Degree Program in FINANCE
Department of ECONOMICS AND FINANCE

Chair of Asset Pricing

Markowitz and Black-Litterman Models in Portfolio Optimization: A Comparative Analysis of Static and Dynamic Strategies

Supervisor:
Prof. Nicola Borri

Co - Supervisor:
Prof. Valerio Marchisio

Candidate:
Vittorio Suarato
ID: 770651

Academic Year 2023/2024

Contents

Introduction

Every scientific discipline has a distinct way of shaping society's development. For economics and finance, this influence is widely recognized: the field's primary mission is to direct scarce resources toward their most productive uses. Among the most prominent manifestations of this goal is the activity carried out by the asset management industry. Financial institutions worldwide oversee trillions of dollars in funds, having to determine where to invest them in order to foster both returns and broader economic value. These vast pools of capital flow into countless ventures across the global economy, and there is a well-established consensus that *how* these resources are allocated—namely, the proportions invested in different asset classes and geographic markets—represents the single most consequential decision facing institutional investors (Black and Litterman, 1992). To guide these decisions, the discipline has introduced a host of theoretical frameworks and quantitative techniques, such as portfolio optimization methods. These methods serve as a guiding framework for managers to determine how best to allocate among different assets in a single fund or portfolio.

Historically, much of the theory in this field can be traced to the seminal work of Markowitz (1952). By focusing solely on mean and variance of returns, Markowitz demonstrated that, for an investor pursuing mean–variance efficiency, positioning a portfolio on the *efficient frontier* offers a higher expected return for the same level of risk. Despite its theoretical elegance, the Markowitz (or mean–variance) model has faced criticism in real-world settings. Researchers such as Michaud (1989) noted its tendency to produce concentrated "corner solutions" and remain highly sensitive to estimation errors—particularly errors in expected returns (Chopra and Ziemba, 1993).

In response to these limitations, Black and Litterman (1991a; 1991b) proposed an alternative paradigm that merges the structure of equilibrium models with investor views. Under the so-called Black–Litterman model, a baseline portfolio derived from market-equilibrium implied returns is updated by blending in subjective forecasts. This Bayesian process tempers extreme allocations by "shrinking" them back toward the market, while still allowing knowledgeable investors to incorporate signals or forecasts.

In this thesis, we compare three portfolio construction approaches: the classical Markowitz (mean–variance) Model , the Black–Litterman Model without views, as proxy for a passive benchmark strategy, and the Black–Litterman Model incorporating investor views. We compare these models under two distinct frameworks. Firstly, a buy-and-hold strategy, where the portfolio is optimised once at the beginning of the out-of-sample (OOS) period and held for five years without rebalancing. Secondly, in a rebalancing framework, in which portfolio weights are updated periodically in response to new market information and changes in the investor's views.

In the static setting, we test the Black–Litterman model's performance by analyzing the entire distribution of portfolio outcomes across 1,000 Monte Carlo simulations of potential views. Each view is generated with uncertainty proportional to the in-sample volatility of the underlying assets, and the assignment of those views fluctuates around a baseline optimistic and pessimistic scenario. This avoids relying on a single set of correct or incorrect views.

Next, we introduce a rebalancing framework, where the inputs in the models and the Black–Litterman views update on a quarterly basis to reflect evolving market information. Here, we examine and compare scenarios in which investor views are more frequently accurate (an optimistic case, with 75% of correct signals and 25% incorrect signals) versus systematically misaligned views (pessimistic case). Overall, these methodologies allow for a comprehensive investigation of static versus dynamic portfolio allocation decisions on the models, as well as the role of different assumptions on expected returns.

This thesis is organized as follows. Chapter 1 lays out the theoretical framework, detailing both the mean–variance and Black–Litterman models, highlighting their core principles and mathematical framework. Chapter 2 explains the empirical methodology: it describes the data, the specific implementation steps for each model, and the backtesting procedures (both no-rebalancing and quarterly rebalancing). Chapter 3 then reports and discusses the findings obtained from these tests. Finally, a concluding section summarizes the key results.

Chapter 1

Theoretical Framework

1.1 Markowitz Mean–Variance Model

1.1.1 Historical Context and Motivations

As previously mentioned, the development of Modern Portfolio Theory (MPT) in the 1950s signaled an important shift toward a more rigorous, quantitative approach to investing. Prior to Harry Markowitz's seminal contribution, many portfolio managers relied on heuristics and so-called "naïve diversification" methods, distributing capital almost equally across a range of assets without a systematic mechanism to evaluate their interdependencies (Michaud, 1989). Such simplistic techniques occasionally provided basic diversification benefits, but they could not adapt to changing market conditions, nor could they precisely account for variations in the volatilities and correlations of individual assets.

One of Markowitz's key insights was that risk and return should be jointly optimized, rather than considered in isolation (Markowitz, 1952). His framework used variance (or standard deviation) to measure the risk of a portfolio and proposed that investors should seek efficient combinations of assets along an "efficient frontier." This concept proved groundbreaking because, unlike earlier methods, it explicitly measured and optimized correlation effects between assets. In doing so, it laid out a mathematical foundation for investors to systematically balance the trade-off between risk and expected return, rather than attempting to mitigate risk purely by holding many different securities (Elton and Gruber, 1997).

The shift from heuristic approaches to robust, data-driven portfolio construction continues to shape how practitioners design, monitor, and rebalance multi-asset portfolios, serving as the cornerstone for subsequent models, including the Black–Litterman framework. Given the importance of Markowitz's principle to subsequent portfolio theory, it is helpful to briefly revisit its underlying mathematical framework.

1.1.2 Mathematical Formulation

A central assumption of the Markowitz Mean–Variance Model is that investors are risk-averse, meaning they require higher expected returns to justify taking on additional uncertainty. The model typically assumes a single-period horizon, within which asset returns follow a normal distribution. Under this framework, investors aim to balance risk and return by examining the trade-off between expected portfolio return and its variance.

Mathematically, let $\mathbf{w}$ be a vector of portfolio weights, $\boldsymbol{\mu}$ be the vector of expected returns for the underlying assets, and $\boldsymbol{\Sigma}$ be the covariance matrix of asset returns. The expected return of the portfolio is given by:

$$\mathbf{w}^\top \boldsymbol{\mu},$$

while its variance is represented by:

$$\mathbf{w}^\top \boldsymbol{\Sigma} \mathbf{w}.$$

Markowitz's mean–variance optimization problem can be structured to either minimize portfolio variance for a given expected return or maximize expected return for a given level of risk. A typical formulation involves choosing $\mathbf{w}$ to:

$$\min_{\mathbf{w}} \quad \mathbf{w}^\top \boldsymbol{\Sigma} \mathbf{w}$$

subject to

$$\mathbf{w}^\top \boldsymbol{\mu} = \mu_p, \quad \sum_i w_i = 1, \quad w_i \geq 0,$$

where μ_p is the target return. By solving this optimization for various levels of μ_p, one traces out the efficient frontier, which illustrates all possible portfolios offering the highest expected return for each risk level. Portfolios lying below this frontier are deemed suboptimal, as they provide lower returns for the same or higher variance.

1.1.3 Utility Functions and Optimal Portfolio Selection

A common approach to capturing the risk-averse behavior mentioned above is to assume an *exponential utility function*, often written as

$$U(W) = -e^{-\lambda W},$$

where W represents the terminal wealth of the investor, and λ is the *risk aversion* parameter. The exponential form implies that as λ increases, an investor becomes more sensitive to uncertainty. Although the original Markowitz model used mean and variance as key decision criteria, incorporating a specific utility function allows analysts to formally link risk preferences to portfolio choices.

In practice, one can combine expected return and variance into a single maximization problem by weighting the negative of variance by λ. That is, let $\mathbf{w}$ be the vector of portfolio weights, $\boldsymbol{\mu}$ the expected returns, and $\boldsymbol{\Sigma}$ the covariance matrix. Then the investor's objective might be stated as:

$$\max_{\mathbf{w}} \left[\mathbf{w}^\top \boldsymbol{\mu} - \frac{\lambda}{2} \mathbf{w}^\top \boldsymbol{\Sigma} \mathbf{w} \right] \quad \text{subject to} \quad \sum_i w_i = 1.$$

This formulation reflects the trade-off between return (first term) and risk (second term), scaled by λ. By differentiating and setting the gradient to zero, one can derive the *optimal portfolio weights* that balance these competing factors.

Interpreting the Markowitz solution involves recognizing that risk-averse investors will hold portfolios on the efficient frontier, tailored by their particular λ value. Higher λ leads to portfolios with lower variance but also lower expected return. The model thus highlights how varying levels of risk aversion can shift an investor's position along the frontier.

1.1.4 Limitations and Possible Remedies

Despite its foundational role in portfolio optimization, the Markowitz model exhibits notable weaknesses that compromise its real-world applicability. A primary concern, underscored by Michaud (1989), is its extreme sensitivity to small input changes in expected returns and variances—such deviations can dramatically alter the optimal portfolio weights. Best and Grauer (1991) further confirm that even minimal changes in mean estimates can result in drastic reallocation, often excluding key securities. We illustrate an example of these sensitivities in Section 3.1.2.

Beyond sensitivity, another pressing problem is the model's poor out-of-sample performance, often attributed to estimation error in expected returns and covariances. Merton (1980) highlights that using historical averages of realized returns without accouting for risk variations could be detrimantal, and suggests that estimators should account for heteroskedasticity of risk, while expected returns should be constrained to nonnegative values for improved accuracy. Empirical studies by Broadie (1993) reveal how limited historical data can inflate such errors, whereas longer time horizons can introduce nonstationarity. Chopra and Ziemba (1993) demonstrate that inaccuracies in means are approximately ten times more detrimental than errors in variances or covariances, a finding that is particularly severe for aggressive investors who weight returns heavily. Meanwhile, Ledoit and Wolf (2004) show that reliance on sample covariance matrices inflates estimation risk, advocating instead for shrinkage techniques to produce more stable portfolios.

Because of these issues, the Markowitz model's out-of-sample performance often compares badly to simpler benchmarks like the naive 1/N portfolio. Jorion (1991) finds that an equally weighted or value-weighted portfolio performs similarly to a minimum-variance (MINV) portfolio. Meanwhile, De Miguel et al. (2009) examine 14 datasets over 60- and 120-month estimation windows and show that none of the MV specifications reliably surpass a naive strategy. They argue that estimation errors—and the lengthy sample periods required for accuracy—erode any potential MV gains.

Other remedies have been proposed. For instance, Jagannathan and Tongshu (2003) suggest adding nonnegativity constraints on portfolio weights, showing through Monte Carlo simulations that these shrink extreme weights and mitigate sampling error, making portfolios with sample covariance matrices comparable to more sophisticated factor–based or shrinkage methods. Similarly, Levy and Levy (2014) introduce Variance–Based Constraints (VBC) and Global Variance–Based Constraints (GVBC), which cap asset weights according to their variances, leading to more stable Sharpe ratios than traditional MV portfolios.

Bayesian statistics constitute another major strand of solutions. Pástor (2000) demonstrates that embedding prior beliefs about asset pricing models yields more robust allocations and can mitigate phenomena like the U.S. ”home bias”. Jorion (1986) employs Bayes–Stein estimators—originally proposed by Stein (1955)—to enhance return forecasts by combining sample data with a normal prior, arguing that the classical sample mean is inadmissible. His simulations confirm that this shrinkage–based approach significantly mitigates estimation errors, thereby improving risk–adjusted returns compared to classical sample means.

A revolutionising approach emerging from Bayesian methods is the Black-Litterman Model, which incorporates prior beliefs and shrinks estimates toward a broader market equilibrium model (Black and Litterman, 1991a; 1991b). The model stands out for blending investor views with a market equilibrium prior, providing a systematic way to adjust return forecasts and mitigate estimation risk.

Although numerous comparisons exist between mean–variance (MV) portfolios, naive allocations, and other methods, relatively few works concentrate on the Black–Litterman (BL) model's out-of-sample performance versus these competitors. Harris et al. (2017) offer an insightful example by extending BL to a dynamic asset allocation framework, revealing that time–varying estimates of return distributions substantially enhance portfolio metrics such as the Sharpe ratio, Value–at–Risk, and Conditional Value–at–Risk. Their findings show that this dynamic BL approach can outperform both naive diversification and purely benchmark–based strategies, partly due to improved downside risk management.

Bessler and Wolff (2015), in turn, examine the addition of commodities within a multi–asset context, testing a range of strategies—such as equally weighted, minimum–variance, mean–variance, risk–parity, and the BL approach. Their out–of–sample tests indicate that BL–based portfolios often surpass naive and strategically weighted solutions in terms of risk–adjusted returns. Similarly, Bessler et al. (2017) compare BL allocations with MV and naive diversification across stocks, bonds, and commodities; they find BL portfolios yield superior Sharpe ratios, lower turnover, and greater stability, particularly during economic downturns. These results highlight how integrating subjective inputs and equilibrium–implied returns—central to the BL model—can mitigate estimation errors inherent in classical MV optimization.

1.2 Black–Litterman Model

1.2.1 Core Principles

In what follows, we draw on the princples and derivations found in Black and Litterman (1991a, 1991b, 1992), He and Litterman (2002), Meucci (2010), Idzorek (2004) and Satchell and Scowcroft (2000).

A central innovation of the Black–Litterman Model is its Bayesian mechanism, which blends an investor's subjective view with a baseline derived from market equilibrium. This stands in contrast to classical mean–variance optimization, not so much for the uncertainty in its input estimates, but because of how these inputs are processed to generate portfolio allocations. As noted by Chopra and Ziemba (1993), the issue in traditional optimization is not just the possible presence of errors in expected returns or covariances, but the way these errors translate into highly unstable and extreme portfolio weights. The Black–Litterman model mitigates this issue by incorporating a Bayesian framework that shrinks estimates toward a more stable equilibrium, reducing the impact of extreme return forecasts on the final allocation. In a portfolio context, this implies that stand-alone ”direct” estimates for asset returns can lead to unstable allocations, whereas moderate adjustments around a common reference point (a diversified market baseline) help mitigate extreme outcomes.

The Black–Litterman framework treats the market-implied returns as a “prior” distribution, reflecting consensus beliefs embedded in the capital markets. It then updates this prior with “view information,” which can be absolute or relative statements about expected returns on specific assets. By specifying both the magnitude of each view (how much one expects an asset to outperform) and its uncertainty, investors effectively control how strongly the new information pulls the final (posterior) estimates away from the prior. The posterior expected returns become a weighted combination of the equilibrium baseline and the stated views, moderated by a covariance structure that penalizes conflicting or highly uncertain assertions.

Therefore, another key aspect of the Black–Litterman models is that it allows for flexible incorporation of expert judgment, which is particularly useful in fast-moving or less efficient markets. As a result, investors obtain a more stable set of return forecasts and can construct portfolios that better reflect both market conditions and informed opinions.

1.2.2 Market Equilibrium Implied Returns

The first step in implementing the Black–Litterman model is the derivation of the *implied returns* from the market itself, a process often called *reverse optimization*. Rather than starting with sample means, the model infers each asset's expected return under the assumption that the market portfolio is already in equilibrium. Practically, one selects a baseline allocation—for example, a global market-cap-weighted portfolio—and applies a risk aversion parameter λ. This parameter governs how strongly investors trade off expected returns for additional risk.

The formula for implied returns takes the form:

$$\boldsymbol{\pi} = \lambda \boldsymbol{\Sigma} \boldsymbol{w}_{market}$$

where $\boldsymbol{\Sigma}$ is the covariance matrix of asset returns, $\boldsymbol{w}_{market}$ is the vector of weights in the market portfolio, and $\boldsymbol{\pi}$ is the resulting vector of equilibrium returns. Conceptually, if the market weight of a particular asset is high, or if its covariance with other assets is low, then the market consensus implies it should earn a higher return in equilibrium. In the absence of specific views, the Black–Litterman Model simply uses $\boldsymbol{\pi}$ as its posterior estimate of expected returns.

By rooting the model in a market equilibrium, Black–Litterman helps resolve the "overfitting" problem of classical optimization, giving the portfolio a starting point that is well diversified. From there, subjective views—if any—can be integrated smoothly, but even without them, as we will see from the emprical exercise, the equilibrium approach often produces allocations that avoid the extreme bets commonly observed in unconstrained mean–variance solutions.

1.2.3 Incorporation of Subjective Views

As mentioned, a defining feature of the Black–Litterman model is that it allows investors to incorporate their own opinions—often called *views*—about future returns into the equilibrium baseline. These views can be *absolute* (e.g., "Asset A's expected return is 5%") or *relative* ("Asset A will outperform Asset B by 2%"). To embed them formally, three main components are introduced:

1. **Pick Matrix (P)**: Suppose the investor has K views involving N assets. Each row of the $K \times N$ pick matrix P encodes how that particular view depends on each asset.

