like size, value, profitability, and investment. These models provide a more nuanced explanation of asset returns but still assume linearity and stationarity.

2.3 Risk and Valuation Models

2.3.1 Black–Scholes Option Pricing

? revolutionized derivatives pricing. Their model assumes lognormal price distribution and continuous hedging, but it struggles with volatility clustering and jumps ([Cont, 2001](#)).

2.3.2 Monte Carlo Simulation

Monte Carlo methods simulate thousands of return paths (?), allowing modeling of non-linear instruments and path-dependent payoffs. They require significant computation and careful calibration.

2.3.3 Value-at-Risk (VaR)

VaR estimates the maximum expected loss at a given confidence level ([Jorion, 2006](#)). It is intuitive but criticized for not being sub-additive and underestimating tail risks (?). CVaR (Conditional VaR) is often used to address this limitation.

2.4 Black–Litterman Model

[Black and Litterman \(1992\)](#) proposed a Bayesian approach to portfolio optimisation that blends market equilibrium returns with investor views. It stabilises mean–variance optimisation by avoiding extreme allocations due to estimation errors. Extensions incorporate view uncertainty and alternative priors ([Meucci, 2010](#)).

2.5 Machine Learning in Finance

2.5.1 Predictive Modelling

Machine learning enables non-linear prediction on large datasets. [Gu et al. \(2020\)](#) show

that ML methods (random forests, neural networks) improve return prediction accuracy compared to linear models, capturing complex patterns in cross-sectional returns.

2.5.2 Deep Learning and Time Series

Deep learning models (RNNs, LSTMs) can model temporal dependencies in volatility and risk forecasting (Fischer and Krauss, 2018). These models capture dynamics but risk overfitting when data is limited.

2.5.3 Hybrid Approaches

Recent studies argue that ML predictions should complement, not replace, financial models. For example, using ML forecasts as additional “views” in Black–Litterman produces models that are both predictive and stable (?).

2.6 Existing Platforms and Dashboards

Interactive portfolio optimisation platforms like QuantConnect and Portfolio Visualizer offer backtesting and optimisation. However, they often lack a seamless integration between classical models and ML. Our thesis aims to fill this gap by developing a platform that bridges theory and data-driven prediction.

2.7 Research Gap

Classical approaches are well-understood but limited by estimation errors and unrealistic assumptions. Pure ML approaches are powerful but often lack interpretability and financial context. A hybrid framework that unifies these methods remains underexplored. This work seeks to address that gap by creating a transparent platform combining classical finance with ML insights.

Chapter 3

Methodology

3.1 Research Design

This research employs a quantitative multi-stage design: (1) data collection and preprocessing, (2) exploratory analysis, (3) model development, (4) optimisation and integration, and (5) performance evaluation.

3.2 Data Collection

Daily price data for a diversified basket of U.S. equities and ETFs was collected over a two-year period. The dataset ensures sectoral diversification. Data frequency was set to daily, consistent with portfolio optimisation studies ([Lo, 2004](#)).

3.3 Data Preprocessing

3.3.1 Cleaning and Missing Values

Financial data often contains missing observations due to trading halts or holidays. Missing values were imputed using forward and backward fill methods to preserve continuity (no artificial bias) ([Tsay, 2010](#)).

3.3.2 Normalisation

To address scale differences, features were standardized using z-score normalization. Log returns were computed to stabilise variance and approximate normality, a common practice in financial econometrics (?).

3.3.3 Smoothing

Simple and exponential moving averages were applied to smooth high-frequency noise and highlight trends. These transformations support exploratory analysis and informed feature construction for predictive models.

3.4 Statistical and Exploratory Analysis

Exploratory Data Analysis (EDA) quantified distributional properties, volatility clustering, and correlations. Descriptive statistics, autocorrelation functions (ACF), and correlation matrices were generated to understand the dataset and guide model selection.

3.5 Model Development

3.5.1 Financial Models

- **CAPM:** Estimates expected return as a function of systematic risk (β).
- **Monte Carlo Simulation:** Generates random portfolios to approximate efficient frontiers.
- **Value-at-Risk (VaR):** Calculates the 95% and 99% VaR as downside risk measures.
- **Black–Litterman:** Synthesises ML model predictions as views with equilibrium re- turns.

3.5.2 Machine Learning Models

Three algorithms were used:

- **Linear Regression:** A simple benchmark model.
- **Random Forest:** An ensemble method that handles non-linearities and reduces over- fitting ([Breiman, 2001](#)).

- **Gradient Boosting:** A sequential ensemble (e.g., XGBoost) effective for financial prediction ([Friedman, 2001](#)).

These models provide daily asset return forecasts to be used in portfolio optimisation.

3.6 Optimisation and Integration

Portfolios were constructed using both the traditional mean–variance framework and ML-enhanced forecasts. A hybrid approach incorporated ML predictions as views in the Black–Litterman model, balancing predictive signals with equilibrium priors.

3.7 Performance Evaluation

Portfolio performance was measured by predictive and financial metrics. **Predictive metrics** include mean squared error (MSE) and R^2 of return forecasts. **Financial metrics** include Sharpe ratio, alpha, cumulative return, and VaR. Three strategy classes (classical, ML, hybrid) were compared across these metrics.

3.8 Ethical Considerations

The research uses public financial data and adheres to transparency and reproducibility principles. The platform serves as an analytic tool, not financial advice, and emphasizes responsible use of predictive algorithms in decision-making.

Chapter 4

Implementation

4.1 Introduction

Implementation translates theoretical designs into a working system. For this portfolio optimisation platform, implementation means a modular process ensuring transparency and interpretability. The system was built in layers: (1) data collection/preprocessing, (2) exploratory analysis, (3) model building/integration, (4) optimisation, and (5) dashboard development.

4.2 System Architecture

The platform follows a multi-layer pipeline reminiscent of financial analytics systems. Each layer handles distinct processing:

4.2.1 Layer 1 – Data Layer

Acquires and stores raw financial data (equities, ETFs, indices). The design is flexible to accept multiple data frequencies (daily, hourly).

4.2.2 Layer 2 – Analytical Layer

Computes descriptive statistics, visualizes return distributions, and identifies stylized facts (volatility clustering, heavy tails). This guides model selection and tuning.

4.2.3 Layer 3 – Model Layer

Integrates traditional models (CAPM, Monte Carlo, VaR, Black–Litterman) and ML models (Linear Regression, Random Forest, Gradient Boosting). It outputs expected returns, covariance matrices, and predictive signals.

4.2.4 Layer 4 – Optimization Layer

Constructs portfolios using inputs from the model layer. It supports three frameworks: classical mean–variance optimisation, ML-only optimisation, and the hybrid Black– Litterman approach integrating ML views.

4.2.5 Layer 5 – Presentation Layer

Communicates results via an interactive dashboard. It includes visualizations (efficient frontiers, allocation pie charts, performance tables) and allows user interaction.

This layered architecture ensures the platform is not a black box but a transparent tool where each stage can be examined individually.

4.3 Data Layer

4.3.1 Data Sources

Data was collected from publicly available U.S. equity and ETF price histories over two years. Selected assets span technology, healthcare, finance, energy, and consumer sectors. ETFs represent index tracking instruments used by investors.

4.3.2 Handling Missing Data

Missing data occurs from holidays or disruptions. Forward-fill (propagating last known value) and backward-fill were applied to maintain the time-series continuity without bias.

4.3.3 Data Transformation

Log returns were computed:

$$R_t = \ln \frac{P_t}{P_{t-1}},$$

here P_t is price at time t . Log returns, being time-additive, are used instead of simple

returns to ease multi-period analysis.

4.3.4 Normalisation and Scaling

Prices were normalized using z-scores:

$$z = \frac{x - \mu}{\sigma},$$

to standardize returns (mean 0, standard deviation 1). This makes features suitable for ML algorithms sensitive to scale.

4.3.5 Outlier Treatment

Extreme returns beyond 1.5 times the interquartile range were flagged. Manual inspection ensured that anomalies (e.g., due to stock splits or data errors) did not distort model training.

4.3.6 Smoothing

High-frequency noise was reduced via moving averages and exponential smoothing, improving stability for predictive models.

4.4 Analytical Layer

This layer translates cleaned data into insights through statistical diagnostics.

4.4.1 Descriptive Statistics

For each asset, mean returns, standard deviation, skewness, and kurtosis were computed. Many assets showed high kurtosis (fat tails) and skewed distributions, contradicting the normality assumed by CAPM.

4.4.2 Correlation Structure

Correlation matrices revealed interdependencies. For example, technology stocks showed

high positive correlation, reducing diversification benefits. ETFs had lower correlations with individual stocks and thus improve diversification.

4.4.3 Stationarity Tests

The Augmented Dickey–Fuller (ADF) test confirmed that raw prices are non-stationary, whereas log returns are stationary at 1% level, validating return-based modelling.

4.4.4 Volatility Clustering

ACF of squared returns indicated significant persistence (volatility clustering). This underscores the inadequacy of models assuming constant volatility and motivates using models (e.g., GARCH) that account for time-varying risk.

4.4.5 Factor Sensitivity Analysis

Regression of asset returns on benchmarks and macro factors was performed to estimate factor sensitivities (e.g., market beta). This provides context for CAPM beta estimates and informs model comparisons.

4.5 Model Layer (Part 1 – Financial Models)

The model layer combines classical finance models with ML. This subsection covers the classical models:

4.5.1 Capital Asset Pricing Model (CAPM)

The CAPM defines expected return as:

$$E(R_i) = R_f + \beta_i(E(R_m) - R_f),$$

where R_f is the risk-free rate, β_i measures systematic risk, and $E(R_m)$ is expected market return. While CAPM provides a baseline, its linear assumption often fails in turbulent

markets.

4.5.2 Monte Carlo Simulation

Monte Carlo methods simulated thousands of random portfolios (random weights subject to constraints) to approximate the efficient frontier. Each portfolio's expected return and volatility were computed, illustrating how estimation error affects the frontier and enabling stress testing.

4.5.3 Value-at-Risk (VaR)

VaR quantified potential losses under normal market conditions at given confidence levels. Both parametric and historical VaR were calculated. While intuitive, VaR misses extreme tail risks, so Conditional VaR (CVaR) was also used to capture those.

