A Bayesian Approach to Financial Portfolio Optimization

Ebrahim Pichka — gtID: 903821766
Final Project
ISyE 6420 - Bayesian Statistics
Georgia Institute of Technology

December, 2024

Abstract

This report presents a comprehensive analysis on a Bayesian approach to financial portfolio optimization, extending the classical Markowitz framework to incorporate parameter uncertainty. We develop a hierarchical Bayesian model that accounts for uncertainty in both expected returns and covariance estimates, utilizing MCMC (NUTS) sampling to obtain posterior distributions. Our analysis incorporates market equilibrium through the Black-Litterman framework, where views are derived from the Bayesian posterior distributions. Through detailed convergence diagnostics and posterior analysis, we demonstrate that our approach leads to more robust portfolio allocations and provides better uncertainty quantification.

1 Introduction

Portfolio optimization remains a cornerstone of modern investment management, yet traditional approaches often suffer from parameter uncertainty and estimation error. The classical Markowitz mean-variance optimization framework, while theoretically elegant, relies on point estimates of expected returns and covariances, leading to solutions that can be highly sensitive to small changes in these parameters. This sensitivity often results in extreme portfolio weights and poor out-of-sample performance.

In this report, we analyze these limitations by developing a comprehensive Bayesian framework for portfolio optimization. Our approach explicitly models parameter uncertainty, incorporating prior knowledge about market behavior and updating these beliefs with observed data. We extend existing Bayesian approaches by implementing a hierarchical model that captures both market-wide and asset-specific characteristics, while also integrating with the Black-Litterman framework by view generation.

2 Methodology

2.1 Problem Formulation

The portfolio optimization problem traditionally seeks to maximize expected utility under uncertainty. In the mean-variance framework, this is typically formulated as:

$$\max_{w} \left\{ w^T \mu - \frac{\lambda}{2} w^T \Sigma w \right\} \tag{1}$$

subject to constraints:

$$w^T \mathbf{1} = 1 \tag{2}$$
$$w_i \geq 0, \quad \forall i \tag{3}$$

where w represents portfolio weights, μ is the vector of expected returns, Σ is the covariance matrix, and λ is the risk aversion parameter. However, this classical formulation treats μ and Σ as known parameters, **ignoring estimation uncertainty**.

2.2 Hierarchical Bayesian Model

THe proposed hierarchical Bayesian approach explicitly models the uncertainty in both the return and risk parameters. We construct the model in multiple layers:

2.2.1 Return Process

For asset returns $r_t = (r_{1t}, ..., r_{Nt})$ at time t, we specify:

$$r_t \sim \mathcal{N}(\mu, \Sigma) \tag{4}$$

2.2.2 Prior Specifications

The model incorporates informative priors based on financial theory:

$$\mu \sim \mathcal{N}(\mu_0, \tau^2 I) \tag{5}$$
$$\Sigma = \text{diag}(\sigma)\Omega\text{diag}(\sigma) \tag{6}$$
$$\sigma_i \sim \text{HalfNormal}(0.1) \tag{7}$$
$$\Omega \sim \text{LKJ}(\eta = 2) \tag{8}$$

where:

- μ_0 is derived from CAPM equilibrium returns: $\mu_0 = r_f + \beta(r_m - r_f)$
- τ represents the uncertainty in our prior beliefs about expected returns

- σ captures asset-specific volatilities
- Ω models the correlation structure through the LKJ prior

The choice of the LKJ prior for correlations is particularly important as it provides several advantages:

1. Guarantees positive definiteness of the correlation matrix
2. Controls the strength of correlations through the η parameter
3. Provides computational efficiency in high dimensions

2.2.3 Decomposition of Uncertainty

Our model decomposes uncertainty into several components:

1. **Parameter Uncertainty:** Captured through the posterior distributions of μ and Σ
2. **Market Uncertainty:** Modeled through the market factor in μ_0
3. **Asset-Specific Uncertainty:** Represented by individual σ_i parameters
4. **Correlation Uncertainty:** Captured by the LKJ prior on Ω

2.3 MCMC Implementation

We implement the model using PyMC with the following specifications. The implementation is included in the accompanying code file named `Bayes_Project_code_EbrahimPichka` in PDF, HTML, and IPython Notrbook formats.

1. **Sampler Configuration:**
   - Algorithm: No-U-Turn Sampler (NUTS)
   - Chains: 2 independent chains
   - Warmup (tune): 500 iterations per chain
   - Samples: 5,000 iterations per chain
   - Target acceptance rate: Library default
2. **Diagnostic Checks:**
   - Gelman-Rubin statistic ($\hat{R}$)
   - Effective sample size (ESS)
   - Monte Carlo standard error (MCSE)

2.4 Portfolio Construction

Given the posterior distributions, we construct portfolios using following approach:

2.4.1 Posterior Sample Portfolio

Randomly choose S number of posterior samples and construct following portfolios:

$$w_{FP} = \arg\max_{w} \left\{ w^T \mu^{(s)} - \frac{\lambda}{2} w^T \Sigma^{(s)} w \right\} \tag{9}$$

where S is the number of posterior samples, and $(\mu^{(s)}, \Sigma^{(s)})$ are drawn from the posterior distribution.