For instance, a row of

$$[0, 1, -1, 0, \ldots, 0]$$

would represent a relative view that the second asset outperforms the third.

2. **View Vector (Q)**: For each view $k \in \{1, \ldots, K\}$, a numerical target or forecast appears in the corresponding entry Q_k. If the view is absolute (say "Asset 2 grows by 4%"), then the relevant row in P identifies which asset is involved, and Q_k specifies 4%. If the view is relative (e.g., "Asset 2 will beat Asset 3 by 2%"), Q_k would be 2%.

3. **View Uncertainty Matrix (Ω)**: Because no forecast is perfect, investors specify the degree of confidence they hold in each view. This confidence is encoded in a diagonal (or more generally, positive-definite) matrix Ω. A lower diagonal entry Ω_{kk} means higher confidence in that view, so the Black–Litterman posterior shifts more strongly away from equilibrium and toward that specific view.

Additionally, a *scaling parameter* τ plays a crucial role. It governs the weight placed on market-implied returns relative to the investor's new information. A small τ indicates heavy reliance on equilibrium, whereas a larger τ amplifies the impact of the views. Balancing τ and Ω is key to obtaining stable, intuitive portfolios, rather than the extreme allocations sometimes seen in classical mean–variance.

1.2.4 Posterior Distribution and Final Expected Returns

Having established the market equilibrium prior and the incorporation of subjective views, we can now merge these elements to derive the *posterior* distribution of returns in the Black–Litterman framework.

Under the Black–Litterman approach, the original market model posits that returns $\mathbf{X}$ are normally distributed around some unknown mean $\boldsymbol{\mu}$ with covariance $\boldsymbol{\Sigma}$. Because $\boldsymbol{\mu}$ itself cannot be known with certainty, it is modeled as a random variable with prior distribution centered at the equilibrium vector $\boldsymbol{\pi}$, such that:

$$\boldsymbol{\mu} \sim \mathcal{N}(\boldsymbol{\pi}, \tau\boldsymbol{\Sigma})$$

where τ reflects the uncertainty in the equilibrium prior, with $\boldsymbol{\pi}$ typically representing CAPM-implied excess returns. When the investor introduces linear views

$$P\boldsymbol{\mu} \sim \mathcal{N}(\mathbf{Q}, \boldsymbol{\Omega}),$$

Bayesian updating *shrinks* the prior $\boldsymbol{\pi}$ toward the stated views $\mathbf{Q}$, weighted by the precision of each view.

Following Bayes' rule, the *posterior distribution* for $\boldsymbol{\mu}$ (the expected returns) emerges as a weighted combination of equilibrium and views:

$$\boldsymbol{\mu}_{BL} = \left[(\tau\boldsymbol{\Sigma})^{-1} + P^{\top}\boldsymbol{\Omega}^{-1}P\right]^{-1}\left[(\tau\boldsymbol{\Sigma})^{-1}\boldsymbol{\pi} + P^{\top}\boldsymbol{\Omega}^{-1}\mathbf{Q}\right].$$

In practice, this closed-form solution let us interpret how subjective information reconciles with market equilibrium.

Once the model yields $\boldsymbol{\mu}_{BL}$, a *final expected returns* vector is available for standard mean–variance optimization. These posterior parameters incorporate equilibrium beliefs and subjective forecasts, and the subsequent optimization step usually avoids corner solutions. Instead, it yields portfolios that tilt away from the neutral market weighting , reflecting investor views in a controlled manner. In essence, the Black–Litterman posterior highlights how *expected returns* become a mixture of equilibrium and views, while the overall *covariance* structure adapts to reflect the uncertainty in those views.

Chapter 2

Empirical Methodology

2.1 Data Description

The empirical analysis relies on a dataset of nine major market indices spanning both equity and fixed-income asset classes, along with a proxy for alternative investments. All price data were retrieved from Bloomberg and Refinitiv, ensuring standardized and reliable sourcing from two major financial data providers.The selected indices provide monthly price observations from 29 September 2000 through 31 December 2024, resulting in a dataset of approximately twenty-four years. The choice to work at a monthly frequency is common in portfolio optimization research and reflects standard practice when estimating medium to long-term risk-return characteristics.

The nine indices represent distinct slices of the global investment universe. First, **SPBDU3TT** (the S&P U.S. Treasury Bill Current 3-Month Index) acts as a cash or near-risk-free *money market* proxy, measuring the performance of three-month U.S. Treasury Bills. For U.S. government bonds, **MERG0Q0** (ICE BofA U.S. Treasury Index) captures the performance of USD-denominated sovereign debt issued by the U.S. government (excluding bills, inflation-linked debt, and strips). Meanwhile, **MEREG00** (ICE BofA Euro Government Index) tracks EUR-denominated sovereign debt issued by eurozone members, providing a high-level representation of European government bonds (converted in USD). Turning to equities, the **S&P 500 (SPX)** serves as a benchmark for large-cap U.S. equities, capturing roughly 80% of total U.S. market capitalization. **MSCI Europe (MXEU000U)** provides exposure to a basket of large and mid-cap European equities, while **MSCI Japan (MXJP)** isolates the performance of Japanese equities. The **MSCI Pacific ex-Japan (MXPCJ)** index captures other major developed markets in the Asia-Pacific region, excluding Japan. To incorporate emerging market exposure, the **MSCI EM (MXEF)** index tracks large and mid-cap equities across diverse economies such as China, Brazil, and India. Finally, the **Wilshire Liquid Alternative (WLIQAT)** index aggregates several hedge-fund-type strategies, giving a glimpse into "liquid alternatives" that differ structurally from traditional stock-and-bond holdings.

A key motivation for selecting these indices is that they collectively approximate the primary asset allocation choices available to institutional portfolio managers. This approach aligns with standard portfolio allocation literature, which frequently groups global equities, government bonds, and alternative investments to form a broad market portfolio. Each index is widely regarded as sufficiently liquid, investable, and representative of its respective asset class. For instance, **MERG0Q0** and **MEREG00** serve as dedicated U.S. and European bond benchmarks, while the MSCI family of indices dominates institutional benchmarking across developed and emerging equity markets. Including the Wilshire Liquid Alternative Index ensures the coverage of hedge-fund-like or alternative exposures, increasingly used in institutional portfolios for diversification and potential downside protection.

Since most data were sourced directly from Bloomberg and Refinitiv in the form of historical monthly prices, extensive cleaning or preprocessing was not required. The time series is already aggregated at the monthly level, eliminating the need for manual resampling. As described in the subsequent methodology section, log returns are computed from these monthly prices before annualizing to align with common industry and academic conventions. With this data configuration, the study captures a broad time span that includes both bull and bear markets—spanning events such as the Dot-Com Bust, the 2008 Financial Crisis, and the more recent market turbulence. Such a wide sample helps ensure that the final results and comparisons between models reflect different market regimes and volatility conditions, thereby providing more robust conclusions about optimal portfolio construction in a global context.

2.2 Model Implementation

This section outlines how the theoretical framework described in Chapter 1 is operationalised in practice using MATLAB. The focus here is on translating the mean-variance optimization and the Black-Litterman model into concrete steps involving data inputs, parameter estimation, and model configuration, while deferring any discussion of results to Chapter 3.

2.2.1 Markowitz Model

Data Preparation and Parameter Estimation

As a first step, monthly price series for each asset class are transformed into monthly log returns by taking their natural logarithm and computing the differences across consecutive months. The sample mean of these monthly returns serves as the estimate for each asset's expected monthly return. These estimates are converted to an annual scale by multiplying by 12. The covariance matrix is derived from the same monthly returns and similarly scaled by a factor of 12 to reflect annual risk levels. The resulting set of annualized expected returns and annualized covariance matrix form the core inputs to the Markowitz optimization.

Building the Portfolio Object and Solving the Optimization Problem

We employ the `Portfolio` object from the MATLAB Financial Toolbox to streamline the optimization. First, we create a new portfolio, specifying the annualized expected returns and covariance matrix via the `setAssetMoments` function. Next, we define constraints with `setDefaultConstraints`, which enforces full investment and non-negativity of weights by default. We then use the built-in methods `estimateFrontier` and `plotFrontier` to compute and visualize a discrete set of efficient portfolios (50 by default). Each portfolio on this frontier reflects a different risk-return combination, effectively illustrating the trade-off Markowitz emphasized.

2.2.2 Black-Litterman Model

This subsection describes how the Black–Litterman (BL) model is translated into an empirical procedure. Although the steps below mirror some aspects of the Markowitz implementation, the key differentiator here is the introduction of implied equilibrium returns and subjective views, both of which modify the expected returns and covariance of the portfolio.

Data Preparation and Covariance Estimation

The BL approach begins with reading the same monthly price data used for the Markowitz model. As before, each asset's price series is converted into log returns, yielding the monthly returns matrix. We then compute a sample covariance matrix from these returns and annualize it by multiplying by 12. Although this empirical covariance matrix is ultimately input into the BL framework, the Black–Litterman approach itself replaces the need to *directly* rely on historical mean returns.

Reverse Optimization for Implied Returns

A central innovation of Black–Litterman is the derivation of *implied equilibrium returns* from a chosen market portfolio. In the MATLAB code, this step is represented by multiplying the annualized covariance matrix by a risk-aversion parameter, λ (set to 3 in our model) and the vector of **market weights**. By doing so, we assume that each asset's weight in the market portfolio reflects the market's collective belief about its expected returns.

Incorporating Investor Views for Posterior Estimates

Another key feature of Black–Litterman is the integration of subjective views on expected returns. In the code, these views are encoded in the matrix **P**, which identifies how each view relates to the assets, and **Q**, which specifies the magnitude of the predicted returns or spreads. The parameter Ω (the "view uncertainty" matrix) captures the investor's confidence in each stated view. A scaling factor (τ) then governs how heavily to weight the market equilibrium versus subjective view components. Once we specify our investor views via matrices **P** and **Q**, and define the uncertainty of these views with Ω, the code applies the formula defined in Section 1.2.4 to generate the posterior estimates.

Constructing the Black–Litterman Portfolio

Once the posterior estimates have been computed, the workflow parallels that of Markowitz optimization. In the code:

1. **Portfolio Object Creation:** A portfolio object is initialized with `setAssetMoments`, using the posterior expected returns *and* the posterior covariance.

2. **Constraints:** Default constraints (`setDefaultConstraints`) enforce full investment and non-negative weights.

3. **Efficient Frontier Calculation:** The step of generating an efficient frontier uses the same MATLAB functions as the standard Markowtiz approach (`estimateFrontier`, `plotFrontier`). Each point on this frontier represents a distinct risk–return combination under the new, BL-adjusted estimates.

As with the Markowitz procedure, we defer any presentation of numerical results or performance comparisons to Chapter 3. However, laying out these steps here clarifies how the Black–Litterman model is *empirically implemented* within a coding environment, bridging the theoretical background from Chapter 2 to the subsequent discussion of outcomes.

2.3 Backtesting Framework

2.3.1 Buy-and-Hold Strategy

In this subsection, we describe how the Black–Litterman (BL) and Markowitz models are implemented under a single-period, *no-rebalancing* framework. The goal is to investigate how each model performs when the portfolio is formed *once* at the start of the out-of-sample window and held until the end, under different *optimistic* and *pessimistic* view assumptions.

Overview of the Setup

We divide the dataset into an in-sample portion (80% of monthly returns) and an out-of-sample portion (the final 20%). The out-of-sample period consists of 59 months, approximately five years. During the *in-sample* phase, we estimate the mean and covariance matrix for Markowitz and compute implied returns for Black–Litterman. We then form *one set of portfolios* at the beginning of the out-of-sample period. This neglects any re-optimization over time, allowing us to assess how each model's initial estimates and views influence performance over an extended horizon.

The BL model with *no views* (using only equilibrium returns) serve as benchmark. We compare the standard Markowitz model to a BL model with views under two distinct scenarios:

1. **Optimistic Views**: Where the views align with the *realised* out-of-sample performance for one year ahead.
2. **Pessimistic Views**: Where the views align to the opposite of the realised 1 year-ahead performance (the views are reversed compared to the first case).

We include BL with no views to isolate the added value (or detriment) of incorporating explicit views. If the model with incorrect views performs worse than its benchmark, it suggests that the manager's "insight" adds negative value. Conversely, if a favorable scenario outperforms, it highlights the degree of potential benefit of accurate forecasts.

Scenario Design and View Generation

We now turn to defining more in detail the "optimistic" and "pessimistic" view scenarios for the Black-Litterman model. In particular, we use the first 12 months of the out-of-sample period to form a baseline *relative* view vector, denoted as $\mathbf{Q}_{\text{baseline}}$. Specifically, we sum each asset's next-year (12-month) realized returns and construct a set of differences reflecting absolute and pairwise relative performance. The relative views are structured logically to contrast key global asset classes, such as equities and bonds across different regions. For example, U.S. equities are compared to E.U. equities, while U.S. government bonds are compared to E.U. government bonds . This approach extends to other global asset classes, pairing Japan with Pacific ex-Japan equities and Emerging Markets with

Alternatives to reflect meaningful contrasts. The corresponding relative view matrix, **P**, is structured as follows:

$$P = \begin{bmatrix} 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 1 & -1 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & -1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 & -1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & -1 \end{bmatrix}$$

Each row in P corresponds to a specific relative view (except the first):

- First row: absolute view on the U.S. Money Market.
- Second row: U.S. government bonds vs E.U. government bonds.
- Third row: U.S. equity vs E.U. equity.
- Fourth row: Japanese equity vs Pacific ex-Japan equity.
- Fifth row: Emerging Markets equity vs Alternative investments.

Monte Carlo Simulation

Rather than setting **Q** exactly to $\mathbf{Q}_{\text{baseline}}$ in the *optimistic* scenario (or the negative of it in the *pessimistic* scenario), we introduce *random draws* ($N = 1000$) around these baseline values to conduct a more robust analysis and to reflect the fact that in practice, even a "correct" forecast is rarely perfect. Formally, each element Q_i in our view vector is modeled as:

$$Q_i \sim \mathcal{N}(\mu_i, \sigma_i^2), \tag{2.1}$$

where μ_i is either $\mathbf{Q}_{\text{baseline}}$ (the optimistic case) or its negative (the pessimistic case), and σ_i is proportional to the *in-sample* annualized volatility of the relevant underlying assets. By tying σ_i to the underlying assets' volatilities (or their average, for a relative view involving two assets), we ensure that views on riskier assets fluctuate more widely than those on stable ones. This avoids unrealistically uniform perturbations across all views and provides a more nuanced representation of how forecast uncertainty might scale with asset risk in an optimistic and pessimistic view scenario.

For *each* Monte Carlo draw, we compute a distinct posterior return vector and covariance matrix using the Black–Litterman formula, then solve for an entire efficient frontier. We select three representative portfolios—low (25%), mid (50%), and high (75%) of the maximum risk. These weights are applied to the out-of-sample returns, and we store the resulting performance metrics (annualized mean, total return, volatility, Sharpe ratio). We preserve all 1,000 draw outcomes and represent their performance through percentile tables and histograms. In particular, with this approach we produce:

1. **Percentile Tables** that summarize risk–return metrics (mean return, total return, volatility, Sharpe ratio) at the 5th, 25th, 50th, 75th, and 95th percentiles.

2. **Histograms** illustrating the distribution of each metric, offering insight into the spread of possible performance results.

3. A final **average BL portfolio** (across all simulations) to use in side-by-side and more direct comparisons with Markowitz and the benchmark.

Consequently, we obtain a more robust picture of how input uncertainty in the views translates into a range of final portfolio performances. This enhances comparability with real-world scenarios where even an "accurate" investor forecast is rarely perfect and is subject to asset-specific volatility.

Motivation for Using a Monte Carlo

If we assigned just a single *optimistic* view equal to the realized one-year return, that would imply near-perfect foresight. Similarly, a single *pessimistic* view set exactly to the negative of the realized return might be too pessimistically precise. In reality, even a knowledgeable portfolio manager rarely forecasts the exact difference among assets. Therefore, we *simulate* many possible slight deviations around each baseline to mirror a scenario where the manager is "mostly correct" or "mostly incorrect" but not perfectly so. This Monte Carlo approach yields a range of potential posterior returns and covariances for each draw, ensuring that the final Black–Litterman portfolios reflect uncertainty in the manager's actual forecasting skill. To summerise, we generate two distinct scenarios based on the direction of the Monte Carlo deviations:

1. **Optimistic Scenario:** The random deviations are centered around the *realized* return differentials ($\mathbf{Q}_{\text{baseline}}$), meaning that, on average, the investor's views align with future performance.