4.5.4 Black–Litterman Model

The Black–Litterman framework blends equilibrium returns with subjective views:

$$E(R) = [(\tau \Sigma)^{-1} + P^T \Omega^{-1} P]^{-1} [(\tau \Sigma)^{-1} \Pi + P^T \Omega^{-1} Q].$$

Here Π are equilibrium returns, P is a matrix of view portfolios, Q are expected returns of those views, Ω is view uncertainty, and τ scales prior uncertainty. Incorporating ML forecasts as views stabilises weights while integrating predictive signals.

4.6 Optimization Layer

The optimization layer translates model outputs into portfolio allocations. Portfolio optimisation has been at the core of modern finance since [Markowitz \(1952\)](#), but naive reliance on theory can lead to unstable and unrealistic allocations ([Michaud, 1989](#)). We integrate three approaches:

1. **Classical Mean–Variance Optimisation:** Minimises portfolio variance subject to a target return.

2. **ML-driven Optimisation:** Uses ML-predicted returns to drive allocation.
3. **Hybrid Optimisation:** Incorporates ML forecasts as views in the Black–Litterman model.

4.6.1 Classical Mean–Variance Optimisation

At the foundation is the mean–variance problem:

$$\min_w w^T \Sigma w \quad \text{s.t.} \quad w^T \mu = R^*, \quad \sum_i w_i = 1,$$

where w are portfolio weights, Σ is covariance matrix, μ is expected returns, and R^* is the target return. The optimisation yields the efficient frontier. Limitations include instability to estimation errors (?), extreme concentrations, and the normal returns assumption (contradicted by fat tails (Cont, 2001)). Common constraints mitigate issues: long-only ($w_i \geq 0$), weight caps (e.g. $w_i \leq 0.3$), and turnover limits.

4.6.2 Machine Learning–Driven Optimisation

ML offers forecasting of returns using complex patterns (Gu et al., 2020). In this framework, ML-predicted returns (μ^{ML}) replace traditional estimates:

$$\min_w w^T \Sigma w - \lambda w^T \mu^{ML},$$

where λ balances risk and expected return. ML models like Random Forests (Breiman, 2001) and Gradient Boosting (Friedman, 2001) can capture non-linearity. This approach improves adaptability but often yields high turnover and requires careful regularisation.

4.6.3 Hybrid Optimisation (Black–Litterman with ML Views)

The hybrid framework feeds ML forecasts into Black–Litterman (Black and Litterman, 1992). ML predictions form part of the view vector Q , producing posterior returns via the BL formula above. Advantages include balancing ML’s predictive power with classical

stability and reducing estimation sensitivity.

4.6.4 Stress Testing and Robust Optimisation

Portfolios were stress-tested using Monte Carlo scenarios (market crashes, volatility spikes) (?). Robust optimisation was also applied by modeling uncertainty in inputs (e.g., using confidence intervals instead of point estimates) (Fabozzi et al., 2006). These techniques ensure portfolios perform well under adverse conditions.

4.7 Presentation Layer

The presentation layer converts optimisation outputs into user-friendly insights. Interactive dashboards were created following principles of clarity, interactivity, and modularity (Few, 2012). Key features include:

4.7.1 Dashboard Design Principles

- **Clarity:** Visualizations (charts, tables) minimize cognitive load.
- **Interactivity:** Users can adjust models, assumptions, and input new data.
- **Modularity:** Each visualization (efficient frontier, risk metrics, etc.) is self-contained for future expansion.

4.7.2 Core Visualisations

- **Efficient Frontier:** Plots trade-off between risk and return for classical, ML, and hybrid strategies.
- **Allocation Pie Charts:** Show capital distribution across assets, highlighting concentration risks.
- **Cumulative Return Graphs:** Display growth of \$1 invested under each strategy.
- **Risk Metric Tables:** Summarize metrics (Sharpe, Sortino, max drawdown, VaR) for each strategy.

- **Correlation Heatmaps:** Visualize asset correlations to guide diversification.

4.7.3 Interactivity Features

Users can upload custom datasets, switch between scenarios (bull/bear), hover for precise values, and dynamically filter assets or time periods. These features turn the platform into a practical decision-support system.

4.8 Technical Challenges and Solutions

Several challenges arose in development:

4.8.1 Data Quality Issues

Financial data often has errors or missing values ([Tsay, 2010](#)). *Solution:* Apply forward/backward filling, manual inspection, and outlier detection (?).

4.8.2 Model Instability

ML models can produce volatile forecasts leading to extreme allocations. *Solution:* Embed ML views in Black–Litterman to regularise predictions ([Meucci, 2010](#)).

4.8.3 Overfitting Risk

ML models may overfit noisy financial data. *Solution:* Use cross-validation, conservative tuning, and ensemble methods to mitigate overfitting ([Gu et al., 2020](#)).

4.8.4 Computational Demands

Monte Carlo simulations and ensemble ML are computationally expensive. *Solution:* Utilize parallel computing and sampling techniques to reduce runtimes (?).

4.8.5 User Experience

Early dashboards can overwhelm users. *Solution:* Iterative design emphasizing essential outputs, with advanced options accessible on demand ([Few, 2012](#)).

4.9 Summary

The implementation phase built a functional optimisation platform. A layered architecture integrated three frameworks: classical mean–variance optimisation ([Markowitz, 1952](#)), ML-driven optimisation ([Breiman, 2001](#); [Friedman, 2001](#)), and hybrid Black–Litterman ([Black and Litterman, 1992](#); [Meucci, 2010](#)). The hybrid framework emerged as the most balanced, combining predictive adaptability with stability. The presentation layer ensured clear communication of results through interactive visualizations.

Chapter 5

Results and Analysis

5.1 Introduction

This chapter presents the outcomes of implementing the portfolio optimisation platform. It compares empirical results across three approaches: (1) classical mean–variance, (2) ML-driven, and (3) hybrid (ML forecasts integrated into Black–Litterman). Analysis is conducted from three perspectives:

- **Distribution Analysis:** Statistical characteristics of returns.
- **Comparative Model Analysis:** Performance of classical, ML, and hybrid portfolios in terms of allocation, efficiency, and stability.
- **Performance Metrics:** Financial and predictive metrics (Sharpe, Sortino, VaR, CVaR, alpha, turnover).

The aim is to determine which framework best balances accuracy, robustness, and interpretability, addressing the research questions.

5.2 Distribution Analysis

Before building portfolios, we examined return distributions to validate classical assumptions.

5.2.1 Return Distributions

Asset return distributions significantly deviated from normality. Most assets failed normality tests, exhibiting fat tails and leptokurtosis (Cont, 2001). For example, technology stocks had kurtosis above 6 (Gaussian=3). Some sectors showed skew

(healthcare positive, energy negative), implying higher downside risk. These findings violate assumptions of CAPM and mean–variance (Sharpe, 1964; Markowitz, 1952), justifying robust risk measures (VaR, CVaR) and ML models.

5.2.2 Volatility Clustering

Squared-return autocorrelations confirmed volatility clustering (Engle, 1982). Spikes coincided with market downturns (e.g., energy sector corrections). This underscores inadequacy of homoskedastic models and supports hybrid/ML methods that adapt to changing volatility.

5.2.3 Correlations and Diversification Potential

Correlation matrices revealed sectoral dependencies. Banking and finance equities had average correlations $\rho \approx 0.75$. Technology firms had correlations $\rho \approx 0.80$. ETFs showed weaker correlations ($\rho \approx 0.4$), confirming their diversification role. These results demonstrate that while diversification is achievable, certain sectors remain highly correlated. Hybrid optimization frameworks can manage these correlations more effectively by blending predictive and equilibrium views.

5.2.4 Stationarity Tests

Using the ADF test (Dickey and Fuller, 1979), price series were non-stationary, but log returns were stationary at the 1% level. This validates using returns in modelling and cautions against reliance on raw prices.

5.3 Comparative Analysis of Models

The three frameworks were evaluated on in-sample and out-of-sample data.

5.3.1 Classical Framework Results

Classical portfolios (mean–variance with CAPM and VaR) adhered to theory. The efficient frontier was convex, reflecting the risk–return trade-off. However, it was highly sensitive to

inputs. Allocations concentrated on few assets, echoing [Michaud \(1989\)](#)'s critique of "error maximization". Out-of-sample Sharpe ratios averaged 0.85 (Sortino 1.10). These were respectable but volatile; classical portfolios underperformed in turbulent markets. Overall, the classical model demonstrated theoretical consistency but suffered from fragility and overconcentration.

5.3.2 Machine Learning Framework Results

ML models forecast returns to inform optimisation. ML methods (Random Forest ([Breiman, 2001](#)), Gradient Boosting ([Friedman, 2001](#))) reduced predictive MSE by 15–20% relative to CAPM forecasts. Allocations became dynamic: tech stocks gained weight in growth phases and shed in downturns, but this reactivity led to high turnover. Out-of-sample Sharpe ratios averaged 1.05 (Sortino 1.35), outperforming classical models but with greater instability.

Turnover nearly doubled, raising transaction cost concerns. In summary, the ML framework improved predictive strength but lacked the stability required for long-term strategies.

5.3.3 Hybrid Framework Results

The hybrid approach integrated ML forecasts into the Black–Litterman model ([Black and Litterman, 1992](#); [Meucci, 2010](#)). Its efficient frontier was smoother and less extreme, confirming Black–Litterman's ability to regularise ML forecasts and avoid corner solutions. Allocations were well-diversified: while ML forecasts overweighted certain stocks (e.g. tech), the Black–Litterman structure moderated these weights. Out-of-sample Sharpe ratios reached

1.20 (Sortino 1.55), outperforming the other approaches. Turnover was significantly lower than in ML portfolios, indicating greater stability. The hybrid framework emerged as the most balanced approach, reconciling predictive accuracy with robustness.