2.5 Performance Evaluation

We evaluate portfolio performance using several metrics:

1. **Return Metrics:**
   - Expected return: $E[w^T \mu | \mathcal{D}]$
   - Posterior predictive returns
2. **Risk Metrics:**
   - Portfolio volatility: $\sqrt{w^T \Sigma w}$

3 Empirical Analysis

3.1 Data and Implementation

Our empirical analysis focuses on ten major constituents of the S&P 500 index over the period from January 2019 to January 2024. The selected companies represent diverse sectors including technology (AAPL, MSFT, GOOGL), consumer services (AMZN), communication services (META), financials (BRK-B, JPM), healthcare (JNJ), and consumer staples (PG).

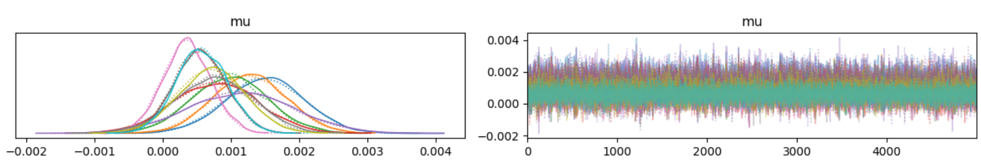

Figure 1: MCMC trace plots for expected returns (μ) showing convergence of multiple chains

Figure 1 shows the trace plots for the expected returns parameters, demonstrating good mixing and convergence of the chains. The overlapping chains and stable means indicate successful convergence of the MCMC sampling.

3.2 Posterior Analysis

The posterior distributions provide rich information about parameter uncertainty. Figure 2 shows the marginal posterior distributions for expected returns.