2. **Pessimistic Scenario:** The deviations are centered around the *negative* of $\mathbf{Q}_{\text{baseline}}$, reflecting a systematic bias where the investor's views contradict realized out-of-sample performance.

Choice of $\tau = 0.05$ and $\Omega = 0.05$

These parameters govern how strongly the posterior leans away from implied returns and toward the manager's views. In both scenarios, the Black-Litterman scaling factor (τ) and the view uncertainty matrix (Ω) are set to 0.05. This calibration is chosen to ensure that posterior return estimates remain stable while still allowing for meaningful updates based on the investor's views and comparability across scenarios.

Model Implementations

Three models are tested:

1. **Markowitz (Mean–Variance)**
   We estimate the in-sample expected returns and covariance matrix, then generate 50 portfolios along the efficient frontier. To facilitate comparison, we identify three *representative* portfolios at low (25%), mid (50%), and high (75%) of the maximum risk on the frontier.

2. **Black–Litterman (No Views)**
   To benchmark the BL approach, we set no explicit views, relying purely on implied equilibrium returns. As with Markowitz, we form an efficient frontier of 50 portfolios and choose the same three risk levels for comparison.

3. **Black–Litterman (With Monte Carlo Views)**
   We incorporate the baseline $\mathbf{Q}_{\text{baseline}}$ in a Bayesian update, but with *random draws* around that baseline (optimistic scenario) or its negative (pessimistic scenario). For each draw, we re-compute the posterior returns and covariance, then solve for the efficient frontier. We extract the same three risk-level portfolios, ultimately averaging the final weights across all draws to form *representative* low, mid and high risk allocations. In essence, we create a distribution of possible BL outcomes, then compare the average or representative outcome to Markowitz and the benchmark, while also displaying the performance across all the simulations.

2.3.2 Quarterly Rebalancing Approach

We now examine a quarterly rebalancing framework. In real-world asset management, practitioners typically update their portfolios over time to incorporate fresh market information or revised forecasts. This approach allows us to see whether dynamic re-estimation of Markowitz and Black–Litterman portfolios yields different outcomes compared to a one-shot allocation.

Overview of the Setup

Rather than holding a portfolio for the entire out-of-sample horizon, we divide the out-of-sample data into three-month quarters and re-optimize at each quarter boundary. In particular, we follow the below process:

1. We implement an *in-sample expanding window* approach, meaning that at each rebalancing date, we include all historical data up to that point (starting with the original in-sample plus any out-of-sample months that have already passed).

2. We use the newly estimated expected returns and covariances (Markowitz) or updated implied returns and views (Black–Litterman) to create an efficient frontier of candidate portfolios.

3. As in the previous approach, we pick three representative portfolios for each frontier at low (25%), mid (50%), and high (75%) target risk.

4. Over the following quarter, we apply the chosen portfolio weights, measure the realized returns, and store them.

5. After the quarter concludes, the in-sample period is extended, and the process repeats.

By the end of the out-of-sample period, we have a sequence of quarterly returns for each model. We then compute aggregated performance metrics, including annualized average return, volatility, and Sharpe ratio.

Scenario Design and View Generation

As with the no-rebalancing case, we test two view accuracy scenarios for Black–Litterman:

- **Scenario A (Mostly Correct)**
  - **75% of the time**, the manager's view equals *half* of the true realized returns for the upcoming quarter.
  - **25% of the time**, the manager's view equals *half* of the negative of the realized returns for the next quarter.
  - This ensures the manager is "right more often than not," but never perfectly correct.

- **Scenario B (Mostly Incorrect)**

– This scenario simply reverses those probabilities or signs, so that most of the time the manager is "partially wrong."

We scale the manager's view to 50% of the next quarter's true realized returns—whether aligned or inverted—so the portfolio tilt is neither extremely bullish nor extremely bearish. In each quarter, we apply a deterministic sequence of "good" vs. "bad" quarters. In other words, we predefine which quarters have views that align with the realized returns. This ensures a deterministic sequence of view accuracies while avoiding large run-to-run variability.

Key Parameter Choices

- $\tau = 0.05$ **and** $\Omega = 0.05$: We keep the same scaling and view-uncertainty levels as in the no-rebalancing case to maintain consistency across all tests.
- **Quarter Length = 3**: A three-month frequency is chosen for demonstration. One could easily adapt the code for monthly or semi-annual rebalancing.
- **Partial Scale (50%)**: We adopt 50% of the realized next-quarter return to avoid extreme overconfidence. This ensures that even in "correct" quarters, the manager's view is moderate.

Chapter 3

Results and Discussion

This chapter presents the empirical findings derived from the methodological framework established in Chapter 2. Our objective is to examine how each of the models performs under different scenarios and parameter choices.

We begin by illustrating the individual model outcomes on the full dataset, highlighting key properties such as allocation sensitivity (in Markowitz), market equilibrium exposures and view-driven adjustments (Black–Litterman). These demonstrations also incorporate a brief examination of how the Black–Litterman model responds to changes in τ (the prior-uncertainty scaling parameter) and Ω (the view-confidence matrix), providing insight into the model's flexibility.

Next, the main backtest results are divided into two parts: a buy-and-hold scenario, where the portfolio is formed once and held for the entire out-of-sample horizon, and a quarterly rebalancing scenario, where the portfolio is updated every three months. In each backtest, we compare how the three models perform under optimistic or pessimistic view assumptions, as well as across low, mid, and high-risk portfolio choices. By examining annualized returns, volatilities and Sharpe ratio, we obtain a comprehensive view of how each approach balances risk and reward under realistic market conditions.

3.1 Exploratory Analysis with the Full Dataset

3.1.1 Markowitz model

To illustrate the Markowitz mean–variance model's properties, we begin by applying it to the entire historical dataset (from 2000 through 2024) in a single optimization. Our objective is to demonstrate how the model builds its efficient frontier using the full set of observed monthly returns. In particular, we aim to show:

1. **Baseline Allocations**: How Markowitz assigns portfolio weights across nine asset classes as we move along the risk spectrum from minimum-variance to high-volatility portfolios.

2. **Sensitivity to Small Input Changes**: How even minor adjustments to the model's expected returns can significantly alter the resulting frontier allocations—a well-documented shortfall of mean–variance optimization.

This isolates the model's basic behavior—and underscores the challenges that motivate alternatives like Black–Litterman in subsequent sections.

Baseline Frontier & Stacked Area Chart

Figure 3.1 provides an illustration of how the Markowitz model allocates weights across the nine asset classes, given the historical expected returns and risks shown in Table 3.1.

Asset Class	Expected Return (%)	Risk (%)
US Money Market	1.9	0.6
U.S. Govt. Bonds	3.3	4.9
Euro Govt. Bonds	4.0	10.7
U.S. Equity	5.8	15.4
Euro Equity	1.6	18.5
Japan Equity	1.1	15.7
Pacific ex-Japan Equity	3.3	20.1
Emerging Markets Equity	4.2	21.1
Alternatives	2.4	4.0

Table 3.1: Annualized expected returns and standard deviations for each asset class, estimated from the entire historical dataset.

The table summarizes our annualized inputs—return and standard deviation—for each asset class. These parameters are fed into the Portfolio object in MATLAB to generate fifty points along the efficient frontier, ranging from the minimum-variance portfolio on the left (lowest standard deviation) to the highest-risk allocation on the right.

In the stacked chart below, each color band corresponds to an asset class, whose share of the portfolio varies as we move from low to high risk. For instance, blue (US Money Market) is more prominent at the extreme left, while purple (U.S. Equity) grows toward the right. Orange (U.S. Government Bonds), yellow (Euro Government Bonds), and other colored segments represent the remaining asset classes, as labeled in the legend.

At the low-risk end of the frontier (far left), the chart is dominated by US Money Market (blue) and U.S. Government Bonds (orange). This confirms their role as conservative, lower-volatility holdings, consistent with their modest returns but relatively small standard deviations (Table 3.1). As the portfolio seeks higher return (and tolerates greater risk), it pivots toward equities— in particular U.S. Equity (purple).

3.1.2 Demonstrating Sensitivity to Input Changes

A well-known characteristic of Markowitz mean–variance optimization is its *sensitivity* to the precise values of expected returns and covariances. Even modest deviations in these

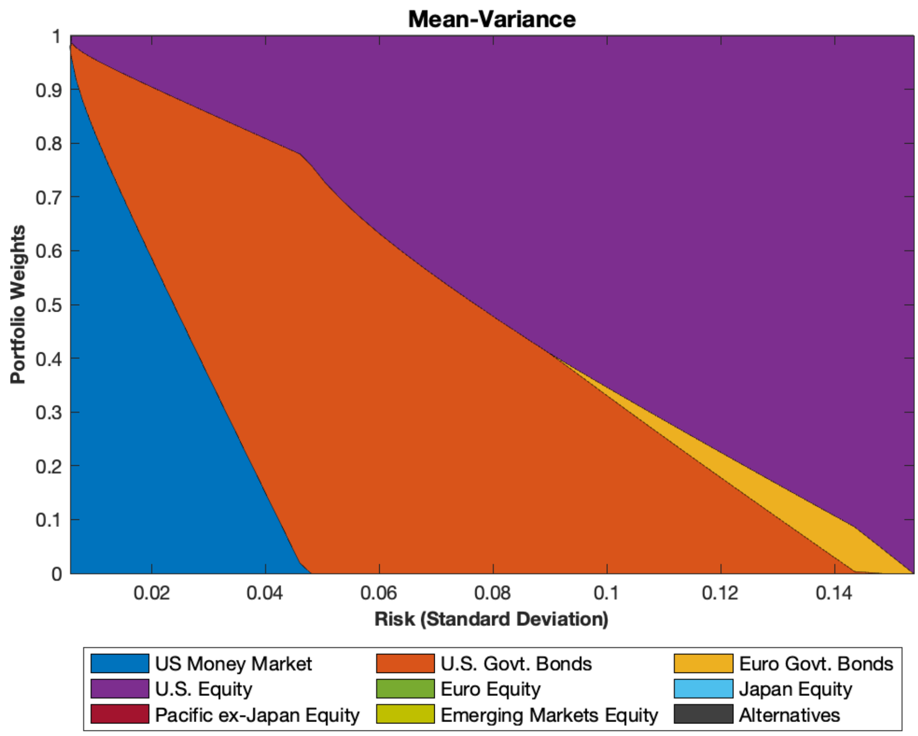

Figure 3.1: Mean–variance model frontier illustrating portfolio weights versus risk, based on the expected returns and standard deviations shown in Table 3.1.

estimates can lead to disproportionately large shifts in portfolio allocations. To illustrate this, we hold the original risk values and constraints fixed, but tweak the annualized expected returns of each asset class by ±1% (as shown in Table 3.2). We then re-run the optimization to generate the new efficient frontier, plotted in Figure 3.2.

The results highlight the model's lack of robustness and tendency toward overconcentration—often yielding *corner solutions* where most of the allocation funnels into only one or two assets. A one-percent increase in Emerging Markets Equity (or a slight reduction in US Equity) swings a significant portion of the frontier toward that asset class. Comparing Figures 3.1 and 3.2, we see large reallocations from U.S. bonds and money markets into alternatives and EM equity, despite relatively small changes in their expected returns. In real settings, a 1% shift in a portfolio manager's forecast is fairly normal and does not necessarily contradict the broader trend; yet observing these dramatic changes in allocation is disconcerting. It raises the question: *Can we trust a model that so heavily amplifies a small basis-point revision in opinion?* This phenomenon arises because mean–variance optimization seeks to exploit any perceived return advantage, often "piling into" a single asset or sector whose expected return appears even marginally higher, a pattern frequently criticized in both academic literature and practitioner surveys for its lack of robustness.

Asset Class	Original Return (%)	New Return (%)	Risk (%)
US Money Market	1.9	0.9 (-1%)	0.6
U.S. Govt. Bonds	3.3	2.3 (-1%)	4.9
Euro Govt. Bonds	4.0	5.0 (+1%)	10.7
U.S. Equity	5.8	4.8 (-1%)	15.4
Euro Equity	1.6	2.6 (+1%)	18.5
Japan Equity	1.1	2.1 (+1%)	15.7
Pacific ex-Japan Equity	3.3	4.3 (+1%)	20.1
Emerging Markets Equity	4.2	5.2 (+1%)	21.1
Alternatives	2.4	3.4 (+1%)	4.0

Table 3.2: Annualized expected returns and risk after a 1% shift in each asset class's original expected return given in Table 3.1.

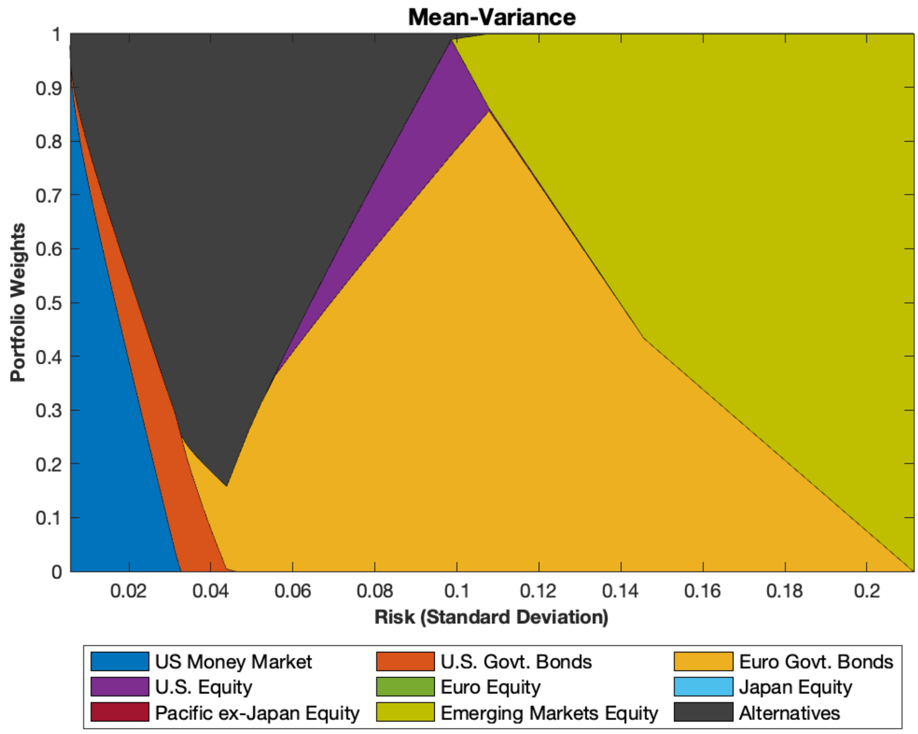

Figure 3.2: Mean–variance frontier illustrating how a 1% shift in expected returns affects portfolio weights and risk.

3.1.3 Black-Litterman (No views)

In this section, we employ the same dataset and annualized covariance matrix used in the Markowitz procedure, ensuring a consistent basis for comparison. However, rather than relying on sample-based historical means, we calculate the Black–Litterman *implied returns* by performing a reverse optimization on the market weights. Because we are operating in a setting without incorporating views, these implied returns become our posterior expected returns. Likewise, the posterior covariance matrix remains identical to the original Σ. As a result, the main difference from Markowitz is the shift from sample-mean estimates to market-implied returns, while all other inputs (e.g., covariances, constraints) remain the same.