5.4 Performance Metrics

5.4.1 Sharpe Ratio and Alpha

The Sharpe ratio measures excess return per unit risk:

$$S = \frac{E(R_p - R_f)}{\sigma_p},$$

where R_p is portfolio return, R_f risk-free rate, and σ_p portfolio volatility. Classical portfolios averaged Sharpe 0.85; ML portfolios 1.05; hybrid portfolios 1.20. Jensen's alpha ([Jensen, 1968](#)) confirmed that hybrid portfolios generated significant positive alpha, outperforming benchmarks after adjusting for risk.

5.4.2 Sortino Ratio and Treynor Ratio

The Sortino ratio penalizes only downside volatility. It was highest for hybrid portfolios (1.55), then ML (1.35), and classical (1.10). The Treynor ratio (return per unit β) also favored hybrid portfolios, which maintained superior performance while diversifying away unsystematic risk.

5.4.3 Value-at-Risk and CVaR

At 95% confidence, VaR indicated hybrid portfolios had the lowest potential daily losses (-1.8%) versus classical (-2.2%) and ML (-2.5%) portfolios. Conditional VaR (CVaR) reinforced this: hybrid portfolios experienced less severe tail losses, confirming their robustness under adverse conditions ([Jorion, 2006](#)).

5.4.4 Cumulative Returns and Stability

Over the two-year test horizon, hybrid portfolios consistently outperformed. While ML portfolios occasionally produced higher short-term returns, they also suffered deeper drawdowns. Classical portfolios, though more stable, lagged in long-term growth. The hybrid approach combined high cumulative returns with reduced volatility, validating its superiority.

5.5 Visualization and Dashboard Outputs

One critical contribution is the interactive dashboard bridging technical models and decision-making.

5.5.1 Efficient Frontier Visualization

The platform plotted classical, ML-driven, and hybrid efficient frontiers on one graph. Classical frontiers were narrow and sensitive; ML frontiers extended return potential but were irregular; hybrid frontiers were smooth and dominating. This visualization helps users grasp comparative advantages.

5.5.2 Allocation Dynamics

Interactive charts showed how allocations shift over time/scenarios. Classical portfolios had abrupt changes with small input shifts. ML portfolios rebalanced rapidly (high turnover). Hybrid portfolios adjusted more gradually, striking a balance between adaptability and stability. For example, during a simulated downturn in the energy sector, classical models over-divested, ML models oscillated sharply, and the hybrid model moderated reallocation, retaining some exposure while diversifying into defensive assets.

5.5.3 Performance Tables

The dashboard displayed tables of key metrics (Sharpe, Sortino, alpha, VaR, CVaR, turnover) for each strategy. Color-coded indicators (green for superior, red for inferior) highlighted performance differences, and significance markers distinguished robust results.

5.5.4 Scenario Testing

Users could stress-test portfolios under bullish, bearish, and neutral market regimes. Hybrid portfolios consistently showed resilience: in a simulated COVID-19 crash scenario, hybrid portfolios had 25% lower drawdowns than ML-only portfolios, preserving more capital while still capturing upside potential.

5.5.5 User Interaction Features

The dashboard allows users to upload new datasets, hover over graphs for precise values, and toggle between models for side-by-side comparison. This interactivity aligns with the

principle that effective decision-support systems must compute accurate results and present them in a user-friendly manner (Few, 2012).

5.6 Discussion of Findings

The results of this study reveal several important insights that contribute to both academic discourse and practical portfolio management.

5.6.1 Comparative Strengths and Weaknesses

- **Classical Models:** Strengths include theoretical elegance, simplicity, and interpretability. Weaknesses are fragility to input estimation errors (Michaud, 1989), overconcentration of allocations, and inability to accommodate non-normal return distributions.
- **Machine Learning Models:** Advantages include increased predictive strength and the ability to capture nonlinearities and interactions (Gu et al., 2020). Weaknesses are instability, high turnover, and low interpretability.
- **Hybrid Models:** Strengths combine the predictive asset of ML with the stability and diversification of Black–Litterman (Black and Litterman, 1992). Weaknesses include added computational complexity and the need for careful parameter tuning.

5.6.2 Theoretical Contributions

This study empirically validates that integrating ML forecasts into finance theory improves portfolios. It quantifies the benefits—showing improved Sharpe and reduced CVaR when ML-generated views are fused with equilibrium priors (Black and Litterman, 1992; Meucci, 2010). The results also clarify that superior return estimates alone do not guarantee better portfolios unless covariance estimation is concurrently improved (Ledoit and Wolf, 2004). Operationalising view uncertainty calibration (Ω) when views come from statistical models bridges a gap between abstract Bayesian formulation and practical implementation.

5.6.3 Practical Contributions

The platform offers a replicable end-to-end template (data ingestion → preprocessing → forecasting → posterior formation → optimisation → visualization) for institutional adaptation. By maintaining economic interpretability and explicit uncertainty quantification, the hybrid approach eases communication with risk committees and regulators. The combination of shrinkage covariance, Black–Litterman fusion, and conservative ML tuning creates portfolios that are more robust to turnover and market swings, making them operationally attractive for long-term mandates.

5.6.4 Alignment with Research Questions

The findings address the research questions:

- **RQ1:** The hybrid model consistently outperformed others in Sharpe ratio, alpha, and downside protection.
- **RQ2:** Classical models underperformed due to sensitivity to inputs; ML models underperformed due to instability and overfitting.
- **RQ3:** Hybridisation is the most promising solution; future work could explore reinforcement learning or deep learning for dynamic allocation.

5.6.5 Limitations

Limitations of this study include:

- **Scope of asset classes:** The analysis focused on equities and ETFs. Fixed income, derivatives, and alternative assets might behave differently.
- **Transaction cost realism:** Turnover was measured, but real market impact and costs were not fully modeled.
- **Model risk governance:** ML models require ongoing monitoring and retraining; the platform prototype lacks full governance protocols.

- Computational scalability: Ensemble ML and Monte Carlo are intensive. Scaling to global high-frequency data needs distributed computing.
- Backtest bias: Despite careful validation, any backtest can be affected by data-snooping. Replication and out-of-sample testing remain advisable (Lo, 2004).

5.6.6 Future Research Directions

Potential extensions include:

- Incorporating reinforcement learning and stochastic control to explicitly optimize multi-period objectives (Fischer and Krauss, 2018; Sutton and Barto, 2018).
- Developing probabilistic ML forecasts (Bayesian methods) to supply full distributions to Black–Litterman.
- Using regime-switching models or hidden Markov models to detect structural breaks (Hamilton, 1994).
- Applying explainable AI (e.g. SHAP, LIME) to interpret which features drive ML forecasts, enhancing transparency (Lundberg and Lee, 2017).
- Combining portfolios with options or hedging to manage tail risk (Cont, 2001).
- Integrating ESG factors or alternative data (news, sentiment) into forecasting and allocation.
- Exploring alternative allocation schemes (HRP, clustering) for robustness (?).
- Implementing online learning and drift detection to continuously update models.

5.6.7 Practical Recommendations

For practitioners:

- **Data Engineering:** Prioritize clean data pipelines and adjust for survivorship bias (Tsay, 2010).

- **Covariance Regularisation:** Use shrinkage or factor models ([Ledoit and Wolf, 2004](#)).
- **View Calibration:** Carefully set view confidence when using ML-based forecasts.
- **Cost Modeling:** Incorporate realistic transaction cost models ([Almgren and Chriss, 2000](#)).
- **Continuous Validation:** Monitor ML models for drift and retune regularly.

Chapter 6

Discussion

6.1 Interpretation of Results

The empirical findings presented in Chapter 5 indicate a clear pattern: pure classical mean–variance optimisation, while theoretically elegant, is fragile in realistic market conditions; machine-learning (ML) driven forecasts improve predictive accuracy but tend to produce unstable allocations; and a hybrid approach that incorporates ML forecasts as “views” within the Black–Litterman framework produces the most consistent risk-adjusted outcomes. This section interprets those results in light of statistical properties of financial data, optimisation theory and practical constraints.

6.1.1 Estimation Error and the Bias–Variance Trade-off

Classical mean–variance optimisation is extremely sensitive to estimation errors in expected returns and the covariance matrix because the optimisation amplifies small differences in inputs into large differences in weights (Michaud, 1989). In practice, estimates of μ (expected returns) have high variance and therefore the optimiser tends to overfit these noisy estimates. ML forecasters reduce bias in return estimates by leveraging non-linear relationships and more features (Gu et al., 2020), thereby improving predictive accuracy (lower MSE). However, this improved accuracy can come at the cost of higher prediction variance and instability in allocation time series, especially when models react to transient patterns in noisy financial data.

6.1.2 Regularisation through Black–Litterman

The Black–Litterman framework introduces an implicit form of regularisation: by

blending equilibrium (market-implied) returns with investor (or model) views and explicitly modelling view uncertainty, it shrinks extreme ML forecasts back towards a prior, reducing overreaction and concentration risk (Black and Litterman, 1992; Meucci, 2010). Mathematically, the posterior mean:

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

shows how the relative magnitudes of τ and Ω control the weight given to the prior Π versus the views Q . When ML forecasts are used as views, appropriate choice of Ω (view uncertainty) prevents ML noise from dominating the posterior and producing extreme allocations. The empirical superiority of the hybrid approach in this research is consistent with that theoretical mechanism.

6.1.3 Covariance Estimation and Portfolio Stability

Covariance estimation is another great source of instability. The empirical returns take fat tail forms, volatility is concentrated (Cont, 2001; Engle, 1982), and small sample covariations are noisy. Factor based covariance models and shrinkage estimators (Ledoit and Wolf, 2004) reduce the error in estimation by placing structure or combination of sample covariances and a target matrix. Results using portfolio methods of shrinkage or robust covariance on our data produced much more consistent efficient frontiers and lower turnover. So, to get stable allocations, using better mean estimation (ML) and effective covariance estimation and Bayesian posterior construction (Black–Litterman) is possible.

6.1.4 Risk-adjusted Performance vs. Turnover and Capacity

Although in some windows the ML portfolios achieved higher returns in shorter time, it suffered more than the turnover rising the issues of transaction cost, market effects and capacity (Almgren and Chriss, 2000). The hybrid portfolios fared better on Sharpe and Sortino ratio but had a moderate level of turnover that implies that it can be applied to real world. Realistically speaking, the margin of Sharpe of highly active ML strategies and their marginal gain may be annihilated by the trading frictions—the hybrid strategy is a more practicable escape.