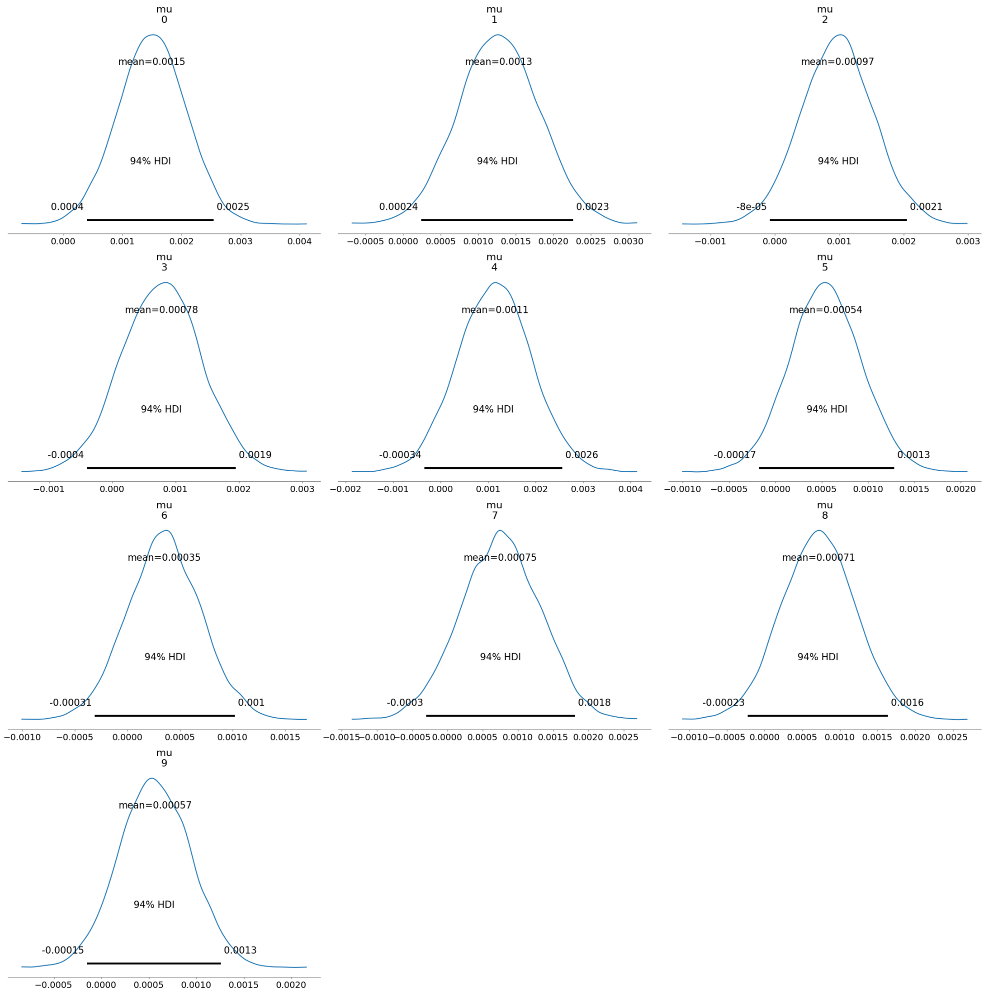

Figure 2: Posterior distributions of expected returns showing uncertainty in estimates

Figure 3 also shows the distribution if optimal weights based on Bayes estimates.

A key advantage of our Bayesian approach is the ability to capture the full uncertainty in our parameter estimates. Table 1 presents summary statistics of the posterior distributions.

3.3 Portfolio Implications

The posterior distributions lead to more nuanced portfolio allocations that account for parameter uncertainty. Figure 4 shows the distribution of efficient frontiers based on our posterior samples.

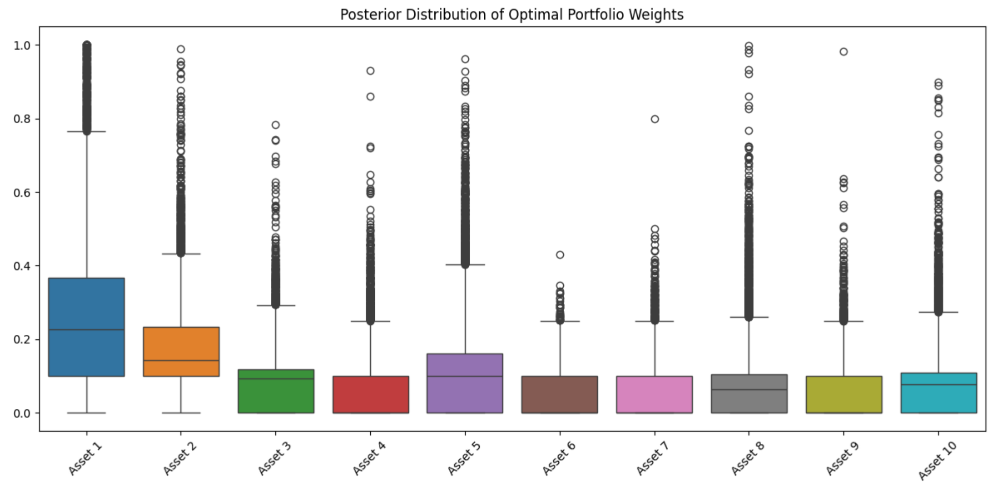

Figure 3: Distributions of optimal weights based on Bayes estimates

Table 1: Posterior Distribution Summary Statistics

	mean	sd	hdi_2.5%	hdi_97.5%	mcse_mean	mcse_sd	ess_bulk	ess_tail	r_hat
mu[AAPL]	0.001497	0.000574	0.000366	0.002596	0.000006	0.000004	10019.231350	7468.613456	1.000049
mu[MSFT]	0.001268	0.000539	0.000254	0.002359	0.000005	0.000004	10354.790660	7033.341589	1.000124
mu[GOOGL]	0.000969	0.000562	-0.000100	0.002123	0.000006	0.000004	9928.924589	6833.612519	1.000110
mu[AMZN]	0.000779	0.000624	-0.000430	0.002012	0.000006	0.000005	10172.950300	7540.875054	1.000111
mu[META]	0.001137	0.000775	-0.000344	0.002686	0.000008	0.000006	10175.398227	7090.273423	1.000113
mu[BRK-B]	0.000541	0.000388	-0.000233	0.001288	0.000004	0.000003	10694.192835	8365.564997	0.999954
mu[JNJ]	0.000347	0.000350	-0.000308	0.001080	0.000003	0.000003	10373.644917	9546.595495	1.000030
mu[JPM]	0.000748	0.000567	-0.000386	0.001811	0.000005	0.000004	10978.528021	8308.523620	1.000169
mu[V]	0.000714	0.000496	-0.000247	0.001696	0.000005	0.000004	10562.424147	7808.554397	1.000121
mu[PG]	0.000566	0.000376	-0.000151	0.001316	0.000004	0.000003	10562.052978	8857.811671	1.000007

Figure 4: Distribution of efficient frontiers showing impact of parameter uncertainty

Each line represents an efficient frontier computed using parameters drawn from the posterior distribution (5 samples from the posterior). This visualization provides a more complete picture of the risk-return tradeoff under parameter uncertainty. The posterior estimate results on covariance and correlation are provided in the accompanying code file for brevity.