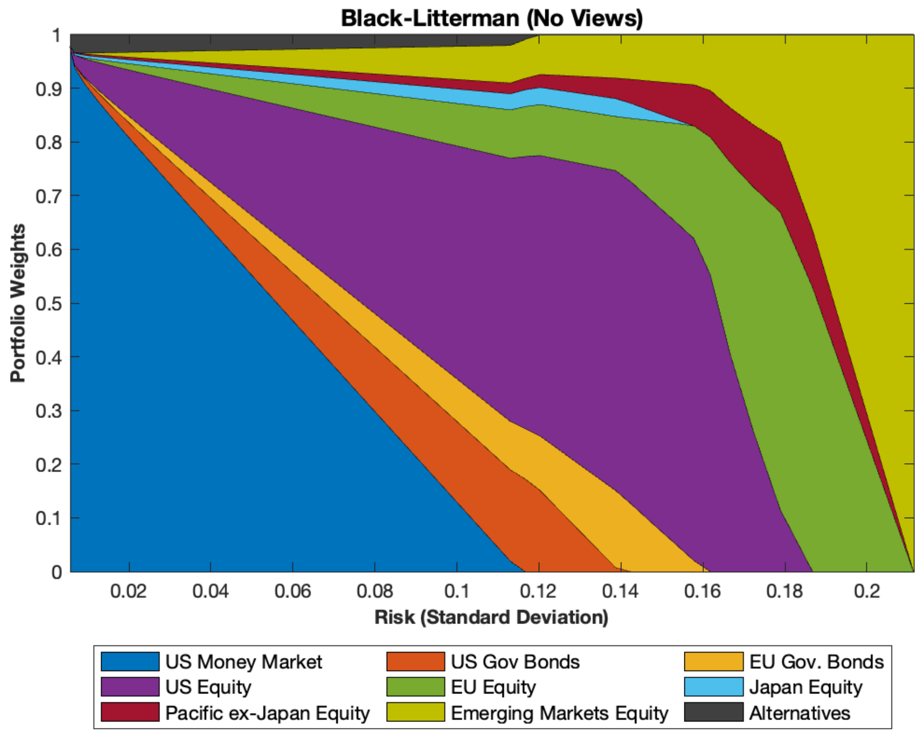

Figure 3.3: Black–Litterman (No views): Mean–variance frontier illustrating how equilibrium implied returns determine portfolio weights across varying risk levels.

Unlike previous Markowitz allocations, which produced abrupt corner solutions, we can immediately see that in Figure 3.3 the progression appears smoother, though still featuring concentration in a few asset classes at the extreme ends of the profile. This suggests that relying on market-implied returns generates a baseline that is more stable.

For investors primarily seeking a passive strategy—one that aligns allocations with broad market exposures—this equilibrium-based approach can offer a reasonable starting point and help maintain a stable, diversified portfolio. However, many portfolio managers require an active stance, introducing their own perspectives on market segments, asset classes, or broader economic conditions. In the absence of explicit views, this approach cannot capture such subjective beliefs, making it less effective in settings where investors wish to tilt allocations based on forward-looking insights. In the next section, we will therefore explore the full Black–Litterman model with views, illustrating how the incorporation of subjective opinions can further shift the posterior returns away from the market-implied baseline and reshape the resulting optimal portfolio allocations.

3.1.4 Black-Litterman Model (With Views)

In this section, we extend the Black–Litterman framework by incorporating specific investor beliefs about the future performance of certain assets. Instead of relying solely on equilibrium returns, a set of view matrices (**P**) and view returns (**Q**) adjusts the implied returns to produce a new, *posterior* distribution of expected returns. Compared to the previous specification of the model, the resulting portfolio allocations can differ significantly, depending on how strongly the investor holds these convictions (the choice of Ω and τ). In Table 3.3, we specify a set of arbitrary investor views for the Black–Litterman model:

1. US Bonds **underperform** EU Bonds by 2% (relative view)
2. EU Equity **overperforms** US Equity by 4% (relative view)
3. Alternatives **expected return** equal to 5% (absolute view)
4. **No views** on remaining 3 asset classes

Table 3.3: Arbitrary investor views for the Black–Litterman model, specifying both relative and absolute expectations.

In the stacked area chart below (Figure 3.4) , we see how incorporating explicit views shifts the allocation compared to the market equilibrium weights of the previous section. With $\tau = 0.05$ and 5 percent uncertainty in each view (Ω), the stacked area chart shows how these inputs reshape the efficient frontier relative to the market-implied baseline. US Equity is weighted notably less than EU Equity at mid-range risk levels, consistent with the stated underperformance view, while Alternatives receive a boost due to their absolute return target. Although the model still transitions toward riskier assets at higher volatility levels, the presence of explicit views clearly alters the portfolio composition by giving an extra push to EU markets and Alternatives at various points along the frontier. Overall, the chart illustrates how even modestly confident views (i.e., 5 percent uncertainty) can produce allocations that deviate from the equilibrium baseline, reflecting the model's ability to incorporate the investor's specific forward-looking perspectives into the final portfolio.

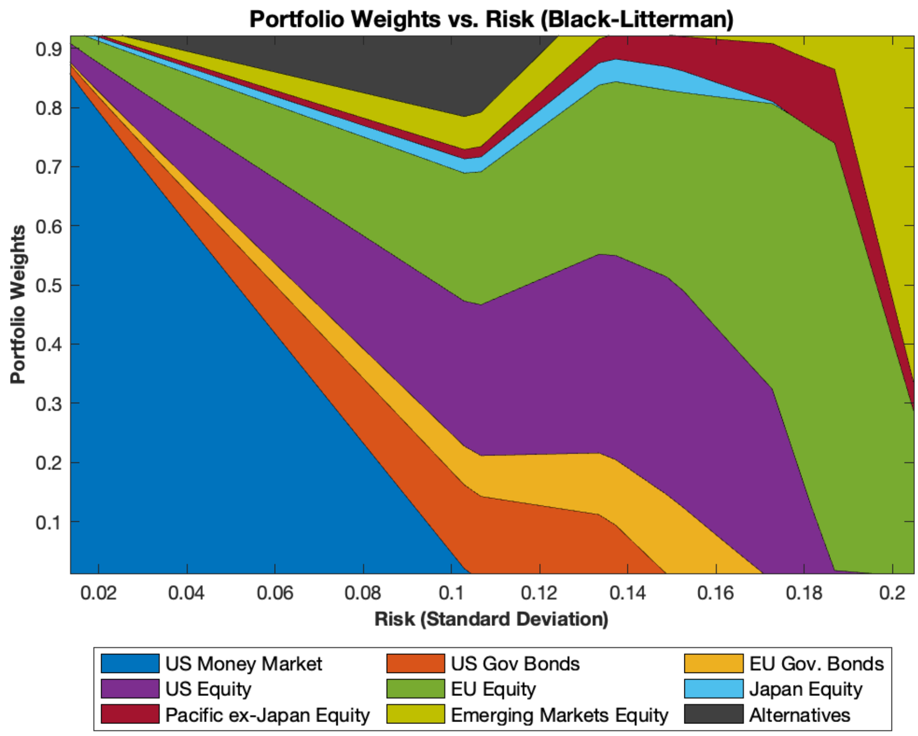

Figure 3.4: Black–Litterman Model: Portfolio weights vs. risk, illustrating the impact of the investor views from Table 3.3. The parameter values are $\tau = 0.05$ and $\Omega = 0.05$.

3.1.5 Parameter Sensitivity: Varying τ

Starting from the BL model under the views assigned in Table 3.3, Figure 3.5 shows how portfolio allocations evolve as τ increases from very small values (10^{-10})—where the model heavily favors the market-implied equilibrium—to extremely large values (10^9)—where the assigned views dominate. At lower τ values, the efficient frontier for each risk level looks identical to the market implied returns case, reflecting no tilt from the investor views. In contrast, at higher τ, the portfolio weights increasingly converge toward the allocations implied by the views, often culminating in more pronounced shifts. In particular, we see additional emphasis on EU Equity (which is favored over US Equity) and a stronger commitment to Alternatives at a 5 percent expected return.

Table 3.4 underscores this transition by reporting the *posterior returns* under two extreme τ values. For $\tau = 10^{-10}$, posterior estimates are equivalent to the implied returns, while at $\tau = 10^9$, they effectively align with the view vector.

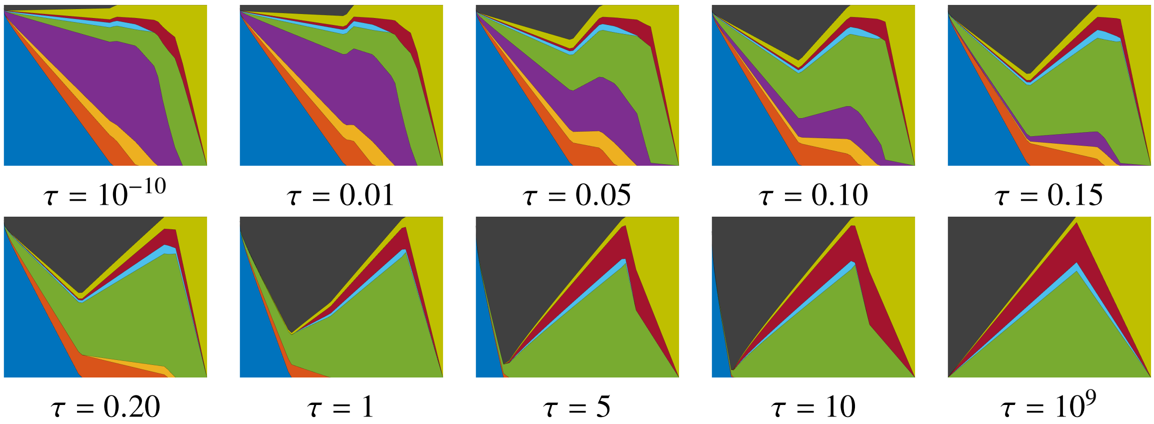

Figure 3.5: Black-Litterman Portfolio Allocations for Varying τ Values

Asset	**Posterior Return** ($\tau = 10^{-10}$)	**Posterior Return** ($\tau = 10^9$)
US Bonds	0%	0.1%
EU Bonds	1.8%	2.1%
US Equity	5.0%	14.6%
EU Equity	5.8%	18.9%
Alternatives	1.1%	5.0%

Table 3.4: Convergence of Posterior Returns from Equilibrium to Views as τ Increases

3.1.6 Parameter Sensitivity: Varying Ω

Once again, starting from the model under the views assigned in Table 3.3, Figure 3.6 illustrates how changing Ω—the view-confidence matrix—affects the final allocations. With a *small* Ω (left panel), the model effectively places *high* confidence in the views, resulting in a portfolio allocation that tilts vividly away from the market prior and instead resembles a mean–variance style concentration. This is similar to setting τ very low: the model relies more on user-defined views and less on the equilibrium baseline.

By contrast, a *large* Ω (right panel) signals *low* confidence in those views, causing the allocations to gravitate back toward the original market-implied weights. Even though views are specified, the high degree of uncertainty prevents them from substantially altering the equilibrium returns.

Thus, for practitioners, selecting an appropriate Ω is crucial. Over-confidence (too small Ω) can replicate the sensitivity issues of classical mean–variance, whereas excessive caution (too large Ω) may leave the portfolio virtually unchanged from the market benchmark.

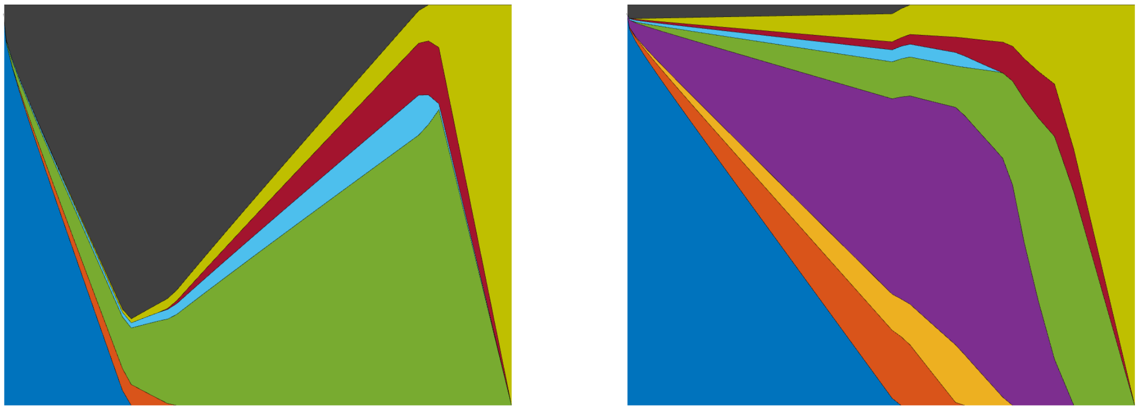

(a) Portfolio Allocation with $\Omega = 1\%$ (b) Portfolio Allocation with $\Omega = 50\%$

Figure 3.6: Effect of Different Ω Values on Portfolio Allocation

3.1.7 Key Takeaways and Practical Observations

Overall, incorporating views through Black–Litterman demonstrates both enhanced flexibility and the potential for *overreach* if parameters like τ and Ω are pushed to extremes. As shown in Figures 3.5 and 3.6, a too-large τ or too-small Ω, respectively, can reproduce some of the corner-solution behaviors observed in the classical mean–variance framework, while the opposite case may render investor insights nearly irrelevant.

The main advantage here is the model's ability to *blend* equilibrium market weights with subjective beliefs, offering a more intuitive balance than either pure Markowitz or the market prior alone. However, the final outcome depends heavily on the user's confidence—too much conviction in personal forecasts can override diversification benefits, while too little confidence leaves one simply "hugging the index".

3.2 Backtesting: Buy-and-Hold strategy

3.2.1 Introduction and Setup

This subsection presents our single-period, *no-rebalancing* backtest, in which each portfolio is formed once at the start of the out-of-sample window and held for the entire horizon (approximately five years). We assess three models: (1) Markowitz (Mean–Variance) using in-sample estimates, (2) Black–Litterman (Market-Implied) relying solely on reverse-optimized returns, and (3) Black–Litterman incorporating investor forecasts that either align with or deviate from realized out-of-sample returns (optimistic or pessimistic scenario).

We include the "market prior" version of Black–Litterman not only to isolate the impact of explicit views, but also because practitioners often treat equilibrium allocations as a baseline or benchmark. Comparing all three allows us to determine whether adding moderately accurate (or inaccurate) views offers a significant advantage—or introduces drawbacks—relative to an equilibrium-based approach or a traditional Markowitz portfolio. In the results that follow, we will examine:

1. A set of representative portfolios from the Monte Carlo (MC) runs, to compare side by side with Markowitz and the benchmark.

2. The *full distribution* of outcomes from 1,000 MC draws, analyzing percentiles and histograms of risk–return metrics to capture the variability in performance across different possible view realizations.

3.2.2 Out-of-Sample Asset Class Performance

Before turning to the portfolio allocations in the two scenarios, Figure 3.7 displays the out-of-sample returns of each asset pair that received a *relative* view, observed over the full five-year horizon. It is important to notice that our baseline views are derived from the *first 12 months* of this window—where short-term performance differences do not always align with the longer-term trends observed at the end of the period. For instance, Emerging Markets initially *outperform* Alternatives around the 12-month mark, yet lag behind them by year five. A similar reversal is evident for U.S. vs. E.U. Government Bonds, whereas Japan vs. Pacific ex-Japan Equity and U.S. vs. E.U. Equity appear to align with the correct directional view at 12 months but *underestimate* the magnitude of divergence over the full horizon.

Since portfolio weights are derived based on these 12-month views but evaluated over the entire five-year out-of-sample period, their predictive power is inherently limited. As a

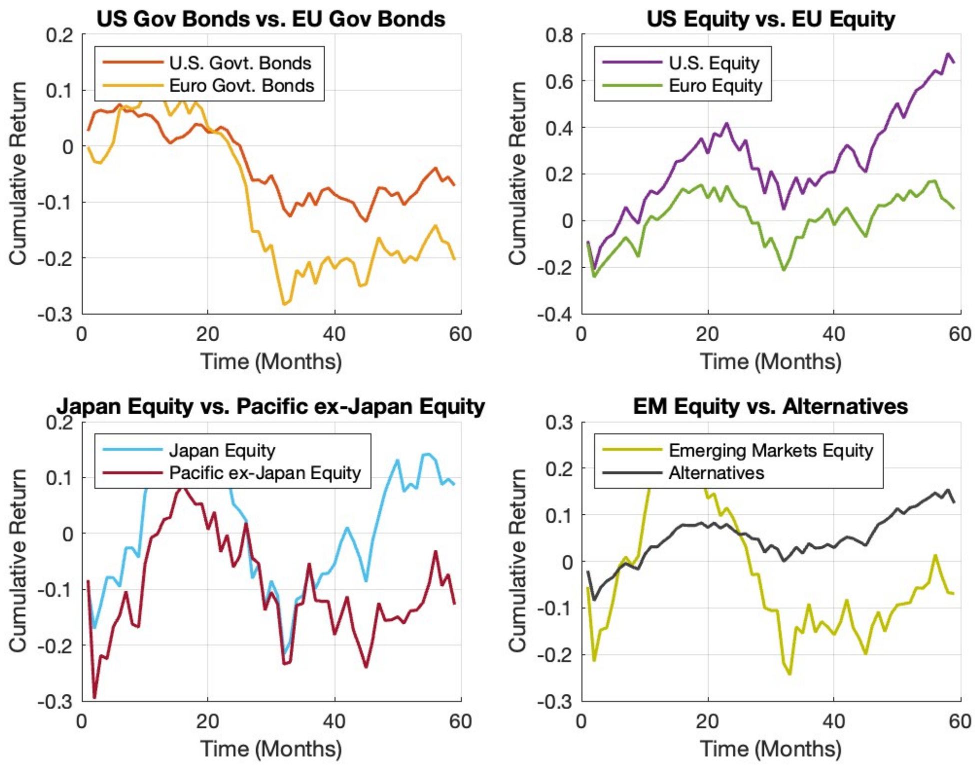

Figure 3.7: Out-of-sample return trajectories for the assets subject to relative views over the five-year horizon.

result, even when the views appear correctly specified in hindsight—given they are drawn from realized one-year-ahead returns—their ”informativeness” for long-term allocation decisions is weak. Consequently, the out-of-sample performance of these portfolios remains modest and, as we will see, will even underperform the benchmark.