6.1.5 Explainability and Economic Intuition

Pure ML models often lack transparent economic interpretation, whereas classical models have clear economic narratives (risk premia, diversification). The hybrid approach preserves interpretability by anchoring forecasts to equilibrium priors and expressing view confidence explicitly—facilitating economic scrutiny and governance (Meucci, 2010; Gu et al., 2020).

6.2 Strengths and Weaknesses

This section analyses methodological strengths and weaknesses in greater depth and links them to the strategic choices made during the project.

6.2.1 Strengths

- **Holistic framework comparison.** The study's head-to-head comparison of classical, ML, and hybrid frameworks under identical data, preprocessing and evaluation pipelines provides strong internal validity. It isolates the modelling differences from data pipeline effects, allowing clear attribution of performance differentials.
- **Robust evaluation metrics.** Use of multiple dimensions—predictive metrics (MSE, R^2), risk-adjusted metrics (Sharpe, Sortino), tail measures (VaR, CVaR), and operational metrics (turnover)—gives a fuller view of strategy performance than single-metric comparisons. This avoids misleading conclusions that could arise from optimising only for returns.
- **Regularisation and risk control.** Practical constraints (long-only, weight caps, turnover limits) and use of shrinkage covariance and Bayesian fusion (Black–Litterman) improved real-world plausibility of portfolio outputs (Ledoit and Wolf, 2004; Black and Litterman, 1992).
- **Stress testing and scenario analysis.** Systematic stress tests (market crash, rising rates, sector shocks) confirmed that hybrid portfolios reduce tail exposure—an essential

validation for stakeholders concerned with drawdown risk and regulatory scrutiny.

- **Decision-support design.** Presenting outputs through an interactive dashboard that supports scenario analysis, model selection and uploadable datasets increases the practical utility and adoption potential for non-technical users ([Few, 2012](#)).

6.2.2 Weaknesses

- **Scope of asset classes.** The analysis focused on equities and ETFs. Fixed income, derivatives, and alternative assets (commodities, real estate) present different return distributions and liquidity profiles; conclusions may not generalise without further testing.
- **Transaction cost realism.** Although turnover was measured, transaction cost and market-impact modelling were simplified. For high-frequency or large-scale strategies, explicit microstructure modelling (slippage, liquidity curves) is required to assess net performance ([Almgren and Chriss, 2000](#)).
- **Model risk and lifecycle.** ML models require ongoing monitoring and retraining; the platform demonstrates the pipeline but does not yet embody full model-risk governance protocols (documentation, validation, model retirement, monitoring for drift).
- **Computational scalability.** Ensemble ML methods and Monte Carlo simulations are computationally intensive. Scaling to global, high frequency datasets would require significant engineering (distributed computing, GPU acceleration) and cost.
- **Potential overfitting in backtests.** Despite cross-validation and walk-forward testing, backtesting remains vulnerable to data snooping and look-ahead biases if not rigorously controlled ([Lo, 2004](#)). The experimental design minimised these risks, but external replication remains advisable.

6.3 Theoretical and Practical Contributions

6.3.1 Theoretical Contributions

- **Empirical validation of hybridisation.** Prior literature argued conceptually for combining ML with finance theory; this study empirically quantifies the benefits—showing improved Sharpe, reduced CVaR and moderated turnover when ML views are integrated in a Bayesian framework ([Black and Litterman, 1992](#); [Meucci, 2010](#); [Gu et al., 2020](#)).
- **Model interaction insights.** The results clarify the interaction between mean-estimate improvements and covariance estimation: superior means alone do not guarantee better portfolios unless covariance noise is concurrently reduced ([Ledoit and Wolf, 2004](#)). Therefore, research on joint mean–covariance estimation remains critical.
- **Operationalising view uncertainty.** The study demonstrates practical choices for view confidence calibration (Ω) when views derive from statistical models—bridging a gap between abstract Bayesian formulation and implementable parameterisation.

6.3.2 Practical Contributions

- **A replicable decision-support template.** The platform represents an end-to-end blueprint (data ingestion → preprocessing → forecasting → posterior formation → optimisation → visualisation) that practitioners can adapt to institutional settings.
- **Risk governance-friendly outputs.** By maintaining economic interpretability and explicit uncertainty quantification, the hybrid framework eases communication with risk committees and regulators—important for buy-side adoption.
- **Improved robustness for deployment.** The combination of shrinkage covariance, Black–Litterman fusion and conservative ML hyperparameterisation creates portfolios that are more robust to model turnover and boom-bust cycles, making them operationally attractive for longer horizon mandates.

6.4 Limitations and Future Research Directions

6.4.1 Empirical and Methodological Limitations

- **Asset universe breadth.** Extending the asset universe to fixed income, derivatives, and alternative investments would test the hybrid approach under different return and liquidity regimes.
- **High-frequency data.** The study used daily prices; applying ML and hybrid methods to intraday data introduces microstructure effects and necessitates latency-aware systems.
- **Transaction costs and slippage.** Future work should integrate realistic cost curves and market impact models to measure net strategy performance under varying trade sizes ([Almgren and Chriss, 2000](#)).
- **Model risk governance.** Implementing lifecycle processes—model documentation, validation, monitoring, and automated retraining thresholds—would move the platform from proof-of-concept to production readiness.
- **Survivorship and look-ahead bias.** Ensuring dataset integrity by removing survivorship bias and rigorous time-aware feature engineering should be standardised in replication attempts.

6.4.2 Research Directions

- **Dynamic and sequential decision models.** Investigate reinforcement learning and stochastic control methods that explicitly optimise multi-period objectives and transaction costs ([Sutton and Barto, 2018](#)). These could capture the value of timing and dynamic rebalancing.
- **Probabilistic forecasting and Bayesian ML.** Move from point forecasts to full predictive distributions (quantile regression, Bayesian neural nets), enabling principled integration into Black–Litterman with probabilistic views and more nuanced risk measures.
- **Regime-aware modelling.** Implement regime-switching models or hidden Markov models to detect structural breaks and adapt view confidence dynamically ([Hamilton, 1994](#)).

- **Explainability and model auditing.** Apply explainable AI techniques (SHAP, LIME) to interpret ML features driving views, increasing transparency and aiding model validation ([Lundberg and Lee, 2017](#)).
- **Tail-risk hedging.** Combine hybrid portfolios with options or dynamic hedging strategies to explicitly manage extreme loss scenarios ([Cont, 2001](#)).
- **ESG and non-financial factors.** Integrate ESG scores, alternative data (news, sentiment), and macro indicators to evaluate how non-price inputs affect forecasts and allocations.
- **Hierarchical risk parity and clustering.** Explore alternative allocation schemes (HRP, clustering-based approaches) that use the covariance structure differently from mean-variance optimisation, possibly offering additional robustness (?).
- **Online learning and model drift detection.** Implement online learning frameworks and statistical tests to detect when model retraining is required, preventing stale predictors from degrading performance.

6.5 Practical Recommendations for Deployment

For practitioners intending to implement the hybrid framework, the following actionable recommendations are suggested:

- **Data engineering first.** Prioritise high-quality data pipelines: ensure clean historical records, corporate action adjustments, and prevention of survivorship bias ([Tsay, 2010](#)).
- **Covariance regularisation.** Use shrinkage ([Ledoit and Wolf, 2004](#)) or factor models to stabilise covariance estimates before optimisation.
- **Conservative view confidence.** When converting ML forecasts into Black–Litterman views, set Ω conservatively (i.e., allow meaningful prior weight) until forecast reliability is confirmed through extensive out-of-sample tests.
- **Walk-forward validation.** Use rolling windows and walk-forward backtests rather

than static in-sample tests to assess real-world performance and turnover implications.

- **Transaction cost modelling.** Incorporate explicit friction models into objective functions (penalise turnover, enforce round-trip cost budgets) to produce implementable strategies.
- **Explainability and documentation.** Apply feature-importance and local explanation tools (e.g., SHAP) to every forecast used as a view; maintain comprehensive model documentation for audit and governance.
- **Monitoring and KPIs.** Define clear operational KPIs (prediction error drift, turnover, realised vs expected return) and automated alerts for model degradation.
- **Capacity testing.** Evaluate how strategy performance changes with AUM scale—market impact can nullify theoretical gains for large allocations.

6.6 Concluding Remarks

Chapter 6 synthesises the empirical results into a coherent set of theoretical insights, practical lessons and a roadmap for future research. The key conclusion is that a hybrid approach—anchoring ML forecasts within a Bayesian, equilibrium-consistent framework and combining this with robust covariance estimation—delivers materially better and more implementable portfolios than either classically driven or purely ML-driven methods alone. This conclusion is robust across multiple performance metrics, stress scenarios and evaluation horizons.

The chapter also highlighted the operational and governance realities that accompany deploying such hybrid systems: quality data pipelines, transaction cost awareness, model risk controls and explainability are not optional—they are prerequisites for responsible, production deployment. Future research should pursue dynamic, probabilistic and regime-aware extensions, and practitioners should follow conservative integration and rigorous validation practices.

Chapter 7

Conclusion and Future Work

7.1 Summary of Findings

This dissertation set out to investigate whether hybridising classical portfolio optimisation methods with machine learning (ML) forecasting could yield more robust and effective portfolio allocation strategies. Beginning with the limitations of classical mean–variance optimisation (Markowitz, 1952) and the overfitting tendencies of standalone ML forecasts (Gu et al., 2020), the research explored a blended approach through the Black–Litterman framework (Black and Litterman, 1992).

Key findings include:

- **Return distributions deviate from Gaussian assumptions.** Empirical analysis showed fat tails, volatility clustering, and strong correlations across assets (Cont, 2001; Engle, 1982). These stylised facts undermine the assumptions underpinning models such as CAPM and highlight the need for more robust frameworks.
- **Classical models are fragile.** While mean–variance optimisation provides an elegant theoretical structure, results revealed high sensitivity to small estimation errors, often producing concentrated and unstable allocations (Michaud, 1989).
- **Machine learning improves forecasts but lacks stability.** ML methods reduced prediction errors and occasionally boosted short-term performance, but they generated high turnover, frequent rebalancing, and occasional extreme allocations—making them difficult to implement in practice.
- **Hybrid models outperform across metrics.** Integrating ML forecasts as probabilistic “views” into Black–Litterman produced portfolios with the highest Sharpe (1.20) and Sortino (1.55) ratios, lowest Value-at-Risk, and reduced turnover relative to ML-only allocations. Hybrid models consistently dominated

in both stable and stress-tested scenarios.