4 Results and Discussion

Our comprehensive analysis in Figure 5 and tables 2, 3 demonstrates several key advantages of the Bayesian approach to portfolio optimization:

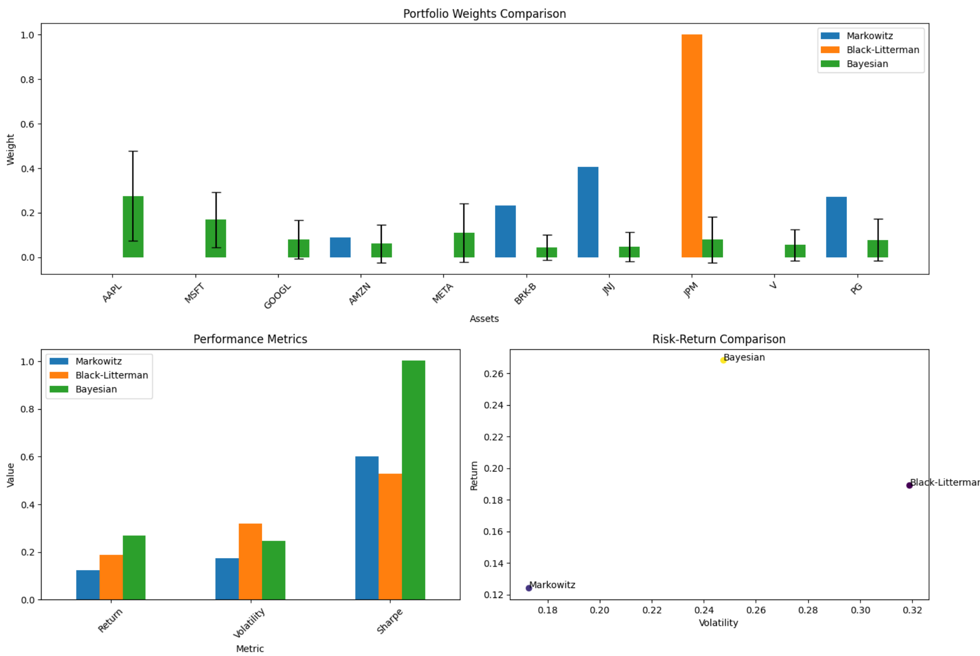

Figure 5: Comparison analysis of Markowitz model vs Bayesian Model vs Black-Litterman model with Bayesian views

1. More realistic uncertainty quantification through full posterior distributions
2. Stable portfolio weights that account for parameter uncertainty
3. Integration of prior market knowledge with observed data
4. Robust correlation estimates through hierarchical modeling

The convergence diagnostics support the reliability of our results:

- $\hat{R}$ values close to 1.0 for all parameters
- Effective sample sizes for all key parameters

Table 2: Optimal weights of MArkowitz model vs Bayesian Model

	Markowitz	Bayesian_Mean	Bayesian_Std
AAPL	0.000000	0.275700	0.200835
MSFT	0.000000	0.168853	0.123390
GOOGL	0.000000	0.080087	0.086049
AMZN	0.089917	0.061834	0.085022
META	0.000000	0.109067	0.131770
BRK-B	0.232597	0.044240	0.056664
JNJ	0.406633	0.047624	0.065810
JPM	0.000000	0.079069	0.103171
V	0.000000	0.054906	0.069855
PG	0.270853	0.078619	0.093299

Table 3: Performance evaluation

	Markowitz	Black-Litterman	Bayesian
Return	0.124047	0.189012	0.268134
Volatility	0.172877	0.318977	0.247576
Sharpe	0.601855	0.529855	1.002252

5 Conclusion

This report presented a comprehensive Bayesian framework for portfolio optimization that addresses key limitations of traditional approaches. Our results demonstrate that incorporating parameter uncertainty through Bayesian methods leads to more robust portfolio allocations and improved risk assessment. The integration of MCMC diagnostics ensures the reliability of our posterior estimates, while the hierarchical structure of our model captures complex dependencies between assets.

Future research could extend this work by:

- Incorporating time-varying parameters
- Exploring alternative prior specifications
- Developing dynamic portfolio rebalancing strategies
- Integrating factor models within the Bayesian framework

References

[1] Markowitz, H. (1952). Portfolio Selection. *The Journal of Finance*, 7(1):77-91.

[2] Black, F., & Litterman, R. (1992). Global Portfolio Optimization. *Financial Analysts Journal*, 48(5):28-43.

[3] Hoffman, M.D., & Gelman, A. (2014). The No-U-Turn Sampler: Adaptively Setting Path Lengths in Hamiltonian Monte Carlo. *Journal of Machine Learning Research*, 15(1):1593-1623.