3.2.3 Portfolio Allocations - Optimistic Scenario

Table 3.5 summarizes the baseline investor views for each pair. Note that, in practice, the actual views used in the model are drawn via Monte Carlo simulations around these baseline figures, with the standard deviations of each view depending on the volatility of its underlying assets. Specifically, if a relative view involves two assets, we take the average of their in-sample volatilities to scale the uncertainty in that view. Thus, the allocations we see represent an average of 1,000 simulated draws rather than a single deterministic set of views. This approach better captures how high-volatility assets produce views with wider dispersion, while stable assets have more tightly clustered forecasts.

Tables 3.6, 3.7, and 3.8 report the low, mid, and high-volatility portfolio allocations for each model, with Black-Litterman based on the views in Table 3.5 . In the low-risk

1. US Money Market **expected return** equal to 0.6% (absolute view)
2. US Bonds **underperform** EU Bonds by 6.6% (relative view)
3. US Equity **overperforms** EU Equity by 9.9% (relative view)
4. Japan Equity **overperforms** Pacific ex-Japan Equity by 6.2% (relative view)
5. Emerging Markets Equity **overperforms** Alternatives by 18.9% (relative view)

Table 3.5: Optimistic Views (Baseline)

case, Mean–Variance heavily favors U.S. Government Bonds and Money Market; the BL (No Views) allocates more to Money Market and some U.S. Equity, while BL (With Views) re-directs a notable share away from North American companies in favor of the Japan equity market, consistent with that view's outperformance. A similar pattern holds at medium volatility, where the BL (With Views) model more aggressively weights Japan Equity and Emerging Markets compared to the other two. At high volatility, the BL (With Views) allocation leans even further into EM Equity, reflecting its strong *assumed* out-of-sample potential—an expectation that, in hindsight, does not fully materialize over the entire evaluation period.

Asset Class	**MV**	**BL (No Views)**	**BL (Views)**
US Money Market	29	49	63
US Gov Bonds	60	9	2
EU Gov Bonds	1	5	3
US Equity	11	26	13
EU Equity	0	4	1
Japan Equity	0	2	12
Pacific ex-Japan Equity	0	1	1
Emerging Markets Equity	0	3	5
Alternatives	0	2	0

Table 3.6: Portfolio Weights (Low Volatility)

Asset Class	MV	BL (No Views)	BL (Views)
US Money Market	0	0	27
US Gov Bonds	61	17	4
EU Gov Bonds	33	9	6
US Equity	5	50	24
EU Equity	0	9	3
Japan Equity	0	3	23
Pacific ex-Japan Equity	0	2	3
Emerging Markets Equity	2	7	10
Alternatives	0	2	0

Table 3.7: Portfolio Weights (Mid Volatility)

Asset Class	MV	BL (No Views)	BL (Views)
US Money Market	0	0	5
US Gov Bonds	33	0	2
EU Gov Bonds	67	0	3
US Equity	0	41	24
EU Equity	0	35	6
Japan Equity	0	0	31
Pacific ex-Japan Equity	0	2	5
Emerging Markets Equity	0	22	24
Alternatives	0	0	0

Table 3.8: Portfolio Weights (High Volatility)

3.2.4 Metrics and Cumulative Returns - Optimistic Scenario

To assess each portfolio's longer-term behavior, we compare both quantitative metrics (Table 3.9) and cumulative return trajectories (Figure 3.8) across the low, mid, and high volatility portfolios.

Performance Metrics

From Table 3.9, it is evident that Markowitz (Mean–Variance) lags behind in most categories. Its total return (TR) and mean return (MR) are notably lower or even negative in the mid- and high-volatility portfolios. This underperformance stems primarily from *poor diversification*: having derived expected returns purely from the in-sample average, the model often concentrates heavily on a few assets, leaving it more vulnerable when those picks do not materialize out-of-sample.

By contrast, the benchmark (BL (No Views)) shows more balanced allocations and delivers higher Sharpe Ratios (SR) in each risk segment. The BL (Views) model performs similarly well, especially at low and mid volatily—though the difference in SR relative to the benchmark is slightly lower. This outcome makes sense given that the short-term views used by the model do capture some early trends; however, since the evaluation runs for five years, those initial insights lose power over the latter part of the period.

Low-Volatility Portfolio				
Model	**MR**	**TR**	**VOL**	**SR**
Markowitz	1.2	5.5	4.8	0.25
BL (No Views)	4.3	22.0	7.0	0.61
BL (Views)	3.3	17.1	5.7	0.59

Mid-Volatility Portfolio				
Model	**MR**	**TR**	**VOL**	**SR**
Markowitz	-1.6	-8.8	7.6	-0.21
BL (No Views)	5.9	27.7	13.8	0.43
BL (Views)	4.1	19.1	11.2	0.36

High-Volatility Portfolio				
Model	**MR**	**TR**	**VOL**	**SR**
Markowitz	-3.1	-15.9	9.0	-0.35
BL (No Views)	6.0	24.6	17.6	0.34
BL (Views)	4.1	15.8	15.5	0.26

Table 3.9: Performance metrics for low, mid, and high volatility portfolios in the optimistic scenario. (MR = Mean Return, TR = Total Return, VOL = Volatility, SR = Sharpe Ratio. MR, TR, and VOL are in percentage form.)

Cumulative Returns

Figure 3.8 shows the same story in graphical form:

- **Low Volatility**: Markowitz (blue line) remains well below the other two. BL (No Views) in green and BL (With Views) in red track each other closely and both finish with higher cumulative returns.
- **Mid Volatility**: Markowitz dips significantly under zero before partially recovering. Both BL strategies trend higher, showing similar trajectories and metrics.
- **High Volatility**: All three exhibit greater fluctuations, but Markowitz endures prolonged negative performance. BL (No views) and BL (With Views) oscillate

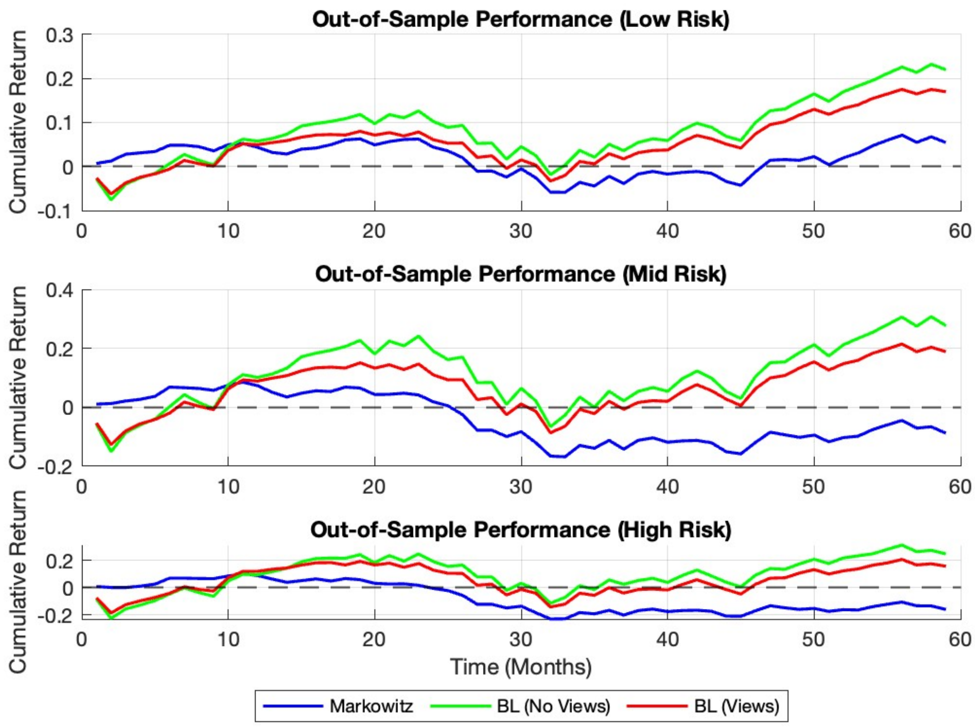

Figure 3.8: Out-of-sample return trajectories for low, mid, and high-volatility portfolios in the optimistic scenario.

around similar trajectories, with the latter gaining an early edge then converging toward the benchmark later on.

Overall, these results underscore that while neither Black–Litterman variant greatly outdistances the other in five-year performance, they both benefit from more balanced allocations than Markowitz. The incorporation of views may seize extra upside if early-period trends persist, yet its advantage narrows if those trends fade. Consequently, while the short-term forecasts can provide a mild boost, the equilibrium-based baseline remains a strong competitor over longer horizons.

3.2.5 Performance Percentiles - Optimistic Scenario

Before examining the the performance results of the Monte Carlo simulation, we briefly revisit how the views themselves vary. As shown in Table 3.10, each baseline forecast (Q_{baseline}) is associated with a view-specific standard deviation σ, derived from the in-sample volatility of the corresponding assets. For instance, the absolute view on U.S. Money Market has a relatively small standard deviation, $\sigma \approx 0.53\%$, whereas the more volatile equity pairs exhibit standard deviations exceeding 10–15%.

View	σ (%)	Q_{baseline} (%)	**68% Range**
1	0.53	0.6	[0%, 1%]
2	7.51	-6.6	[-14%, 1%]
3	16.4	9.9	[-7%, 26%]
4	17.6	6.2	[-11%, 24%]
5	12.7	18.9	[6%, 32%]

Table 3.10: Monte Carlo view perturbations around the baseline. The standard deviations (σ) for each of the 5 views are shown in the second column. The baseline views (Q_{baseline}) in the third column remain fixed, while the 68% range is computed as $[\, Q_{\text{baseline}} - \sigma,\ Q_{\text{baseline}} + \sigma]$ and each bound is approximated to the nearest integer percentage.

Table 3.11 reports the distribution of key performance metrics—annualized mean return (MR), total return (TR), and volatility (VOL)—for the Black–Litterman (BL) portfolios at three different risk targets, based on 1,000 MC runs. A notable finding is that both *mean return* and *total return* show a considerable spread across percentiles, particularly at the mid- and high-volatility levels. This stems from often wide changes in the views and from the BL model's sensitivity to these forecasts. When the views point in the right direction, the model can allocate aggressively enough to generate good returns, but if the signals deviate or conflict with actual out-of-sample patterns, the portfolio generally underperforms.

Low-Volatility Portfolio Percentiles				
Percentile	**MR**	**TR**	**VOL**	**SR**
5	1.7	7.9	4.1	0.3
25	2.6	13.2	4.9	0.5
50	3.2	16.3	5.8	0.6
75	4.1	21.1	6.3	0.7
95	5.3	27.9	7.1	0.8

Mid-Volatility Portfolio Percentiles				
Percentile	**MR**	**TR**	**VOL**	**SR**
5	0.9	1.2	8.0	0.1
25	2.7	11.2	9.9	0.3
50	3.9	17.5	11.5	0.4
75	5.6	27.0	12.5	0.5
95	7.9	40.7	13.8	0.6

High-Volatility Portfolio Percentiles				
Percentile	**MR**	**TR**	**VOL**	**SR**
5	0.6	-2.6	12.3	0.0
25	2.1	4.4	14.5	0.1
50	3.4	11.9	15.9	0.2
75	6.1	25.8	16.9	0.4
95	8.5	41.8	18.0	0.5

Table 3.11: Percentile performance metrics of Monte Carlo simulation on optimistic views in the BL model for low, mid, and high volatility portfolios. (MR = Mean Return, TR = Total Return, VOL = Volatility, SR = Sharpe Ratio. MR, TR, and VOL are expressed in percentage form)

Comparing these outcomes to the single-point results of the benchmark and Markowitz (Table 3.9):

- **Versus Markowitz:** The BL distribution *dominates* Markowitz. Examining the distributions at each risk level alongside the single Markowitz portfolio, it is easy to see that, in all cases, even the 5th percentile of BL exceeds Markowitz's result in both Total Return (TR) and Sharpe Ratio (SR). In the low volatility setting, Markowitz yields -8.8% TR, whereas the 5th percentile for BL is +1.2%. Hence, virtually all BL draws surpass Markowitz in total return. A similar pattern holds for the SR. For mid volatility, Markowitz achieves 5.5% TR, while the 5th percentile of BL is 7.9%—again exceeding Markowitz, implying almost every BL draws outperform Markowitz. Finally, for the high volatility scenario, Markowitz's total return is -15.9%, but BL's 5th percentile stands at -2.6%. Similar conclusion to the first two cases can be drawn.

  These results reflect the fact that, as soon as the Black–Litterman baseline views are even moderately "directionally correct," the large majority (or entirety) of Monte Carlo draws will yield a more favorable risk–return profile than the single-point Markowitz solution.

- **Versus the Benchmark:** The benchmark achieves higher total returns than both Markowitz and the BL model in most simulations. Based on these percentile tables, the BL portfolio at mid-risk surpasses the benchmark's performance in only the top 25% of simulation (judging by the point where BL's percentile meets or exceeds the benchmark's single-point total return). A similar pattern holds for the low-risk and high-risk configurations. Hence, although the BL strategy can occasionally match or exceed the benchmark, it does so few simulations.

3.2.6 Performance Histograms - Optimistic Scenario

Figures 3.9, 3.10, and 3.11 depict histograms of the four main performance metrics (mean return, total return, volatility, and Sharpe ratio) for each risk level. Each figure compiles the results of 1,000 MC runs, illustrating where BL outcomes concentrate and how far they can deviate in the tails.

- **Mean Return and Total Return:** In the low-risk portfolio, mean returns cluster in the 2–4% range, peaking near 3%. The distribution, however, has a long right tail, reaching beyond 5–6% for a small number of runs. Mid-risk and high-risk histograms exhibit even broader spreads, occasionally pushing mean returns above

8–10%. Correspondingly, total return distributions shift rightward but also reveal a left tail that can dip near zero or below for mid/high-risk portfolios.

- **Volatility:** The low-risk portfolio's volatility is mostly confined to the 4–6% region, whereas mid-risk extends from about 8% to 13%. At the high-risk level, volatilities stretch beyond 15–18%.

- **Sharpe Ratio:** For the low-risk portfolio, Sharpe ratios often cluster around 0.4–0.7, with a noticeable concentration near 0.6. A slim but non-negligible portion of runs achieves ratios above 0.8. In contrast, mid-risk and high-risk histograms display lower median Sharpe ratios (0.3–0.4 for mid, around 0.2 for high) and wider spreads. At high volatility, the left tail can drop below 0.0 or even –0.2, indicating occasional negative returns.

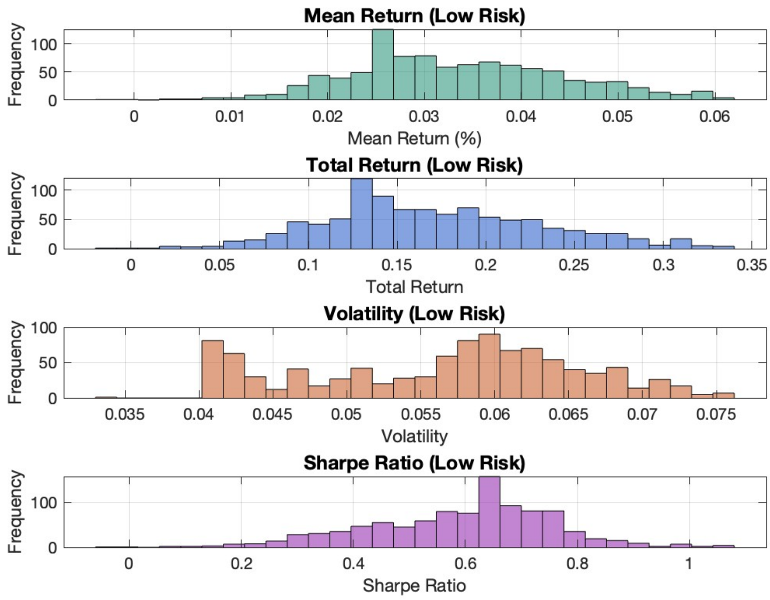

Figure 3.9: Histogram of performance metrics for the low-risk portfolio under the optimistic scenario (BL model with views, 1,000 MC runs).