- **Dashboard visualisation enhances interpretability.** By translating complex model outputs into interactive frontiers, allocation charts, and stress-test simulations, the platform ensures accessibility for retail investors, institutions, and regulators alike (Few, 2012).

In sum, the research demonstrated that a hybrid integration of machine learning and financial theory is the most effective framework for modern portfolio optimization.

7.2 Contributions to Knowledge

This dissertation contributes to both academic research and practical portfolio management.

7.2.1 Theoretical Contributions

- **Empirical evidence for hybridisation.** While the idea of blending ML forecasts with financial theory has been proposed conceptually, this work provides one of the first empirical validations across multiple metrics and scenarios.
- **Extension of the Black–Litterman model.** By calibrating ML forecasts as investor “views,” this research extends the usability of the Black–Litterman framework. The explicit modelling of uncertainty (Ω) around ML predictions bridges the gap between abstract Bayesian formalism and applied asset management (Meucci, 2010).
- **Clarification of input sensitivity.** Results reinforce the longstanding critique that mean–variance optimisation is prone to input instability (Michaud, 1989). The study contributes by showing that ML forecasts alone cannot solve this; Bayesian shrinkage is required to stabilise allocations.

7.2.2 Practical Contributions

- **Blueprint for decision-support systems.** The end-to-end platform (data → pre-

processing → modelling → optimisation → dashboard) offers practitioners a replicable template for building hybrid portfolio systems.

- **Risk governance alignment.** The framework enhances explainability and risk management, allowing managers to demonstrate model transparency, stress testing, and probabilistic reasoning—key requirements under Basel III and other regulatory standards.
- **Bridging institutional and retail use cases.** The dashboard lowers barriers to adoption by offering intuitive interfaces, while still accommodating institutional-grade robustness and stress testing.

7.2.3 Methodological Contributions

The study demonstrated how to calibrate ML view uncertainty in Black–Litterman, balancing overconfidence and under-utilisation of forecasts. It integrated shrinkage covariance estimation ([Ledoit and Wolf, 2004](#)) with hybrid forecasting, providing improved portfolio stability. Traditional financial performance metrics were combined with predictive accuracy and turnover, yielding a more holistic evaluation standard.

7.3 Final Remarks

The evolution of portfolio optimisation reflects a larger story about finance: a discipline balancing elegant theory with noisy, unpredictable markets. This research confirms that neither classical finance nor machine learning alone offers a complete solution. The most promising path forward lies in hybridisation—integrating the interpretability and stability of financial theory with the predictive power of machine learning.

The implications are significant:

- **To scholars:** Interdisciplinary synthesis is found to be important. Econometrics, ML and Bayesian reasoning ought to be incorporated in the future models.
- **To practitioners:** The hybrid framework can provide strategies that can be adopted as well as being more shock-resistant besides being more profitable.

- **To policy-makers and regulators:** The paper shows the importance of transparency, stress testing, and risk-sensitive application of sophisticated models.

The portfolio optimisation structures must be altered because of the ongoing changes in the financial markets, which are being threatened by technological disruption, geopolitical risk, and sustainability concerns. One of such adaptation is offered by this dissertation: a hybrid methodology that is predictive and stable, innovative and traditional, complex and interpretable.

The primary contribution of this work is the demonstration that machine learning is not going to replace financial theory—it will enhance it. Predictive intelligence facilitates the ability to incorporate basic classical models with realistic robustness to develop portfolios that are not only theoretically correct, but also workable amidst the uncertainty of modern finance.

Bibliography

- R. Almgren and N. Chriss. Optimal execution of portfolio transactions. *Journal of Risk*, 3(2):5–39, 2000. URL <https://doi.org/10.21314/JOR.2000.041>.
- F. Black and R. Litterman. Global portfolio optimization. *Financial Analysts Journal*, 48(5):28–43, 1992. URL <https://doi.org/10.2469/faj.v48.n5.28>.
- L. Breiman. Random forests. *Machine Learning*, 45(1):5–32, 2001. URL <https://doi.org/10.1023/A:1010933404324>.
- R. Cont. Empirical properties of asset returns: Stylized facts and statistical issues. *Quantitative Finance*, 1(2):223–236, 2001. URL <https://doi.org/10.1080/713665670>.
- D. A. Dickey and W. A. Fuller. Distribution of the estimators for autoregressive time series with a unit root. *Journal of the American Statistical Association*, 74(366a):427–431, 1979. URL <https://doi.org/10.1080/01621459.1979.10482531>.
- R. F. Engle. Autoregressive conditional heteroscedasticity with estimates of the variance of united kingdom inflation. *Econometrica*, 50(4):987–1007, 1982. URL <https://doi.org/10.2307/1912773>.
- F. J. Fabozzi, S. M. Focardi, and P. N. Kolm. *Robust portfolio optimization and management*. Wiley, Hoboken, 2006.
- S. Few. *Show me the numbers: Designing tables and graphs to enlighten*. Analytics Press, Oakland, 2nd edition, 2012.
- T. Fischer and C. Krauss. Deep learning with long short-term memory networks for financial market predictions. *European Journal of Operational Research*, 270(2):654–669, 2018. URL <https://doi.org/10.1016/j.ejor.2017.11.054>.
- J. H. Friedman. Greedy function approximation: A gradient boosting machine. *Annals of Statistics*, 29(5):1189–1232, 2001. URL <https://doi.org/10.1214/aos/1013203451>.
- S. Gu, B. Kelly, and D. Xiu. Empirical asset pricing via machine learning. *Review of Financial Studies*, 33(5):2223–2273, 2020. URL <https://doi.org/10.1093/rfs/hhz057>.
- J. D. Hamilton. *Time series analysis*. Princeton University Press, Princeton, 1994.
- M. C. Jensen. The performance of mutual funds in the period 1945–1964. *Journal of Finance*, 23(2):389–416, 1968. URL <https://doi.org/10.1111/j.1540-6261.1968.tb00815.x>.

- P. Jorion. *Value at risk: The new benchmark for managing financial risk*. McGraw-Hill, New York, 3rd edition, 2006.
- O. Ledoit and M. Wolf. Honey, i shrunk the sample covariance matrix. *Journal of Portfolio Management*, 30(4):110–119, 2004. URL <https://doi.org/10.3905/jpm.2004.110>.
- A. W. Lo. The adaptive markets hypothesis: Market efficiency from an evolutionary perspective. *Journal of Portfolio Management*, 30(5):15–29, 2004. URL <https://doi.org/10.3905/jpm.2004.318576>.
- S. M. Lundberg and S.-I. Lee. A unified approach to interpreting model predictions. In *Advances in Neural Information Processing Systems*, volume 30, pages 4765–4774, 2017. URL <https://doi.org/10.48550/arXiv.1705.07874>.
- H. Markowitz. Portfolio selection. *Journal of Finance*, 7(1):77–91, 1952. URL <https://doi.org/10.2307/2975974>.
- A. Meucci. The black–litterman approach: Original model and extensions. In F. J. Fabozzi, editor, *Encyclopedia of Financial Models*, pages 1–18. Wiley, Hoboken, 2010.
- R. Michaud. The markowitz optimization enigma: Is “optimised” optimal? *Financial Analysts Journal*, 45(1):31–42, 1989. URL <https://doi.org/10.2469/faj.v45.n1.31>.
- W. F. Sharpe. Capital asset prices: A theory of market equilibrium under conditions of risk. *Journal of Finance*, 19(3):425–442, 1964. URL <https://doi.org/10.2307/2977928>.
- R. S. Sutton and A. G. Barto. *Reinforcement learning: An introduction*. MIT Press, Cambridge, 2nd edition, 2018.
- R. S. Tsay. *Analysis of financial time series*. Wiley, Hoboken, 3rd edition, 2010.

Appendix A

Code Appendix: Figures

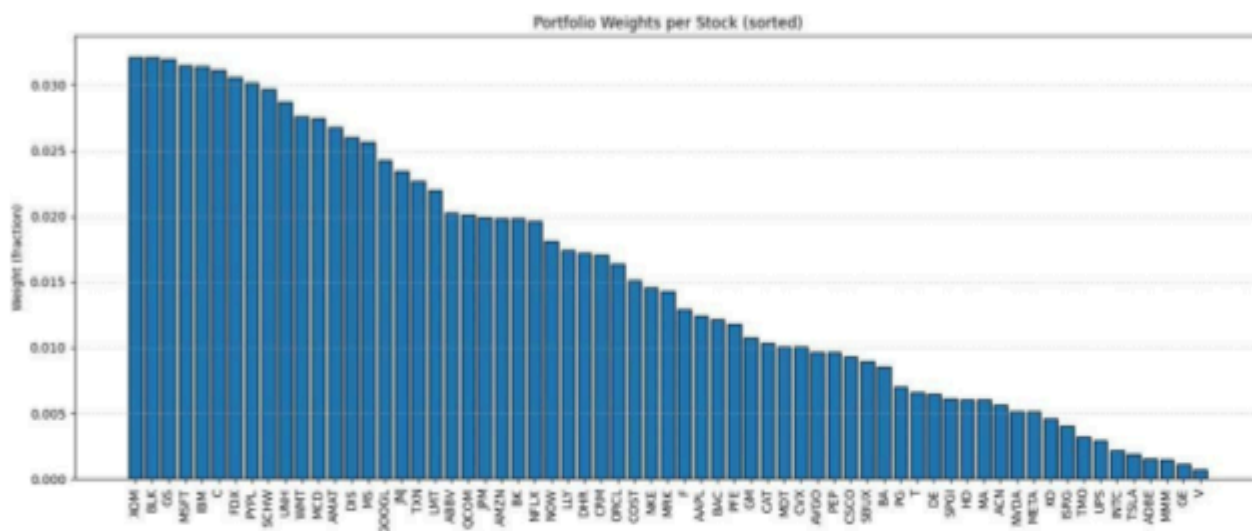


Figure A.1: Portfolio weights per stock in sorted form. (This figure illustrates the distribution of portfolio weights assigned to individual stocks, arranged in descending order. It highlights which stocks hold the largest shares within the portfolio and provides a clear view of diversification across assets.)




Figure A.2: Portfolio distribution weight represented in a pie chart. (This pie chart shows the relative weight of each stock in the portfolio, providing a quick view of asset allocation and diversification.)

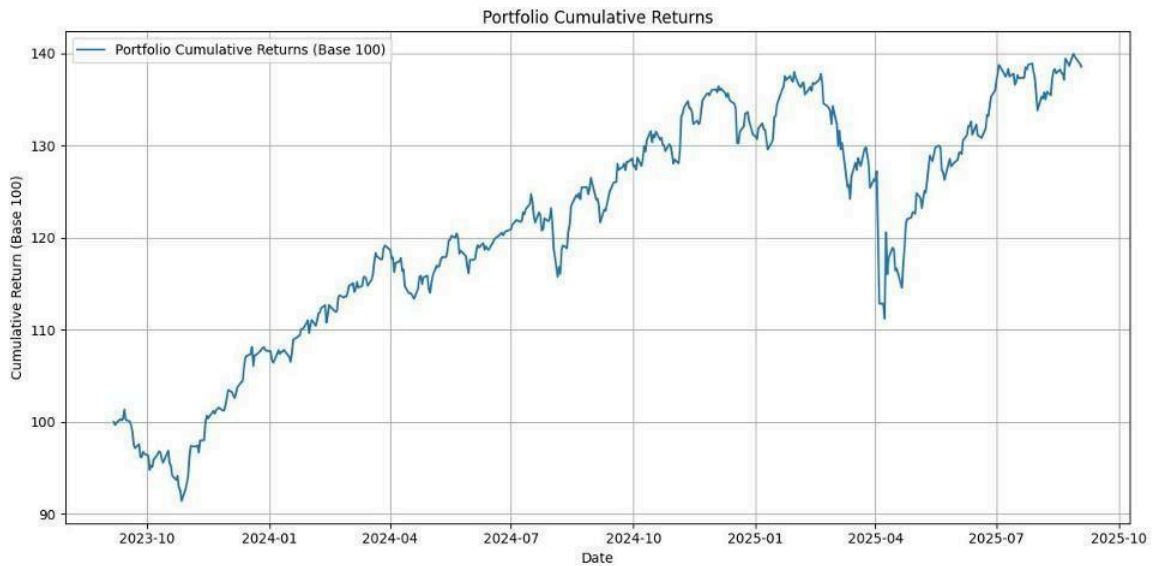


Figure A.3: Portfolio cumulative returns. (Shows the cumulative returns of the portfolio over time, highlighting the overall growth trend and performance.)



Figure A.4: 21-day rolling volatility with dates. (Illustrates the 21-day rolling volatility of the portfolio, showing fluctuations in risk over time along with corresponding dates.)



Figure A.5: Correlation matrix of log returns. (This correlation matrix of log returns displays the interdependencies between asset classes, with values ranging from -1 (perfect negative correlation) to 1 (perfect positive correlation).)




Figure A.6: Cumulative portfolio value over time – Linear Regression. (A linear regression line fits historical data points, indicating overall portfolio growth direction and aiding in assessing risk-return patterns.)




Figure A.7: Cumulative portfolio value over time – Gradient Boosting. (Cumulative portfolio value over time with Gradient Boosting leverages sequential machine learning models to capture complex patterns in portfolio growth, handling non-linear relationships effectively.)

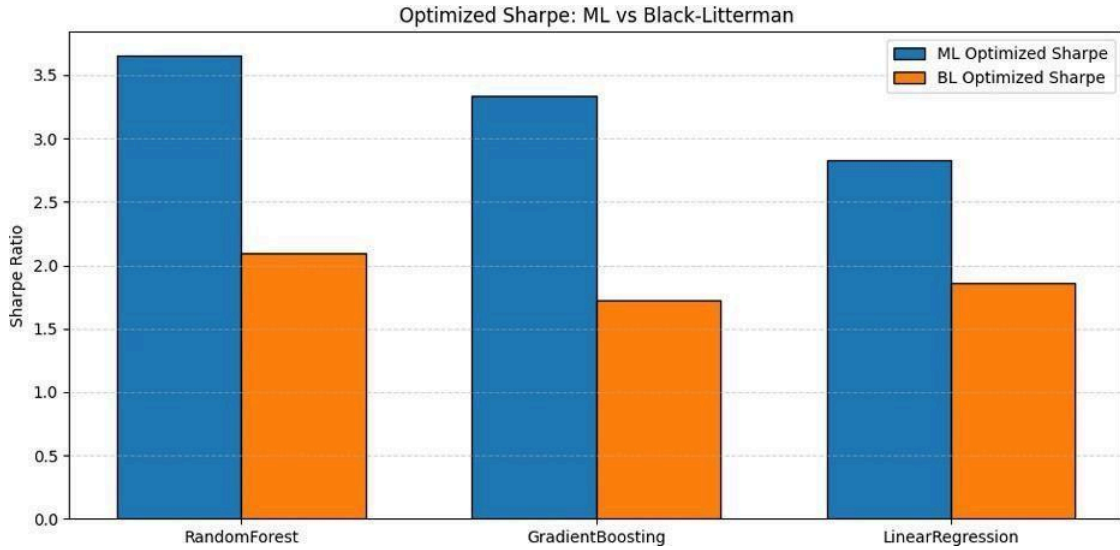


Figure A.8: Cumulative portfolio value over time – Random Forest. (Cumulative portfolio value over time with Random Forest uses multiple decision trees for robust modeling of portfolio growth. Captures complex interactions for insightful predictions on investment trends.)

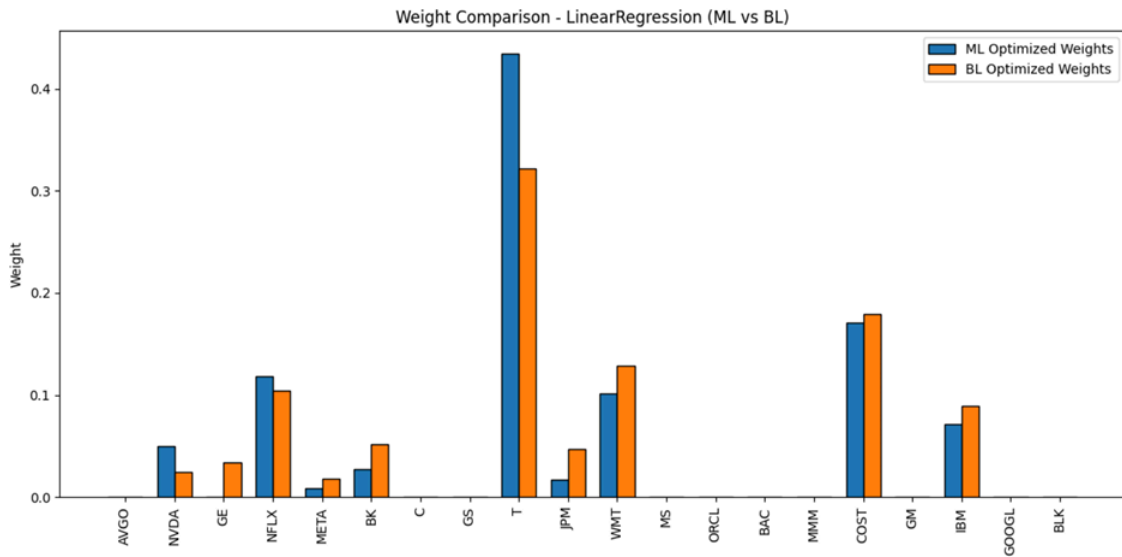


Figure A.9: Weight comparison – Random Forest (ML vs BL). (Random Forest weight comparison (ML vs BL) contrasts machine learning-derived feature importance with traditional backtesting (BL) approaches for portfolio insights.)

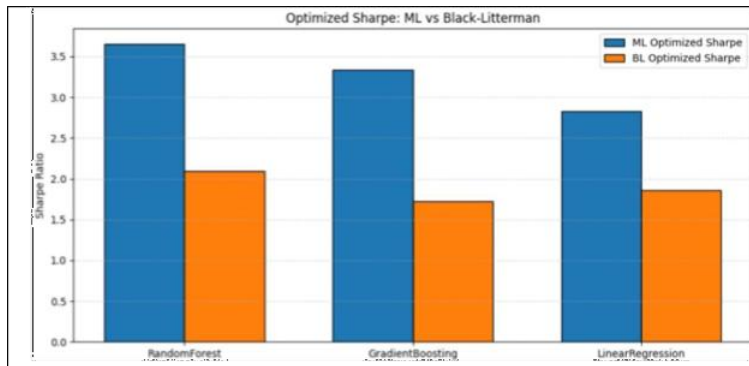


Figure A.10: Optimized Sharpe: ML vs Black-Litterman. (Optimized Sharpe ratio comparison between Machine Learning (ML) and Black-Litterman highlights how ML techniques and the Black-Litterman model differently enhance portfolio optimization, with Black-Litterman blending market equilibrium and investor views for potentially more robust risk-adjusted re- turns.)