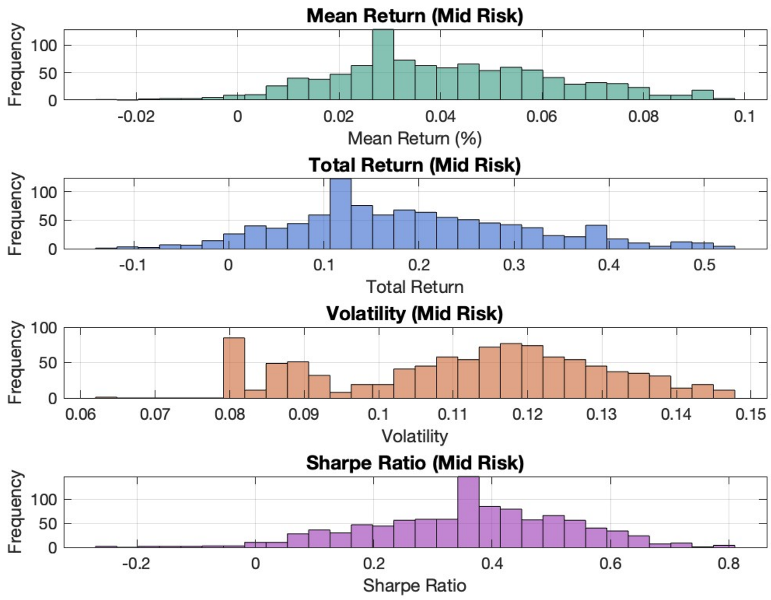

Figure 3.10: Histogram of performance metrics for the mid-risk portfolio under the optimistic scenario (BL model with views, 1,000 MC runs).

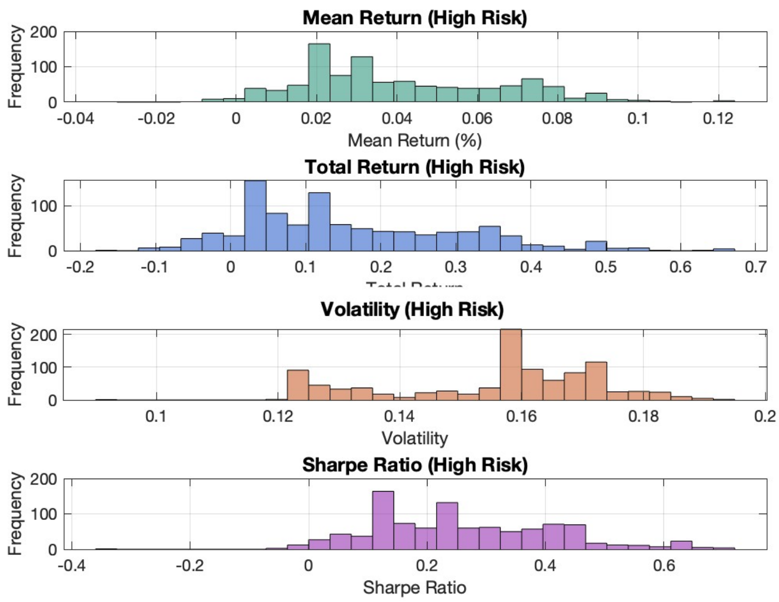

Figure 3.11: Histogram of performance metrics for the high-risk portfolio under the optimistic scenario (BL model with views, 1,000 MC runs).

3.2.7 Portfolio Allocations - Pessimistic Scenario

Table 3.12 lists the *reversed* views used in this pessimistic case, simply inverting the one-year-ahead returns from the optimistic scenario. Hence, where we previously posited overperformance, we now assume underperformance relative to the same reference asset classes, and vice versa.

1. US Money Market **expected return** equal to **-0.6%** (absolute view)
2. US Bonds **overperform** EU Bonds by 6.6% (relative view)
3. US Equity **underperforms** EU Equity by 9.9% (relative view)
4. Japan Equity **underperforms** Pacific ex-Japan Equity by 6.2% (relative view)
5. Emerging Markets Equity **underperforms** Alternatives by 18.9% (relative view)

Table 3.12: Pessimistic Views (Baseline)

Table 3.13 displays how these pessimistic views shape the BL allocations at low, mid, and high risk. In the low volatility setting, unlike the optimistic scenario, the portfolios tilt more heavily toward U.S. Bonds, while equities and emerging markets receive reduced allocations. Money Markets still hold the predominant share. At high volatility, the model leans even more into the bond component, now including EU bonds as well, reflecting a negative stance on equities and other previously favored segments. This strong pivot underscores Black–Litterman's responsiveness to its view inputs: reversing short-term performance assumptions flips the portfolio composition. Markowitz and benchmark allocations remain the same since only the Black–Litterman model incorporates the reversed forecasts.

Asset Class	BL (Low Risk)	BL (Mid Risk)	BL (High Risk)
US Money Market	57	27	12
US Gov Bonds	28	42	42
EU Gov Bonds	4	10	12
US Equity	3	6	7
EU Equity	3	5	11
Japan Equity	3	6	9
Pacific ex-Japan Equity	2	4	7
Emerging Markets Equity	0	0	0
Alternatives	0	0	0

Table 3.13: Black-Litterman Portfolio Weights (Pessimistic Views)

3.2.8 Metrics and Cumulative Returns - Pessimistic Scenario

Just as in the optimistic case, we now analyze how each model performs over five years, but under pessimistic views. Table 3.14 reports the out-of-sample statistics for the low, mid, and high-volatility portfolios, while Figure 3.12 shows their cumulative return trajectories.

Performance Metrics

Low-Volatility Portfolio				
Model	**MR**	**TR**	**VOL**	**SR**
Markowitz	1.2	5.5	4.8	0.25
BL (No Views)	4.3	22.0	7.0	0.61
BL (Views)	1.4	6.9	3.8	0.61

Mid-Volatility Portfolio				
Model	**MR**	**TR**	**VOL**	**SR**
Markowitz	-1.6	-8.8	7.6	-0.21
BL (No Views)	5.9	27.7	13.8	0.43
BL (Views)	0.7	2.4	7.1	0.12

High-Volatility Portfolio				
Model	**MR**	**TR**	**VOL**	**SR**
Markowitz	-3.1	-15.9	9.0	-0.35
BL (No Views)	6.0	24.6	17.6	0.34
BL (Views)	0.7	1.1	10.0	0.01

Table 3.14: Performance metrics for low, mid, and high volatility portfolios in the pessimistic scenario. (MR = Mean Return, TR = Total Return, VOL = Volatility, SR = Sharpe Ratio. MR, TR, and VOL are in percentage form)

With bonds and money market dominating the Black–Litterman allocations, the resulting portfolios exhibit *lower volatility* across all risk categories compared to the optimistic scenario. However, these conservative tilts also lead to more modest returns. In the low-risk portfolio, its total return (TR) of 6.9% is far below the benchmark's 22.0%, but yielding a similar Sharpe Ratio that is able to compensate the decreased returns in favor of lower risk. Interestingly, though, even with these "incorrect" views, the BL solution still outperforms Markowitz in all return metrics, highlighting how Bayesian blending can mitigate pure estimation errors. A similar pattern emerges in the mid- and high-volatility segments: while Markowitz performs poorly overall, the BL model again surpasses it in all metrics.

Cumulative Returns.

Figure 3.12 illustrates these findings visually. The BL (Views) line in red remains flatter than in the optimistic scenario and consistently trails the benchmark (green). Though it is less volatile than Markowitz (blue), it also fails to capture substantial upside, as the bond-heavy composition drags down potential gains. Nevertheless, it retains an edge over Markowitz's heavily concentrated approach, confirming that even wrong views can produce more robust outcomes than reliance on in-sample means and standard MV optimization. By the end of the horizon, the market-implied benchmark allocation again stands ahead on risk–return measures.

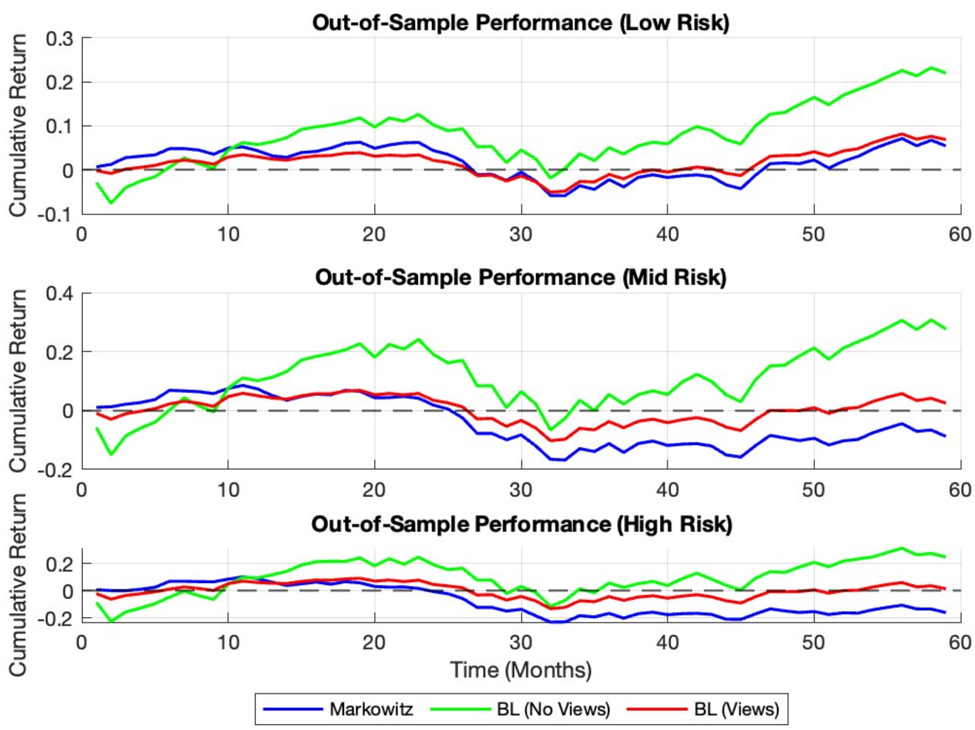

Figure 3.12: Out-of-sample return trajectories for low, mid, and high-volatility portfolios in the pessimistic scenario.

3.2.9 Performance Percentiles - Pessimistic Scenario

Table 3.15 presents the Black–Litterman percentiles under pessimistic views, showing how mean return (MR), total return (TR), and volatility (VOL) behave across the 1,000 Monte Carlo (MC) simulations. At the low-volatility setting, the 5^{th} percentile of total return (−2.8%) is considerably below Markowitz's single-point outcome (+5.5%) reported in 3.14, but by the 50^{th} percentile, BL climbs to 8.5%, surpassing Markowitz. Notably, more than half of all BL simulations in this scenario still generate higher total returns than the single-point Markowitz solution—a testament to how even wrong, when integrated in a Bayesian manner, can produce allocations that typically improve upon classical mean-variance. Still looking at the low volatility setting, only about 5% of BL draws, however, exceed +18% total return, still below the benchmark's +22.0%.

A similar pattern appears in mid- and high-volatility portfolios. Most BL outcomes rapidly overtake Markowitz's strongly negative results (−8.8% and −15.9% respectively)—already by the 25^{th} percentile—demonstrating that the BL model's distribution outperforms Markowitz in a majority of draws, despite the views being pessimistic. Nevertheless, BL rarely matches the benchmark's higher total returns. Consequently, BL remains vulnerable to underperformance at the lower tail but still tends to surpass Markowitz in nearly all simulations. Only a small fraction of draws rival the benchmark.

Low-Volatility Portfolio Percentiles				
Percentile	**MR**	**TR**	**VOL**	**SR**
5	-0.3	-2.8	1.5	0.0
25	0.6	2.4	1.5	0.1
50	1.7	8.5	3.7	0.6
75	1.7	8.8	6.0	1.1
95	3.6	18.1	7.2	1.2

Mid-Volatility Portfolio Percentiles				
Percentile	**MR**	**TR**	**VOL**	**SR**
5	-2.2	-11.4	3.0	-0.3
25	-0.3	-4.0	3.1	-0.0
50	0.7	3.0	7.2	0.2
75	0.8	3.2	11.0	0.2
95	5.9	29.5	12.3	0.5

High-Volatility Portfolio Percentiles				
Percentile	**MR**	**TR**	**VOL**	**SR**
5	-3.2	-16.3	4.5	-0.3
25	-0.3	-4.1	4.6	-0.1
50	-0.3	-2.0	9.2	-0.1
75	1.7	2.9	15.2	0.1
95	9.1	48.1	17.1	0.6

Table 3.15: Percentile performance metrics of Monte Carlo simulation on pessimistic views in the BL model for low, mid, and high volatility portfolios. (MR = Mean Return, TR = Total Return, VOL = Volatility, SR = Sharpe Ratio. MR, TR, and VOL are expressed in percentage form)

3.2.10 Performance Histograms - Pessimistic Scenario

Figures 3.13, 3.14, and 3.15 illustrate the distribution of mean return, total return, volatility, and Sharpe ratio for each risk level under pessimistic views. A few key observations:

- **Yield near-zero or negative returns:** At low risk, for both MR and TR, a sizable portion of simulations yield near–0% or even negative returns, with distinct clusters appearing around 1%. Mid- and high-volatility results exhibit similarly skewed shapes, with left tails near 0 or lower TR. Still, a handful of "lucky" draws (especially in high-vol) push mean returns to around 9% (and total returns to around 50%), though these remain outliers.

- **Sharpe Ratio:** Across all three risk levels, the Sharpe ratio histograms show a predominantly modest or near-zero range, with occasional dips into negative territory (mid and high risk) and a few outliers attaining higher values (low risk). This indicates that, in many simulations, excess returns barely compensate for the associated risk, although certain runs do achieve more favorable outcomes. The distribution skews further downward for higher-volatility portfolios, suggesting that sustaining robust risk-adjusted performance can prove challenging under adverse conditions.

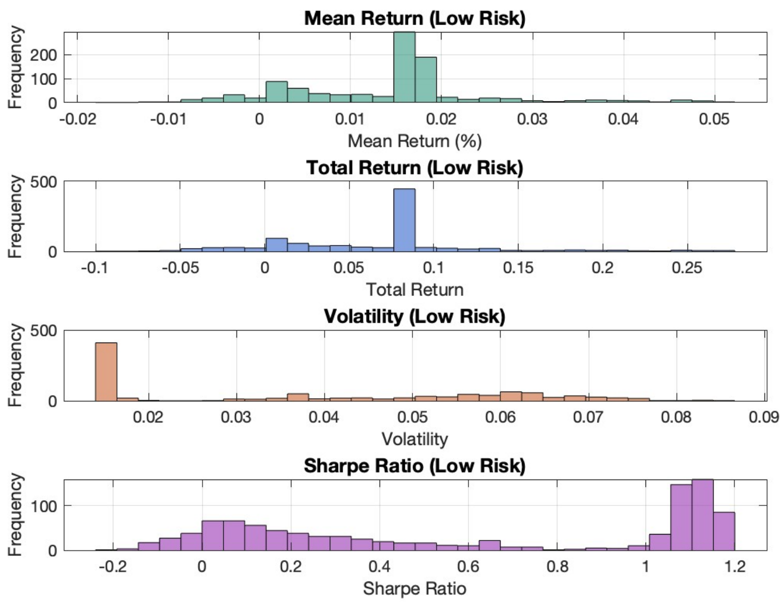

Figure 3.13: Histogram of performance metrics for the low-risk portfolio under the pessimistic scenario (BL model with views, 1,000 MC runs).

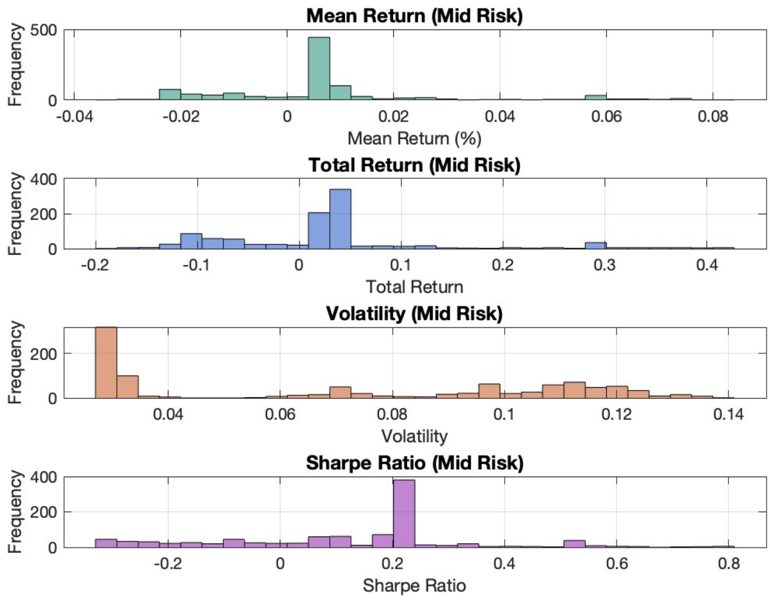

Figure 3.14: Histogram of performance metrics for the mid-risk portfolio under the pessimistic scenario (BL model with views, 1,000 MC runs).

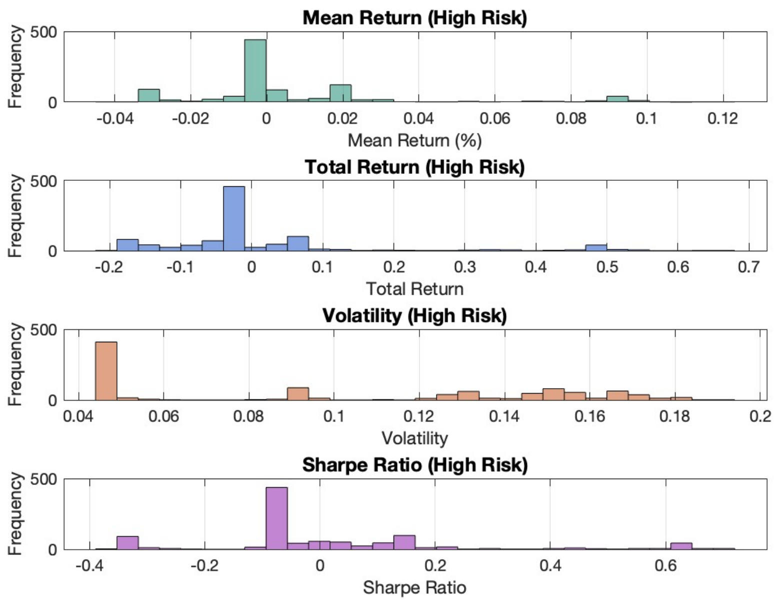

Figure 3.15: Histogram of performance metrics for the high-risk portfolio under the pessimistic scenario (BL model with views, 1,000 MC runs).

Overall, these results suggest that the Black–Litterman framework (whether *market-implied* or *with views*) generally surpasses Markowitz in terms of diversification and risk-adjusted returns. In the *optimistic* scenario, the short-term forecasts provide a mild performance boost—although one that fades if trends shift. Under *pessimistic* views, the BL (views) portfolios become more defensive and still outperform Markowitz in most metrics, but cannot match the market-implied benchmark. This highlights both the strengths and vulnerabilities of relying on short-horizon inputs for a multi- year horizon.

Looking ahead, the next step is to examine whether periodic rebalancing can mitigate or amplify these patterns. By regularly updating expected returns and views (instead of forming a single portfolio and holding it), investors may reduce drawdowns if forecasts are inaccurate or capitalize on prolonged trends if forecasts prove more accurate. The following sections explore this possibility, comparing a quarterly-rebalanced framework across the same three models.

3.3 Backtesting: Rebalancing Scenario

3.3.1 Introduction and Setup

The purpose of implementing a quarterly rebalancing framework is to capture changing market conditions more effectively than a single-period (no-rebalancing) approach. In practice, portfolio managers rarely maintain static allocations over extended horizons; they typically update positions to incorporate newly available information or to adapt to shifting risk–return profiles. By recalibrating portfolio weights every quarter, this approach mirrors real-world investment management, where strategic asset allocation (SAA) establishes a long-term baseline portfolio aligned with an investor's risk tolerance, while tactical asset allocation (TAA) allows for short-term adjustments in response to expected market conditions. This dynamic approach aims to mitigate risks stemming from outdated forecasts and capitalize on emerging trends in asset returns.

This quarterly rebalancing analysis directly extends the no-rebalancing results presented in the previous section. Whereas the single-period approach provided valuable insights into how the Markowitz and Black–Litterman (BL) models perform under fixed allocations, it did not address the reality of ongoing portfolio management. Accordingly, this section evaluates how dynamic adjustments—based on recurring Markowitz estimates or BL posterior updates—affect portfolio performance, particularly under scenarios where the BL views may be mostly correct or mostly incorrect over successive quarters.

3.3.2 Recap of Methodology and Scenario Definition

The quarterly rebalancing procedure follows an expanding-window design. At each quarter, the in-sample portion is extended to include all newly realized returns, and portfolio optimization is subsequently performed on this new data. Specifically, three versions of the model—Markowitz, Black–Litterman (No Views), and Black–Litterman (With Views)—are re-estimated every quarter to derive their respective efficient frontiers. For each of these frontiers, three representative portfolios are selected, corresponding to low (25%), mid (50%), and high (75%) fractions of the maximum risk. Throughout the process, all other parameter settings, including risk-aversion coefficients and view/prior uncertainty (τ and Ω), remain consistent with the earlier no-rebalancing framework.

Within this quarterly setup, two primary scenarios are defined to represent the potential accuracy of the Black–Litterman views:

1. **Mostly Correct (Optimistic)**: In approximately 75% of the quarters that make up the OOS window, the views align directionally with the realized return for the

upcoming quarter (they are set to the realised returns in the next quarter scaled by half), and a 25% chance it is inverted.

2. **Mostly Incorrect (Pessimistic)**: The proportions are reversed. Only 25% of the time does the view match the upcoming return's direction, with 75% of the forecasts pointing the wrong way.

These two scenarios offer a controlled way to assess whether the advantage of partial insight is amplified, neutralized, or offset by periodic portfolio revisions. In practical portfolio management, this mirrors how tactical asset allocation decisions, when executed with a high degree of accuracy, can enhance long-term performance—but when misguided, can erode the benefits of a well-diversified strategic allocation. For a detailed explanation of the methodology used, see Section 2.3.2.

3.3.3 Portfolio Allocations - Optimistic Scenario

It is helpful to recall the out-of-sample returns for each pair of assets, as previously shown in the no-rebalancing section and reported below (Figure 3.16). Here, such returns play an even greater role: the rebalancing approach uses them to incorporate the time-varying dynamics of asset performance, especially in the Black–Litterman (BL) model with views. Unlike the single-period method—where the views were locked in at the start—this quarterly updating allows the model to adjust allocations if the realized returns deviate from earlier expectations.

Figures 3.17, 3.18 and 3.19 present the average portfolio weights obtained from quarterly rebalancing across the out-of-sample (OOS) horizon under the optimistic view scenario. We display average allocations because each quarter yields a distinct optimized weight vector, making it impractical to show all 19 individual charts (corrisponding to the quarters in the testing window). However, the performance metrics (shown next) are based on the cumulative results of investing in each quarterly allocation sequentially, rather than simply computing the returns of these averaged portfolios.

For instance, in the BL (Views) portfolio, Japan Equity receives a higher weight on average than in the benchmark portfolio, reflecting that its relative view remains favorable in most quarters. Similarly, U.S. Government Bonds occupy a greater share than in the benchmark, driven by better relative performance versus other bond categories throughout much of the OOS window. While Japan equity was already favored by the BL model (compared to the benchmark) in the no-rebalancing framework, the more aggressive allocation toward the U.S. bond component now is driven mostly by the effects of quarterly rebalancing.

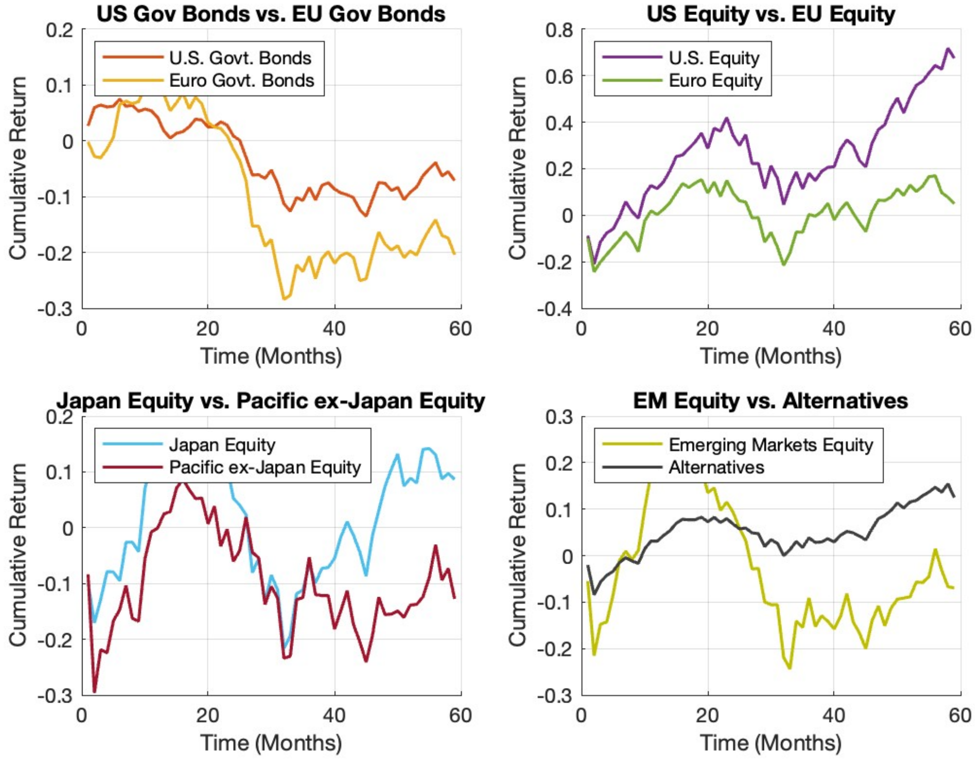

Figure 3.16: Out-of-sample return trajectories for the assets subject to relative views over the five-year horizon.

A clear distinction emerges among the three models. Markowitz, which relies on pure historical estimates, often shows higher sensitivity to recent return estimates, resulting in more concentrated bets. By contrast, the benchmark allocation maintains broader diversification, guided by an equilibrium baseline that helps avoid extreme tilts. Meanwhile, the BL portfolios incorporating views generally respond well to the partial-correct forecasts quarter by quarter, balancing opportunistic reallocations against the baseline ”anchor” that limits overconcentration.

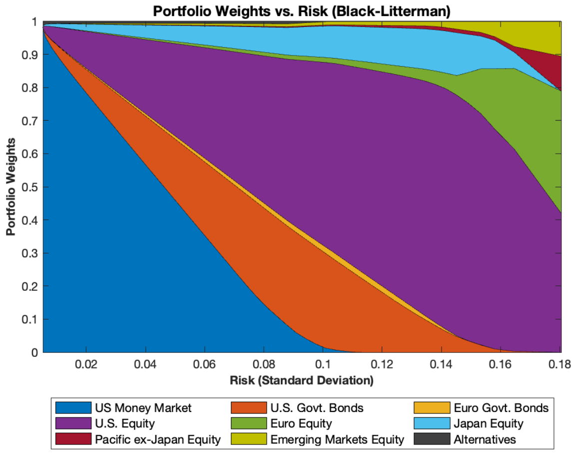

Figure 3.17: Portfolio weights under Black–Litterman with views in the optimistic scenario

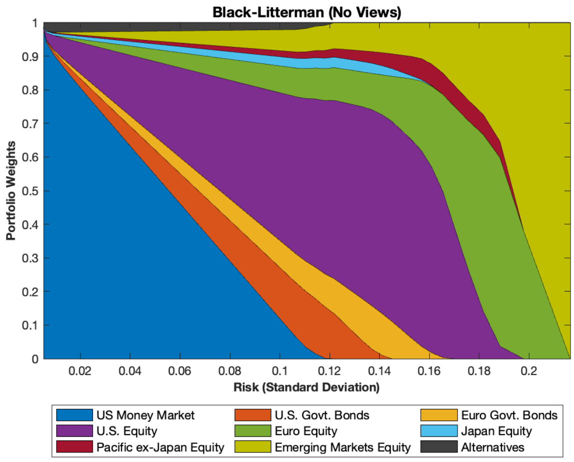

Figure 3.18: Portfolio weights under Black–Litterman (No Views) in the optimistic scenario

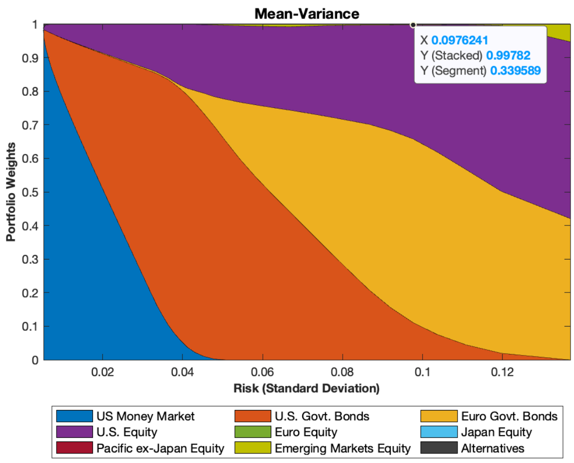

Figure 3.19: Portfolio weights under Markowitz in the optimistic scenario

3.3.4 Metrics and Cumulative Returns - Optimistic Scenario

Table 3.16 compares the three models at low, mid, and high risk under quarterly rebalancing with optimistic views.

Low-Volatility Portfolio				
Model	**MR**	**TR**	**VOL**	**SR**
Markowitz	1.2	4.5	7.1	0.16
BL (No Views)	4.2	20.9	6.4	0.66
BL (With Views)	4.4	21.8	6.1	0.71

Mid-Volatility Portfolio				
Model	**MR**	**TR**	**VOL**	**SR**
Markowitz	0.6	0.7	10.1	0.06
BL (No Views)	5.9	27.4	12.7	0.46
BL (With Views)	6.6	32.8	11.5	0.58

High-Volatility Portfolio				
Model	**MR**	**TR**	**VOL**	**SR**
Markowitz	0.7	-0.3	12.4	0.05
BL (No Views)	7.8	35.4	17.2	0.46
BL (With Views)	9.7	50.7	14.4	0.67

Table 3.16: Performance metrics for low, mid, and high volatility portfolios in the optimistic scenario. (MR = Mean Return, TR = Total Return, VOL = Volatility, SR = Sharpe Ratio. MR, TR, and VOL are in percentage form.)

- **Low-Volatility Portfolio:** Markowitz yields a modest 1.2% mean return and 4.5% total return. Meanwhile, the benchmark achieves significantly higher figures (4.2% mean return, 20.9% total). The BL model tops both in mean return (4.4%) and total return (21.8%). Sharpe ratios follow a similar pattern, with BL (With Views) at 0.71, surpassing Markowitz's 0.16.

- **Mid-Volatility Portfolio:** Markowitz performs poorly again, returning only 0.6% per year and 0.7% total. By contrast, the the benchmark and the full BL model both exceed 5% annualized returns, with views boosting returns at 6.6% (MR) and a total return of 32.8%—far outperforming Markowitz's near-flat outcome.

- **High-Volatility Portfolio:** At higher volatility levels, Markowitz barely breaks even in total return (-0.3%), whereas the BL (with views) reaches an impressive

50.7% total return. The benchmark also sees a notable gain of 35.4%. Hence, the partial-correct views included quarter by quarter in the Black-Litterman model appear especially valuable at higher volatility levels.

A surprising result emerges when comparing the BL model with views to its benchmark. Even though the former model is given correct directional forecasts in 75% of the quarters, the market-implied approach without views consistently delivers strong returns as well. One might anticipate that incorporating views in the Black-Litterman model would decisively outperform because it actively tilts allocations in line with accurate (or mostly accurate) signals; however, the "benchmark" approach trails performance very closely. This outcome suggests that holding the market portfolio can be nearly as effective as sophisticated Bayesian adjustments.

Finally, the line plots in Figure 3.20 confirm this finding visually. For all three risk levels (low, mid, high), the green curve (benchmark) frequently runs close behind or, for certain quarters, even matches the red curve (BL with views), despite lacking explicit directional insights. The distinction is more marked relative to Markowitz (in blue), which underperforms for long stretches, particularly in mid- and high-risk portfolios.