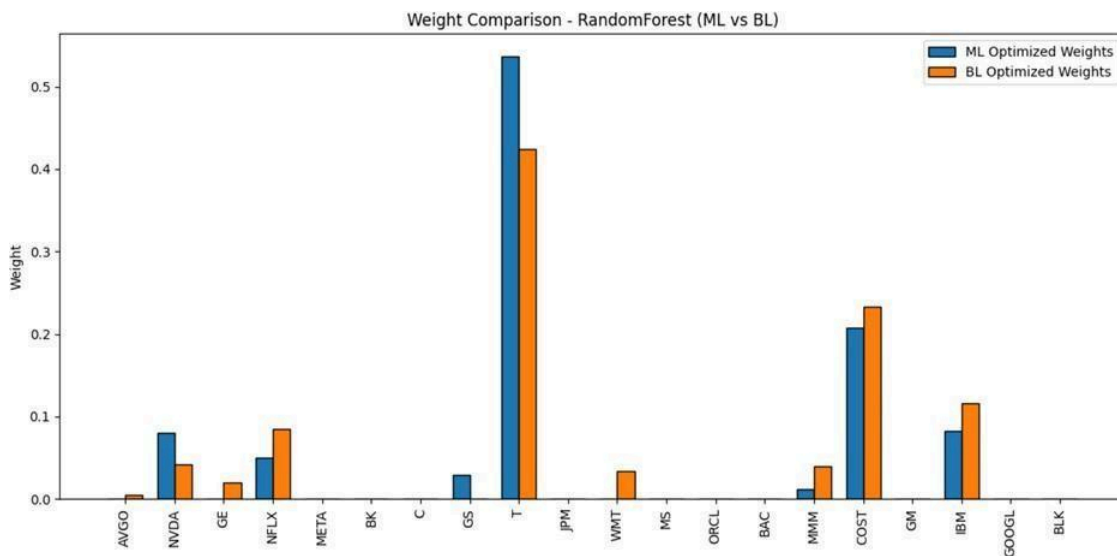


Figure A.11: Efficient frontier & Monte Carlo – Linear Regression (ML). (The efficient frontier paired with Monte Carlo simulations using Linear Regression (ML) generates probabilistic portfolio frontiers to visualize risk-return tradeoffs, aiding investors in optimizing asset allocations based on predicted performance distributions.)

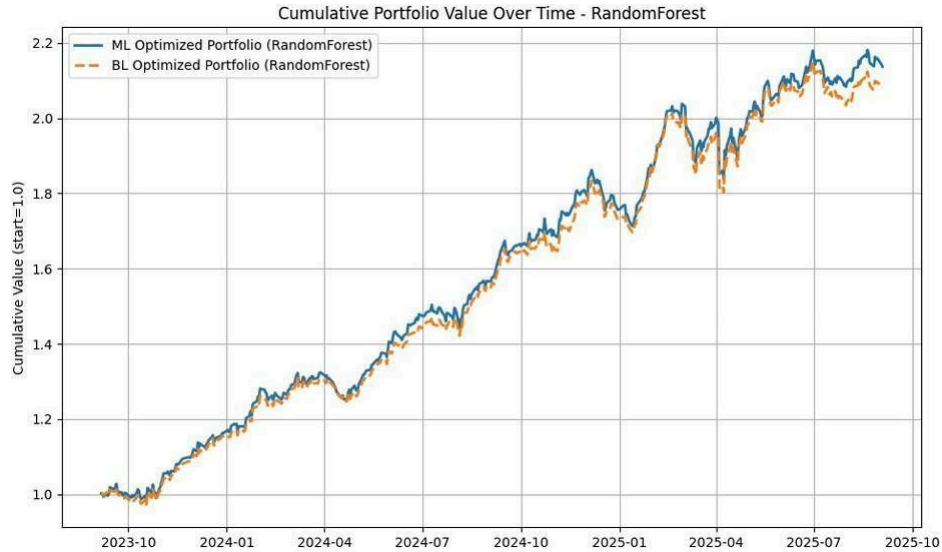


Figure A.12: Efficient frontier & Monte Carlo – Linear Regression (BL). (Efficient frontier & Monte Carlo with Linear Regression (BL) illustrates traditional portfolio optimization boundaries using baseline (BL) methods, showcasing risk-return profiles via simulated asset allocation scenarios.)

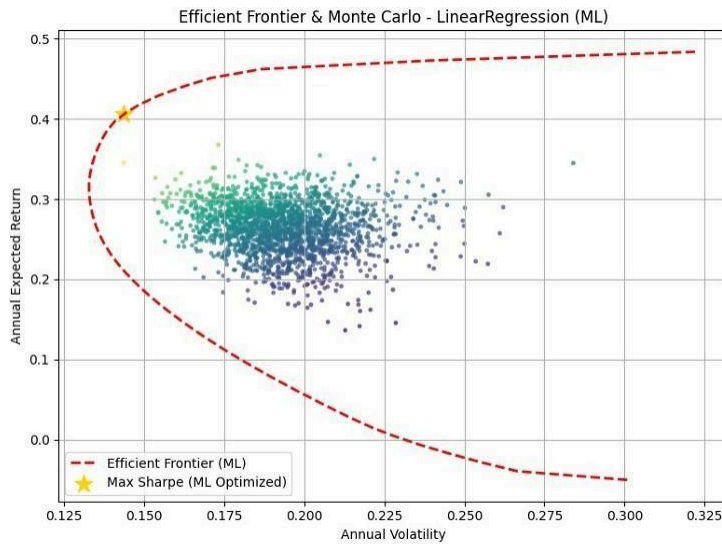


Figure A.13: Efficient frontier & Monte Carlo – Random Forest (ML). (Efficient frontier & Monte Carlo with Random Forest (ML) demonstrates advanced portfolio optimization using machine learning’s non-linear capabilities, potentially uncovering complex risk-return patterns beyond traditional models.)

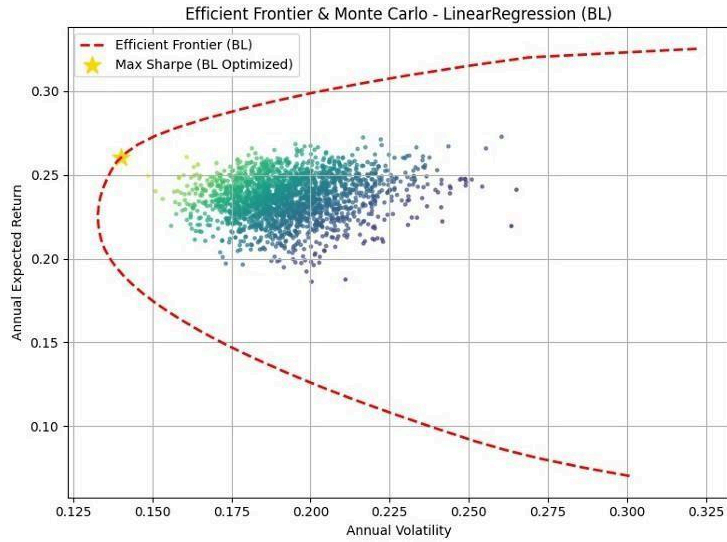


Figure A.14: Efficient frontier & Monte Carlo – Random Forest (BL). (Efficient frontier & Monte Carlo with Random Forest (BL) contrasts baseline (BL) applications of Random Forest in portfolio optimization, highlighting traditional interpretations of risk-return tradeoffs via simulated asset allocations.)

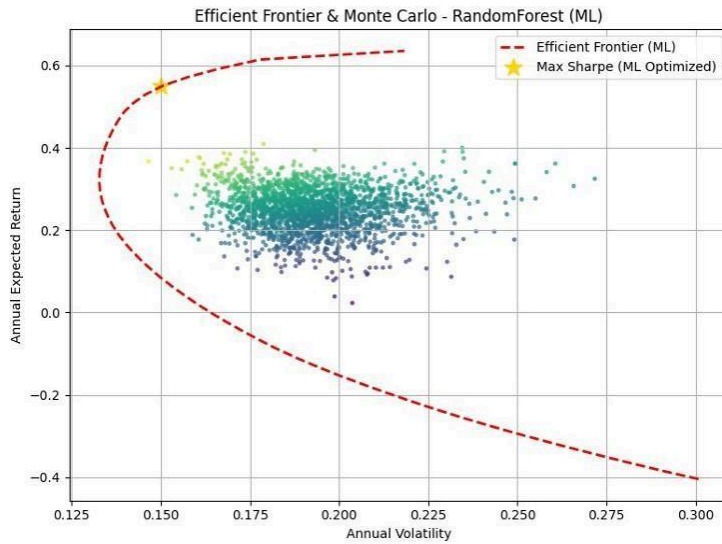


Figure A.15: Efficient frontier & Monte Carlo – Gradient Boosting (ML). (Efficient frontier & Monte Carlo with Gradient Boosting (ML) leverages powerful sequential learning for nuanced portfolio optimization, capturing intricate market patterns to shape risk-return frontiers for potentially enhanced investment strategies.)

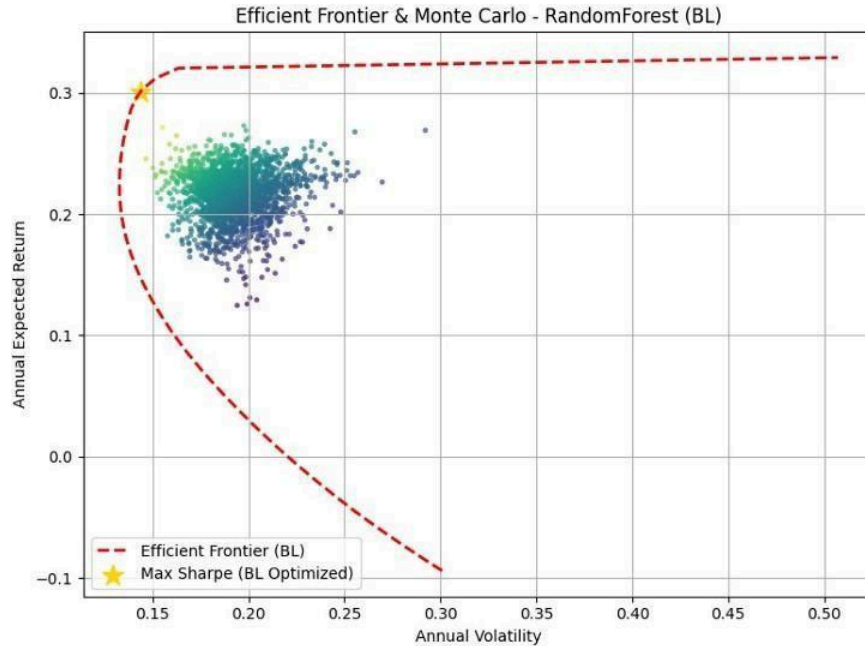


Figure A.16: Efficient frontier & Monte Carlo – Gradient Boosting (BL). (Efficient frontier & Monte Carlo with Gradient Boosting (BL) applies baseline Gradient Boosting techniques for traditional portfolio analysis, focusing on standard interpretations of simulated risk-return outcomes for asset allocation insights.)




Figure A.17: Final cumulative return comparison of all models used. (Final cumulative return comparison of all models synthesizes performance outcomes across machine learning (ML) and baseline (BL) approaches, highlighting relative efficacy in capturing portfolio growth trends and informing investment decisions.)

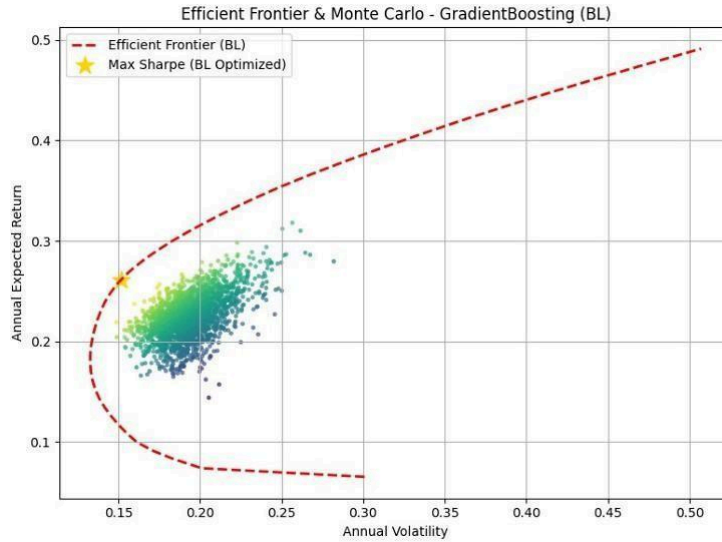


Figure A.18: Cumulative return comparison across models. (Cumulative return comparison across models visually contrasts the terminal wealth outcomes of diverse portfolio strategies, spotlighting performance disparities between machine learning-driven and traditional investment approaches.)

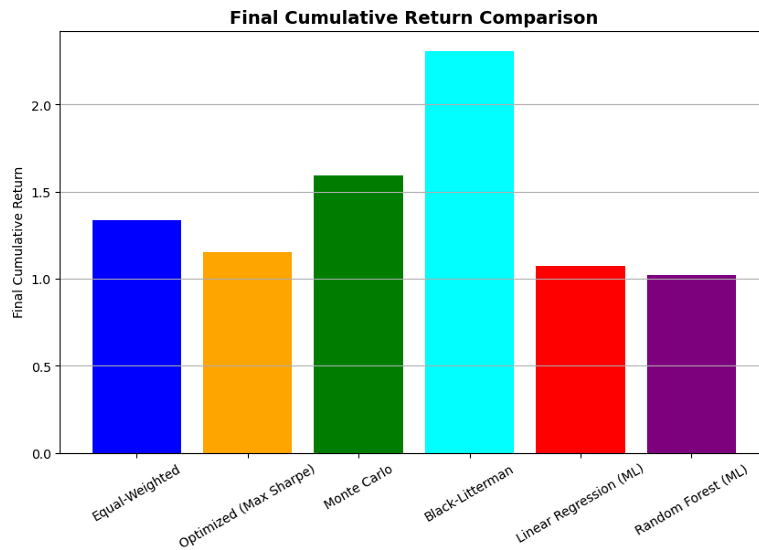


Figure A.19: Prices overview series with their values. (Prices overview series with their values displays historical price trajectories of analyzed assets, providing contextual market data underlying the portfolio optimization and modeling efforts showcased in comparative performance assessments.)

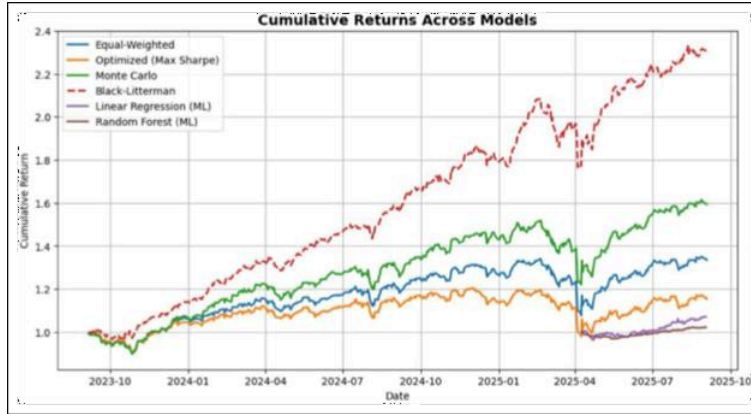


Figure A.20: Return correlation with ticker. (Return correlation with ticker illustrates inter-asset return relationships via correlation matrices, aiding identification of diversification opportunities and portfolio construction insights tied to specific securities.)



Figure A.21: 21-day rolling volatility (Annualized) with date and values. (21-day rolling volatility (Annualized) with date and values tracks short-term risk dynamics of assets over time, spotlighting fluctuations in annualized volatility to inform risk management and asset allocation decisions.)

Return Correlation

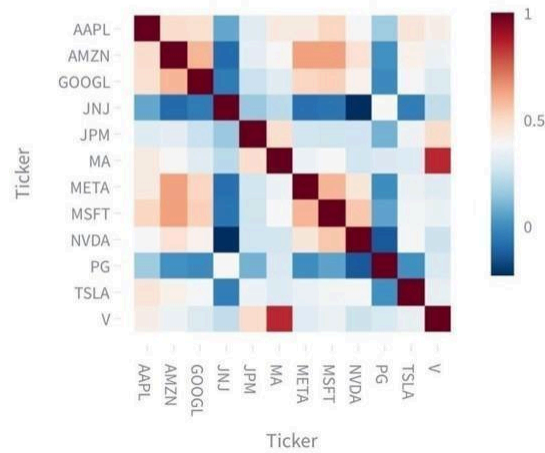


Figure A.22: Financial model Markowitz weights (Non-Negative) values with their assets and weights. (Financial model Markowitz weights (Non-Negative) values with their assets and weights showcases optimal asset allocations from Markowitz mean-variance optimization, highlighting portfolio composition under traditional Modern Portfolio Theory constraints.)

Rolling Volatility

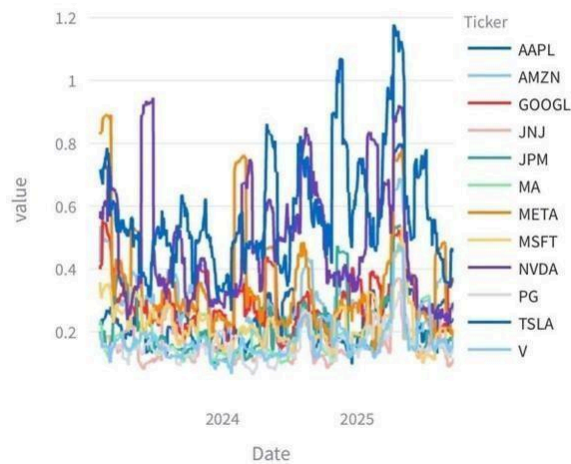


Figure A.23: Financial model Black–Litterman model (Non-Negative) values with their assets and weights. (Financial model Black–Litterman model (Non-Negative) values with their assets depicts portfolio weights blending equilibrium returns and investor views via Black–Litterman, reflecting tailored asset allocations incorporating subjective insights.)

Markowitz weights (non-negative)

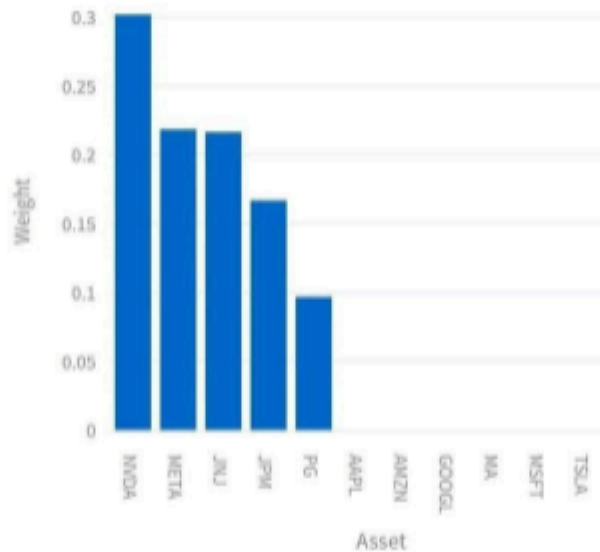


Figure A.24: ML model daily composite views with their assets and weights. (ML model daily composite views with their assets and weights reveals machine learning-derived daily asset views for portfolio steering, indicating data-driven perspectives on relative attractiveness influencing investment positioning.)

Black-Litterman weights (non-negative)




Figure A.25: Out of sample cumulative return of models used with their values and dates. (Out of sample cumulative return of models used with their values and dates contrasts realized performance trajectories of tested portfolio strategies beyond training periods, assessing comparative investment efficacy in unseen market conditions.)

Appendix B

Complete Code Implementation

B.1 Complete Code Link

The complete implementation code for this portfolio optimization platform, including all data preprocessing, model development, optimization algorithms, and dashboard creation, is available at the following Google Colab notebook:

Complete Code Repository:
https://colab.research.google.com/drive/1sT8i3C4_WsNBbfckmJdlRlsuGORnIY2V?usp=sharing

This notebook contains:

- Data collection and preprocessing scripts
- Implementation of classical financial models (CAPM, Monte Carlo, VaR, Black–Litterman)
- Machine learning models (Linear Regression, Random Forest, Gradient Boosting)
- Portfolio optimization algorithms
- Hybrid framework implementation
- Interactive dashboard code
- Performance evaluation metrics
- Visualization and plotting functions

The code is well-documented and can be executed directly in Google Colab environment. Users can modify parameters, upload their own datasets, and experiment with different optimization approaches as described in this thesis.