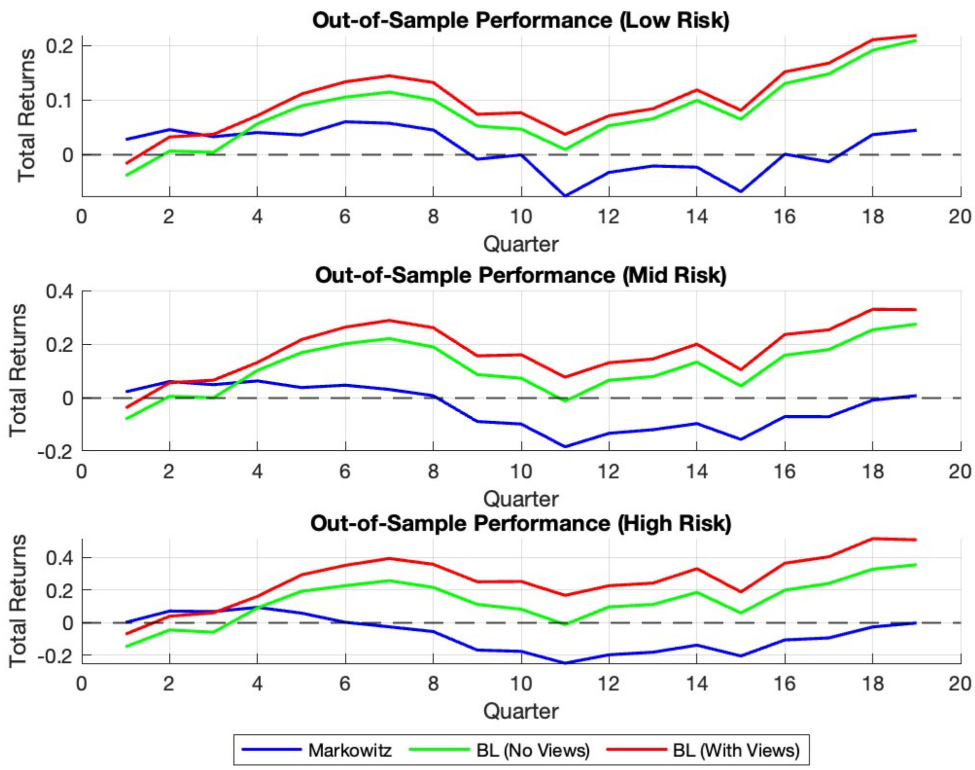

Figure 3.20: Out-of-sample return trajectories for low, mid, and high-volatility portfolios in the optimistic scenario.

3.3.5 Portfolio Allocations - Pessimistic Scenario

Figure 3.21 displays the average efficient-frontier allocation for Black–Litterman when 75% of the quarters are assigned incorrect views. Compared to the optimistic scenario, we see a reduced U.S. equity exposure in favor of European equity (which trailed the former in most of the out-of-sample window, consistent with the view setting scenario in this case). The other two models remain unchanged because they do not incorporate specific views. The resulting Black-Litterman frontier still reflects Bayesian "anchoring" around equilibrium weights which, as we will see, will provide great support despite the mostly incorrect bets in terms of rebalancing.

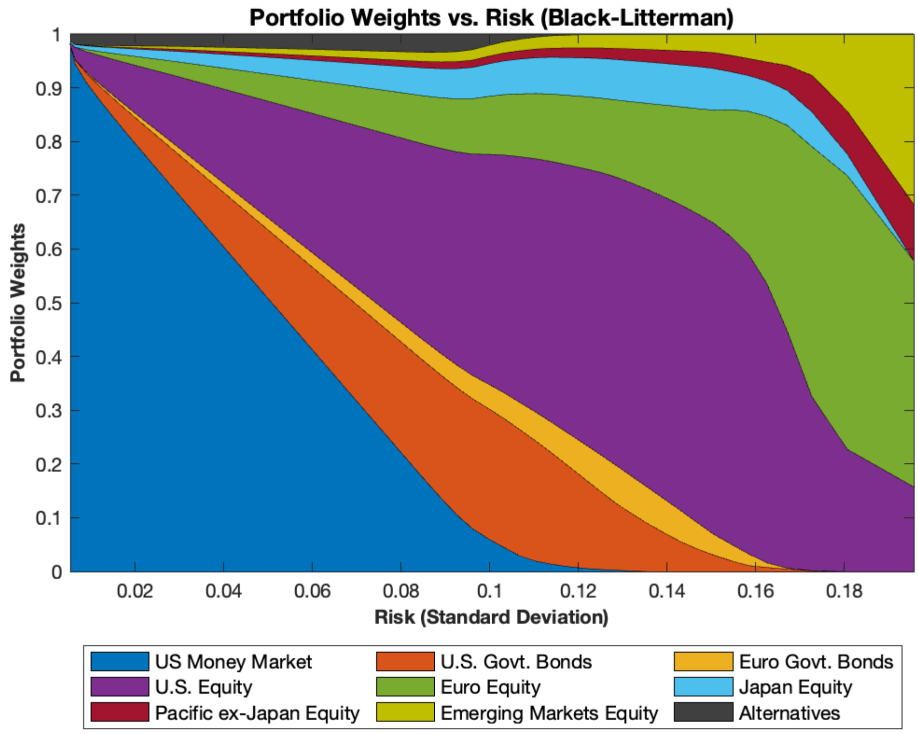

Figure 3.21: Portfolio weights under Black–Litterman with views in the pessimistic scenario

3.3.6 Metrics and Cumulative Returns - Pessimistic Scenario

In Table 3.17, the bottom row presents the new results for BL under pessimistic views. Despite receiving incorrect signals most of the time, this portfolio is able to narrowly trail the benchmark, thanks to the baseline equilibrium that mitigates wrong timed forecasts. Final outcomes remain close—though generally below—those of the benchmark, but well ahead of Markowitz.

Low-Volatility Portfolio				
Model	**MR**	**TR**	**VOL**	**SR**
Markowitz	1.2	4.5	7.1	0.16
BL (No Views)	4.2	20.9	6.4	0.66
BL (With Views)	3.5	17.0	6.6	0.53

Mid-Volatility Portfolio				
Model	**MR**	**TR**	**VOL**	**SR**
Markowitz	0.6	0.7	10.1	0.06
BL (No Views)	5.9	27.4	12.7	0.46
BL (With Views)	4.6	19.4	13.1	0.35

High-Volatility Portfolio				
Model	**MR**	**TR**	**VOL**	**SR**
Markowitz	0.7	-0.3	12.4	0.05
BL (No Views)	7.8	35.4	17.2	0.46
BL (With Views)	6.6	26.5	18.2	0.36

Table 3.17: Performance metrics for low, mid, and high volatility portfolios in the pessimistic scenario. (MR = Mean Return, TR = Total Return, VOL = Volatility, SR = Sharpe Ratio. MR, TR, and VOL are in percentage form.)

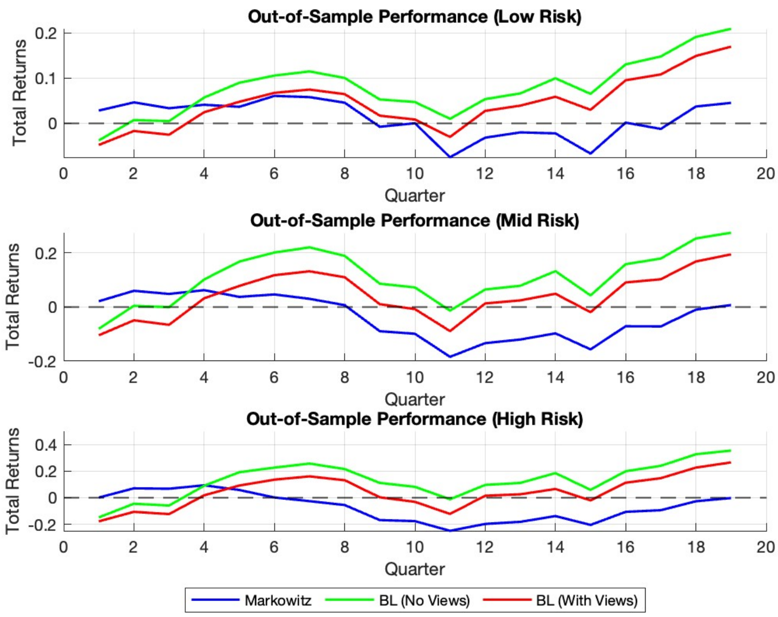

Figure 3.22: Out-of-sample return trajectories for low, mid, and high-volatility portfolios in the pessimistic scenario.

Figure 3.22 confirms these findings in quarterly total returns. The BL incorporating views (red line) rarely outperforms the benchmark (green) but stays near it for extended periods, even if it never decisively pulls ahead. By contrast, Markowitz (blue) lags substantially during the mid-quarters, with some recovery in later periods but still failing to catch either BL approach. Overall, the results reinforce that under mostly inaccurate views, an equilibrium "anchor" keeps BL from crashing too far—yet cannot reliably surpass the less sophisticated but balanced, no-views strategy.

Conclusion

This thesis set out to examine three distinct approaches to portfolio construction—Classical Markowitz (mean–variance), Black–Litterman without views (passive benchmark), and the Black–Litterman model incorporating investor views. These models were applied under two different frameworks. The first framework adopts a buy-and-hold strategy, where the portfolio is optimized just once at the outset and then maintained for five consecutive years. To assess how the Black–Litterman model behaves in such a static environment, a Monte Carlo analysis was conducted, comprising 1,000 simulations with varying investor outlooks. By including uncertainty proportional to each asset's volatility, the simulations captured a *realistic* range of optimistic and pessimistic perspectives, avoiding reliance on any single definitive forecast.

Under this no-rebalancing scenario, the results display two main observations. Over a five-year horizon, Black–Litterman (with views) outperforms Markowitz in almost every simulation, confirming that classical MV struggles out-of-sample due to its lack of diversification. Yet, the BL strategy incorporating investors view *underperforms* the market-implied benchmark in most cases. The latter's well-diversified equilibrium weighting is especially robust in periods of high volatility and trend reversals—both of which occurred in the out-of-sample window (and are generally frequent in real-world market dynamics). Hence, when only one static allocation is maintained for multiple years, a "balanced" or equilibrium-based approach can outperform a tactically correct but short-lived set of views that eventually become misaligned if market trends invert.

A more nuanced picture emerges with the introduction of quarterly rebalancing, where Black–Litterman forecasts can adjust to new market conditions at each interval. In an *optimistic* scenario—75% of the rebalancing windows have "correct" signals— the BL model does outperform the benchmark, but only marginally. This begs the question of whether the slight performance edge justifies the extra complexity and resources necessary required to implement the BL model.

Conversely, in a *pessimistic* view scenario, the "market prior" anchor in Black–Litterman helps avoid catastrophic misallocations, so even systematically inaccurate views do not heavily undermine performance. The Bayesian framework's partial reliance on equilibrium returns keeps the portfolio from deviating too far into losing bets. The results demonstrate that, in a multi-year horizon rebalancing framework, the posterior's pull toward a well- diversified baseline helps stabilize outcomes, even with systematically erroneous short-term signals. Notably, the BL model still outperforms Markowitz, which remains more vulnerable to poor in-sample estimates. Taken together, these findings offer several practical insights.

First, for long-term investors, the buy-and-hold results reinforce that well- diversified *market-implied* allocations can be highly competitive, even against theoretically informed views, as those views would need to be correct for often unrealistically prolonged portions of the horizon.

Second, for more tactical investors willing to reoptimise in shorter intervals, accurate short-term signals can deliver modest outperformance over a passive benchmark. However, the relative gain is not always dramatic; thus the extra cost/effort of frequent re-optimization might not be clearly warranted unless confidence in the forecasts is very high.

Finally, we observed that across all scenarios—static or dynamic, optimistic or pessimistic—Markowitz consistently trails behind the Black- Litterman model because of its sensitivity to estimation error and tendency toward extreme allocations.

In conclusion, this thesis demonstrates how incorporating investor views (optimistic or pessimistic) and employing different backtesting strategies (buy-and-hold vs. quarterly rebalancing) can affect portfolio performance, highlighting that surprisingly often, the simpler *equilibrium* benchmark can match or exceed more complex approaches.

Bibliography

[1] Best, M. and Grauer, R. (1991) 'On the Sensitivity of Mean-Variance-Efficient Portfolios to Changes in Asset Means: Some Analytical and Computational Results', *Review of Financial Studies*, 4, pp. 315–342.

[2] Bessler, W. and Wolff, D. (2015) 'Do Commodities Add Value in Multi-Asset Portfolios? An Out-of-Sample Analysis for Different Investment Strategies', *Journal of Banking and Finance*, 60, pp. 1–20.

[3] Bessler, W., Opfer, H. and Wolff, D. (2017) 'Multi-Asset Portfolio Optimization and Out-of-Sample Performance: An Evaluation of Black-Litterman, Mean Variance and Naïve Diversification Approaches', *The European Journal of Finance*, 23, pp. 1–30.

[4] Black, F. and Litterman, R. (1991) 'Asset Allocation: Combining Investor Views with Market Equilibrium', *Journal of Fixed Income*, 1, pp. 7–18.

[5] Black, F. and Litterman, R. (1991) 'Global Asset Allocation with Equities, Bonds and Currencies', Goldman Sachs and Co., New York.

[6] Black, F. and Litterman, R. (1992) 'Global Portfolio Optimization', *Financial Analysts Journal*, 48, pp. 28–43.

[7] Broadie, M. (1993) 'Computing Efficient Frontiers Using Estimated Parameters', *Annals of Operations Research*, 45, pp. 21–58.

[8] Chopra, V. K. and Ziemba, W. T. (1993) 'The Effect of Errors in Means, Variances, and Covariances on Optimal Portfolio Choice', *Journal of Portfolio Management*, 19, pp. 6–11.

[9] De Miguel, V., Garlappi, L. and Uppal, R. (2009) 'Optimal versus Naive Diversification: How Inefficient is the 1/N Portfolio Strategy?', *The Review of Financial Studies*, 22, pp. 1915–1953.

[10] Elton, E. J. and Gruber, M. J. (1997) *Modern Portfolio Theory and Investment Analysis* (5th ed.), John Wiley & Sons.

[11] Harris, R. D., Stoja, E. and Tan, L. (2017) 'The Dynamic Black–Litterman Approach to Asset Allocation', *European Journal of Operational Research*, 259, pp. 1085–1096.

[12] He, G. and Litterman, R. (2002) The Intuition Behind Black–Litterman Model Portfolios, Goldman Sachs Asset Management.

[13] Idzorek, T. (2004) A Step-By-Step Guide to the Black–Litterman Model: Incorporating User-Specified Confidence Levels.

[14] Jagannathan, R. and Tongshu, M. (2003) 'Risk reduction in large portfolios: why imposing the wrong constraints helps', *Journal of Finance*, 58, pp. 1651–1684.

[15] Jorion, P. (1986) 'Bayes–Stein Estimation for Portfolio Analysis', *Journal of Financial and Quantitative Analysis*, 21, pp. 279–292.

[16] Jorion, P. (1991) 'Bayesian and CAPM Estimators of the Means: Implications for Portfolio Selection', *Journal of Banking & Finance*, 15, pp. 717–727.

[17] Levy, H. and Levy, M. (2014) 'The Benefits of Differential Variance-Based Constraints in Portfolio Optimization', *European Journal of Operational Research*, 234, pp. 372–381.

[18] Ledoit, O. and Wolf, M. (2004) 'Honey, I Shrunk the Sample Covariance Matrix', *The Journal of Portfolio Management*, 30.

[19] Markowitz, H. (1952) 'Portfolio Selection', *Journal of Finance*, 7, pp. 77–91.

[20] Merton, R. C. (1980) 'On Estimating the Expected Return on the Market: An Exploratory Investigation', *Journal of Financial Economics*, 8, pp. 323–361.

[21] Michaud, R. O. (1989) 'The Markowitz Optimization Enigma: Is Optimized Optimal?', *Financial Analysts Journal*, 45, pp. 31–42.

[22] Meucci, A. (2010) 'The Black–Litterman Approach: Original Model and Extensions', in *Encyclopedia of Quantitative Finance*.

[23] Pástor, L. (2000) 'Portfolio Selection and Asset Pricing Models', *Journal of Finance*, 50, pp. 179–223.

[24] Satchell, S. and Scowcroft, A. (2000) 'A Demystification of the Black–Litterman Model: Managing Quantitative and Traditional Portfolio Construction', *Journal of Asset Management*, 1, pp. 138–150.

[25] Stein, C. (1955) 'Inadmissibility of the Usual Estimator for the Mean of a Multivariate Normal Distribution', in *Proceedings of the 3rd Berkeley Symposium on Probability and Statistics*, pp. 197–206.